

THE SCOVILLE SPLICE LIBRARY · VOLUME I

The Iyengar Flavor Index

Eight axes. Every flavor. *All at once.*

The complete standard for measuring — and designing — taste. A receptor-level account of the eight axes, the mathematics that scores them, and the organisms we have tuned to speak each one.

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

Contents

01 Opening essay · Why a flavor needs a number

02 Chapter 1 · The Lock and the Key — taste is a chemistry we can finally read.

03 Chapter 2 · The Half-Maximal Point — why EC_{50} is the only honest unit.

04 Chapter 3 · One Thousand Times the Reference — the formula, derived.

05 Chapter 4 · The Radar and the Bar — how to read a profile.

06 Chapter 5 · Heat — the axis everyone knows.

07 Chapter 6 · Sweetness — sugar's signal without sugar's cost.

08 Chapter 7 · Umami — the taste of being cared for.

09 Chapter 8 · Zero-Bitter — subtraction as a flavor.

10 Chapter 9 · Kokumi — the axis you only notice when it's gone.

11 Chapter 10 · Cool-Burn — temperature that isn't there.

12 Chapter 11 · Pungency — the sharp, fast channel.

13 Chapter 12 · Astringency — the taste that is a texture.

14 Chapter 13 · Full-Spectrum — all eight at once.

15 Chapter 14 · The Combinatorics of Flavor — axes interfere.

16 Chapter 15 · From Score to Sequence — designing the molecule to hit the number.

17 Chapter 16 · The Phenylpropanoid Spine — one pathway under almost everything.

18 Chapter 17 · Forty-Two Hosts — the organism library.

19 Chapter 18 · The Sixty-Four — what could be built.

20 Chapter 19 · Precision as Respect — a number means nothing without its error bar.

21 Chapter 20 · Reading the World — score anything.

Opening essay · Why a flavor needs a number

Every meal our family has ever loved began as a chemical conversation we could not read — a tomato warm from the vine and split at the table, the particular hush of a kitchen when something is almost ready, all of it carried by molecules settling into receptors on the surface of the tongue, a lock and key repeated a billion times a second. We could feel the result exactly. We could describe it only in poetry.

For ten thousand years that was the whole art: grow the plant, taste the fruit, keep the seed of the one that moved you. A lineage of patience, of grandmothers and their one perfect vine, and not once in all of it could anyone say *how much.* How much sweetness. How much warmth. How much of the round, savory depth that makes a broth taste like being cared for. Flavor is the most intimate technology the body owns — the sense wired straight to memory and appetite and the will to be nourished — and it is the only one we never learned to measure. We measured heat in degrees and sound in decibels and light in lumens, built an entire civilization on the premise that what can be quantified can be improved, and left the thing we put in our mouths three times a day to adjectives.

This essay is the instrument. Before a single plant in this standard, before the receptors and the enzymes and the organisms we have tuned to speak each axis, there has to be a unit — a way of saying *how much* that a chili, a tomato, and a cup of tea can all be held to on the same page. What follows is how that unit is built: how a molecule becomes a feeling, why the honest place to measure a taste is the exact middle of its response and not its edge, how the formula falls out of that choice, and how to read a finished profile at a glance. It is arithmetic in the service of reverence. We did the math because we meant to feed people well, and you cannot compose what you cannot count.

The ten-thousand-year silence — *the sense we never learned to measure.*

Consider what we were willing to quantify and what we were not. We pinned down the boiling point of water and the wavelength of the color red; we gave the tremor of the earth a magnitude and the roar of an engine a number of decibels; we learned to say that one star is a precise multiple brighter than another. The whole machinery of improvement runs on measurement — you refine toward a number, and the number tells you whether you are getting closer. Yet taste, the sense that decides what we take into the body and therefore what the body is built from, we left to the tasting note and the family story. A cook could tell you a dish needed "more depth" and be completely right and completely unable to say *how much more,* or depth *of what.*



The founders in the research greenhouse at dawn, where the question that became this company was first asked at a family table.

The silence was not for lack of caring. It was for lack of a reference. A measurement is never an absolute; it is always a ratio against something fixed — a meter was once a bar of metal in a vault, a degree is a fraction of the span between ice and steam. Taste had no bar in the vault. It had opinions. The work of this standard was, before anything else, the work of choosing the fixed things — eight of them, one for each pathway through which a plant actually speaks to the body — and then measuring everything as a ratio against them. The silence ends not with a poem but with a denominator.

Scoville's brave mistake — *he measured heat by how much it took to make it stop.*

There was one brave early attempt, and it is worth honoring because it shows both the hunger for a number and the trap. In 1912 a pharmacist named Wilbur Scoville wanted to rank the pungency of peppers, so he dissolved an extract in sugar water and diluted it, step by step, until a panel of tasters could no longer feel the burn; the dilution at which the heat vanished became the Scoville Heat Unit. It was ingenious, and it was doomed, because the instrument was a person — a panel's fatigue, its tolerance, its mood, baked into every reading.

But the deeper flaw was not the human panel; it was *where on the response* Scoville chose to read. He measured the dilution at which the heat *disappeared* — the vanishing point, the threshold. And the threshold is the single worst place to take a measurement, because it sits in the flat, noisy foot of the response, where a tiny change in concentration produces almost no change in sensation and therefore enormous scatter in the reading. Ask ten people where a sound becomes *inaudible* and you will get ten answers spread across a wide band; ask them which of two clearly audible tones is louder and they will agree. Scoville measured the inaudible edge. His scale survived because the hunger behind it was real, and it survives today in a far better form — an HPLC measurement of capsaicinoid mass, where pure capsaicin is fixed at roughly 16,000,000 Scoville units and every lesser heat is a fraction of it. We still report SHU alongside our own number, out of respect and for the chili community that speaks in it. But its original sin is instructive: it measured heat by *how much it took to make it stop,* which is a confession that no one had a unit for how loud it was to begin with.

The Index corrects the choice of where to read. Instead of the vanishing edge, it anchors to the exact middle of the response — the point where the receptor is driven to half of its full output. That

single decision is the difference between a number that scatters and a number you can design against, and the rest of this essay is, in a sense, the defense of it.

How a molecule becomes a feeling — *the tongue as a bank of listeners.*

To measure a taste honestly you first have to know what is physically happening, step by step, when a food "has" one. Nothing on the tongue perceives a molecule. What the tongue has is a bank of proteins, each tuned to a narrow family of compounds, and each wired so that when its compound arrives, a nerve fires. The feeling is downstream of the firing. Flavor, measured honestly, is the loudness of that firing.

Take heat, the axis everyone knows. Capsaicin does not warm the tissue; it binds **TRPV1**, a four-part cation channel standing in the membrane of the trigeminal nerve endings. In its resting state the channel is shut. Capsaicin slips into a pocket between the channel's third and fourth helices and, by binding, lowers the energy the channel needs to open — the same channel that would otherwise open only above about 43 °C now swings open at the temperature of your mouth. The gate dilates; sodium and calcium ions pour through; the nerve ending depolarizes; an action potential travels to the brain, which has exactly one word for a TRPV1 signal, and the word is *fire.* The sequence is strict and mechanical: **bind → gate opens → ions flow → nerve depolarizes → heat is reported.** No step is a metaphor. A pepper is not hot; it is a key that fits a lock that, when turned, makes the body shout heat.



The assay bench where loudness becomes a number — receptor activation read across a dilution series, the exact middle of the curve recorded as the unit.

The other seven axes run the same logic through three different kinds of hardware, and the division matters because it is what makes eight axes *countable* rather than infinite. Three axes are **ion channels** like TRPV1 — fast, direct, a molecule in the pocket and a gate swings: heat through TRPV1, the cool of menthol through TRPM8, the sharp pungency of mustard and wasabi through TRPA1. Four axes are **G-protein-coupled receptors** — slower, and amplified, because binding launches an internal cascade that multiplies one captured molecule into a flood of second messengers: sweetness through the TAS1R2·TAS1R3 pair, the deep savory pull of umami

through mGluR1, bitterness through TAS2R38, and the elusive richness of kokumi through the calcium-sensing receptor CaSR. And one axis is neither — astringency is not a receptor event at all. Tannins physically cross-link and precipitate the proline-rich proteins that lubricate the mouth; the brain reads the resulting loss of slip through plain mechanical touch. Eight axes, three mechanisms, one countable set:

AXIS	NAME	RECEPTOR	TRANSDUCTION	WHAT IT READS
IFI-H	Heat	TRPV1	ion channel	capsaicin heat and spice
IFI-S	Sweetness	TAS1R2 / TAS1R3	GPCR cascade	sugar-free sweetness
IFI-U	Umami	mGluR1 / TAS1R1·TAS1R3	GPCR cascade	savory, brothy depth
IFI-B	Zero-Bitter	TAS2R38	GPCR cascade	bitterness attenuation (20–65%)
IFI-K	Kokumi	CaSR	GPCR cascade	richness, continuity, mouthfeel
IFI-C	Cool-Burn	TRPM8 / TRPV1	ion channel	cooling / warming
IFI-P	Pungency	TRPA1	ion channel	allicin, wasabi, cinnamon sharpness
IFI-A	Astringency	salivary PRPs	protein precipitation	tannin / polyphenol mouthfeel

This is the whole reason the Index can exist where no index of aroma can. Each of these eight has a single responsible mechanism and a single thing you can measure — the fraction of receptors driven open, or the fraction of lubricating protein thrown out of solution. Measure that fraction as the stimulus climbs, and you have a curve. The curve is the raw material of the number.

The dose that speaks at half — *why the curve, and not the threshold.*

Plot receptor response against the concentration of the stimulus and you do not get a straight line. You get an S — almost nothing at low doses, then a steep climb, then a plateau as every available receptor is occupied and the signal can grow no further. That S-curve is described, across essentially all of pharmacology, by the Hill equation:

$$R / R_{max} = [A]^n \,/\, (EC_{50}^n + [A]^n)$$

where [A] is the stimulus concentration, R/R_max is the fraction of full response, **EC_{50}** is the concentration that produces exactly half of full response, and **n** is the Hill coefficient, a measure of how steep and switch-like the climb is. Set [A] equal to EC_{50} and the equation collapses to one-half, by construction. EC_{50} is the half-maximal point — the amount of a stimulus at which its receptor sits exactly halfway between silent and saturated. A smaller EC_{50} means the receptor answers to less; the molecule is more potent; the signal is louder for the same dose.

Why read there, and not at the threshold Scoville chose, or at saturation? Because the half-maximal point is, provably, the single most reproducible place on the whole curve. Plotted against the logarithm of concentration — the way the nervous system itself encodes intensity — the Hill curve is point-symmetric about EC_{50}, and its slope reaches its one maximum exactly there. The steepest point is the most information-dense point: a small, honest change in the stimulus produces the largest, most legible change in response, so measurement error is smallest relative to signal. Move toward the threshold or toward saturation and the curve flattens; the same measurement error now swamps a vanishing change in signal, and the reading scatters. Scoville read the flattest place on the curve. We read the steepest. That is the entire methodological difference, and it is why our number holds still.

The Hill coefficient does real work too, because it sets how wide a dose range separates a faint taste from a full one. For capsaicin at TRPV1, cooperativity runs near $n \approx 2$. Run the Hill equation from one-tenth to nine-tenths of full response at $n = 2$ and the concentration need only span from $EC_{50} \div 3$ to $EC_{50} \times 3$ — a ninefold window from "barely there" to "nearly saturated." A less cooperative receptor, $n \approx 1$, would spread that same journey over an eighty-onefold window: a gentle, forgiving climb. Steepness is why heat feels like a switch and sweetness like a dial, and it is encoded, honestly, in a single exponent.

With the molecule's own potency as the anchor, the capsaicin curve becomes concrete. TRPV1 reaches half of full output at roughly 0.7 µM capsaicin — and because the molar mass of capsaicin is 305.4 g/mol, that EC_{50} is $0.7 \times 305.4 \approx 214$ µg/L, about **0.214 milligrams per liter.** Two-tenths of a milligram in a liter of water drives half the channels in a nerve ending open. The ninefold activation window for $n = 2$ then runs from about 0.23 µM to about 2.1 µM. These are the real dimensions of the feeling we have spent ten thousand years calling, vaguely, *hot.*

Building the number — *the formula, derived.*

An EC_{50} in micromolar is a true measurement, but it is not yet a *score* — it is bulky, it carries units, and a smaller number means a louder taste, which is backwards from how a person wants to read a scale. So we do to it exactly what every mature measurement does: we fix a reference and take a ratio. Choose a reference stimulus for the axis — for heat, the native wild-type *Capsicum* that evolution handed us — and measure its half-maximal concentration, EC_{50}(reference). Then measure the candidate's, EC_{50}(candidate). Divide the reference by the candidate so that a *more* potent candidate — a smaller EC_{50} — yields a *larger* ratio. Multiply by a thousand to clear the decimals and seat the reference itself at a round, memorable value:

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

That is the whole instrument. The reference, divided by itself, gives one, times a thousand: **the reference scores exactly 1,000 on its own axis, by construction** — it is the origin, the A-440 of taste, the fixed thing in the vault. A candidate twice as loud as the reference scores 2,000; one half as loud scores 500. The ratio is dimensionless, so the micromolars cancel and what remains is pure comparison: a chili, a tomato, and a cup of tea, each a number on the same page, because each is a ratio against its own fixed anchor. It is the same move that makes pH a number instead of an adjective and a magnitude out of an earthquake: pick a reference, measure everything as a ratio against it, and the ineffable becomes the comparable. For the radar scorecards in this library we cap each axis at a normalization maximum of **5,000**, so that the full designed range — from the faint to the engineered extreme — fits one readable dial.

There is one axis where the formula runs, deliberately, in reverse, and naming it is the most honest thing we can do for the reader. Bitterness is the taste we usually want *less* of, so IFI-B, "Zero-Bitter," scores suppression rather than activation: EC_{50}(candidate) is the bitter-ligand concentration needed to reach half-maximal TASR38 response *in the presence of the food's suppressor,* and EC_{50}(reference) is that concentration without it. A food that blunts bitterness shifts the curve rightward — you need more ligand to reach the same half-max — so the ratio climbs above 1,000 precisely when the bitterness has been quieted. Across our portfolio our model predicts 20–65% bitterness attenuation. One axis measures how loud a taste shouts; one measures how well a food has learned to whisper. The formula is the same shape; only the thing in the numerator changes. Astringency keeps its own seam too: its "EC_{50}" is the tannin concentration at which half the lubricating protein precipitates, a measurement of lost slip rather than opened channels. A standard is defined as much by its seams as by its center, and we would rather show them than paper them over.

Reading the heat ladder — *the formula in the greenhouse.*

The formula earns its keep in the growing room. Our ten proof-of-concept superhots — bred not by selecting a phenotype and praying the genotype followed, but by writing the genotype on purpose — climb the heat axis in a deliberate ladder, each rung a cultivar, each cultivar a SKU, a predicted Scoville value, and a predicted IFI-H score:

CULTIVAR	SKU	SHU	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**

Read the ladder closely and it teaches the single most important thing about the Index: **IFI-H is not Scoville.** The two correlate — climb one and you tend to climb the other — but they are not interconvertible, and the ladder shows exactly why. Cinder, at 1.05M SHU, scores IFI-H 2,290; Kindle, with *more* capsaicinoid mass at 1.25M SHU, scores *lower,* at 2,000. More milligrams, quieter signal. The reason is mechanical, not mysterious: SHU counts total capsaicinoid *mass* by HPLC, but the receptor does not weigh molecules — it counts activations, and capsaicin and its cousin dihydrocapsaicin, set by the length and saturation of one acyl chain the plant builds, engage TRPV1 with different potencies and different time-courses. The *shape* of the capsaicinoid blend, not just its weight, sets the loudness. SHU asks how much is in the fruit. IFI-H asks how loudly it speaks. Those are different questions, and only the second one is the taste.

Now turn the formula around and let it tell us what the receptor is seeing. Because IFI-H = 1,000 × EC_{50}(reference) ÷ EC_{50}(candidate), a known IFI-H gives the candidate's half-maximal concentration directly, as a fraction of the reference: EC_{50}(candidate) = EC_{50}(reference) × 1,000 ÷ IFI-H. And because Scoville units convert to capsaicinoid mass at the standard 16 SHU per part-per-million, each rung also carries a mass fraction. Both columns below are *derived* — computed from the canonical SHU and IFI-H by hand, nothing added:

CULTIVAR	SHU	IFI-H	CAPSAICINOID (% DRY MASS)	EC_{50} (× REFERENCE)
native reference	—	1,000	*anchor*	1.000
Ember	866K	1,570	5.4%	0.637
Blaze	1.72M	2,630	10.8%	0.380
Firestorm	2.47M	3,490	15.4%	0.287
Solara	2.72M	3,980	17.0%	0.251
Phoenix	**2.95M**	**4,200**	**18.4%**	**0.238**

The derivation is plain arithmetic a reader can check. Phoenix's model-predicted ~2.95 million SHU, divided by 16, gives 183,622 ppm — **our model places 18.4% of the dried fruit's mass as capsaicinoid,** a load no wild pepper approaches and the reason we had to write the plant gene by gene rather than find it. And its IFI-H of 4,200 means EC_{50}(candidate) = 1,000 ÷ 4,200 = 0.238 of the

reference: **Phoenix's extract is designed to drive TRPV1 to half-maximum at under a quarter of the concentration the native pepper needs — 4.2 times louder in our model.** Four-fold, not forty — because loudness, mercifully, compresses. That is the number doing its job: not a vague "the hottest we grow," but *4.2 times the half-activation potency of the wild type,* a figure you can design toward.



Phoenix (SS-0096) in the controlled chamber — the first cultivar engineered to register on all eight axes at once, and the proof that a flavor you can measure is a flavor you can compose.

How loud, how sure — *error, CV, and the honesty of a bar.*

A number without an error bar is an opinion wearing a lab coat, so every score in this library carries one. The measure we use is the coefficient of variation — the standard deviation as a percentage of the mean — because it says, in one figure, how much a reading wanders from batch to batch. Our engineered cultivars are designed to hold a CV of **≤7%.** The wild type, left to the weather and the open field, wanders at **42%.** That gap is the difference between a seasoning you can formulate around and a lottery.

The ≤7% design target is not luck; it falls out of the formula. The IFI is a ratio of two measurements — reference over candidate — and when you divide two independent measurements, their relative errors add in quadrature: $CV(IFI) = \sqrt{(CV(EC_{50},ref)^2 + CV(EC_{50},cand)^2}$). Two assays each held to a clean 5% give $\sqrt{(25 + 25)} = \sqrt{50} \approx$ **7.07%.** The whole-number target and the measurement discipline are the same fact seen twice. Put it on a real score and it becomes a confidence interval you can state out loud: Phoenix's predicted IFI-H of 4,200 at CV 7% has a standard deviation near 294, so a 95% interval runs roughly **3,624 to 4,776.** The wild type, at 42%, would smear a nominal reading of 1,000 across an interval from below 180 to above 1,800 — the same uselessness Scoville's panel suffered, now quantified.

That error bar also sets an honest limit on what the Index can resolve, and we would rather say it plainly than oversell the instrument. Two scores closer together than about twice the CV cannot be reliably told apart. Solara at 3,980 and Phoenix at 4,200 differ by 5.4% — *inside* the noise. On the heat axis alone, they are the same pepper. What truly separates them is not loudness but *breadth:* Solara speaks on six axes, Phoenix on all eight. The lesson sits at the heart of the whole standard — past a certain point you stop making a flavor louder and start making it *fuller,* and only an eight-axis instrument can see the difference. Two more honest limits

ride alongside: the Index reads *acute* activation, not the slow fade of receptor desensitization that makes the third bite of a chili milder than the first; and it reads a standardized extract, where real food shifts with fat, pH, and temperature, all of which change how much capsaicin ever reaches the receptor. We measure the stimulus cleanly and tell you so, rather than pretend the kitchen is a cuvette.

How to read a profile — *eight numbers at a glance.*

Set eight axes side by side and a food stops being an adjective and becomes a chord. The reference origin is the plainest possible profile — eight anchors, each defined to read exactly 1,000 on its own axis, the still point every real food is a departure from:

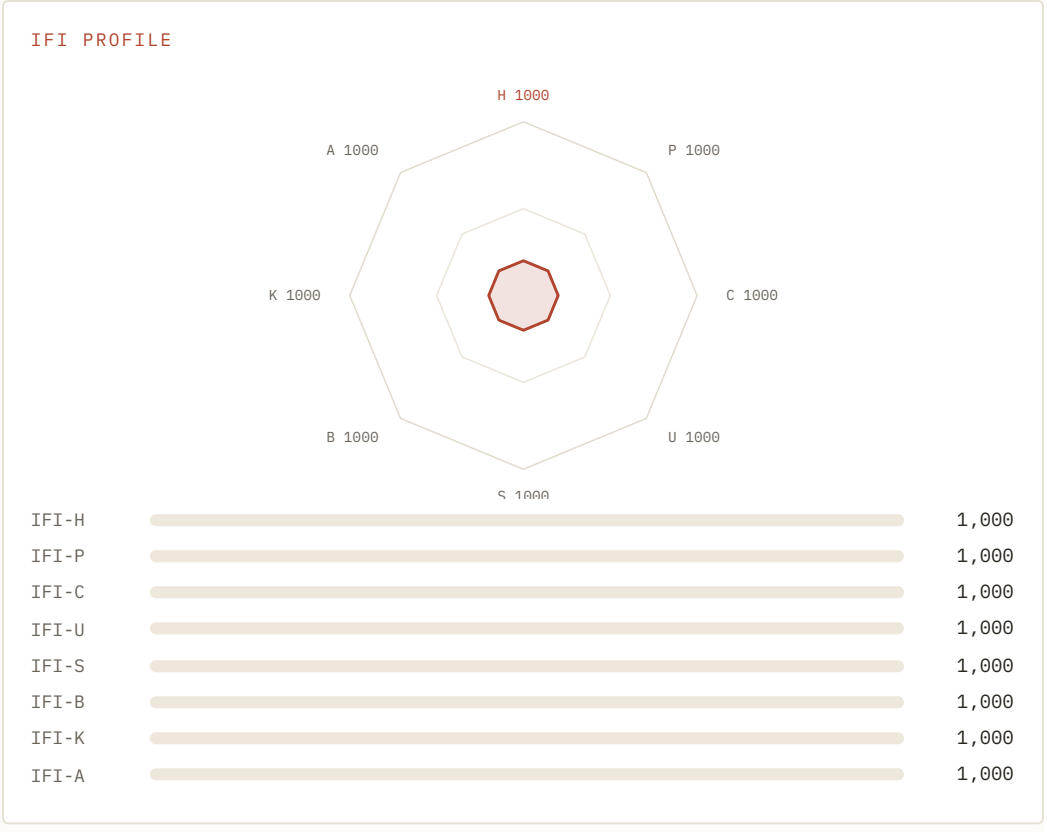

It is not a food you could eat — each 1,000 belongs to a different anchor substance, the wild pepper for heat, a reference sweetener for sweetness, and so on — but it is the zero from which every profile is read. Now set the same eight axes to Phoenix, the first cultivar engineered to register on all of them at once:

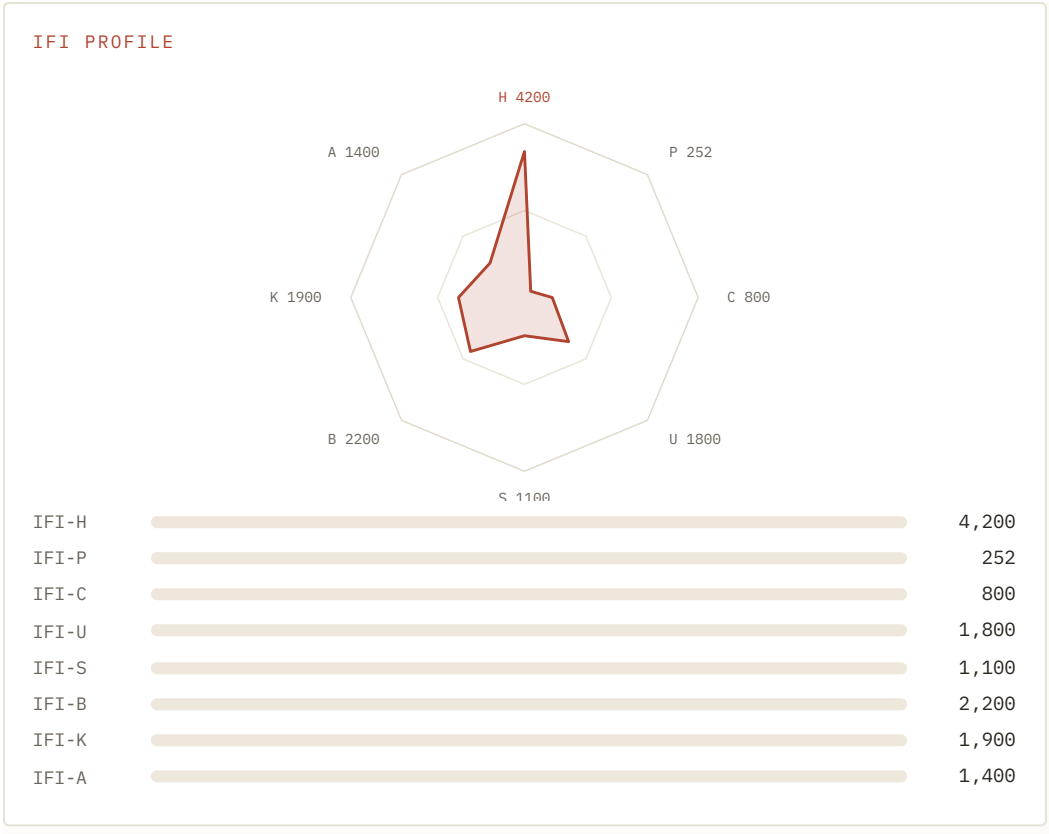

Read it like a chord. Heat towers at 4,200 — 4.2 times the wild type — but it does not stand alone. Beneath it runs a predicted 2,200 of bitterness-quieting, a 1,900 of kokumi richness that gives the burn a body to sit on, 1,800 of umami depth, and harmonics of sweetness and astringency lifting above the reference — pungency and cool modulated to a quieter register, integrated into the whole rather than amplified on their own. This is what *full-spectrum* means, and why it is the pinnacle of the design: not a single axis screamed to its ceiling, but eight voices balanced so the heat arrives inside a flavor rather than instead of one. A reader who learns to read these eight numbers here can read every profile in this library and every book built on it — a tomato's quiet two-note chord of umami and sweetness, a cup of tea's astringent-bitter restraint, a designed strawberry's impossible warmth — because they are all written in the same notation, against the same eight anchors, on the same dial that tops out at 5,000.

Where the number stops — *aroma, adaptation, and the edges we kept.*

A standard is defined as much by its edges as by its center, and ours has a hard one. Most of what we loosely call the "flavor" of a food is not taste at all — it is aroma, hundreds of volatile molecules rising up the back of the throat to the olfactory bulb, a sense with its own vast combinatorial code of some four hundred receptors, any handful of which can fire in concert to spell a smell. We left all of it off the Index, on purpose. Aroma is real and glorious and almost infinitely various, and that variety is exactly why it resists a clean scale: there is no small fixed set of axes, no half-dozen references, that could anchor it without arbitrariness. Taste proper — what the tongue's own receptors report — is different. It runs through a countable handful of pathways, each with a single responsible mechanism, and that countability is the only reason it is measurable at all. The Index is the taste the tongue reads, scored where it can be scored honestly, and it stops at the edge of smell rather than pretend to a precision it could not keep. Eight axes is not all of flavor. It is all of flavor that holds still enough to be a standard.

The honest edges do not end there, and holding them in view is part of the reverence. The reference anchors are a *convention,* not a law of nature — choose a different wild pepper for the heat anchor and every heat score rescales, which is why the anchors are fixed once, in these pages, and never quietly redefined. Because each axis anchors to its own distinct reference substance, IFI scores across axes are not positions on a single shared scale: a heat score of 2,000 and a pungency score of 2,000 each mean twice the potency of that axis's own reference, not equal intensities. The axes are also not perfectly independent when a real mouth meets a real food: heat and cool share a channel; kokumi lends body to umami; one axis can mask or magnify the next, which is a composition problem the standard takes up in its own right rather than a flaw in the measuring. And the single-stimulus assay, clean as it is, will always be a simplification of the living bite, where temperature and fat and the order of the courses all lean on the receptors at once. We state these not as apologies but as the terms of use. A number you can trust is a number whose limits are printed on the box.

The cornerstone of a library — *one standard, sixty more books.*

We did not build the Index to replace the poetry. The split tomato is still the split tomato; the warmth still arrives on the tongue like a held note. We built it because you cannot design what you cannot measure, and we had set ourselves a harder task than describing flavor. We meant to *compose* it — to decide, on purpose, that a fruit should carry this much savor and that much warmth and no bitterness at all, and then to write the plant that would grow into exactly that. A composer needs a staff and a clef. This is ours. The ten-rung heat ladder and the eight-voice chord of Phoenix are only the first proof that it works; the method behind them — scoring a receptor, inverting it against a reference, writing the enzyme that moves the number — transfers to any plant with a sequenced genome. **Forty-two host species are live on these eight axes today; sixty-four more are on the roadmap,** from a volcanic tomato to a sweet-heat strawberry, every one scored the same way, every one grown toward the fullest flavor it can carry.

This volume is also the cornerstone of a library. The eight axes and the shared index laid down here run on through some sixty more books — a deep volume on each axis and each major organism, cookbooks and cuisine books, the frontier books of what could still be built — and all of them resolve against the taxonomy defined in these pages: one umami, one kokumi, one Phoenix, set down once and referred to everywhere. The standard goes on from here to take the eight axes one at a time, down to the named enzymes and the organisms tuned to each; to the composition problem of what happens when axes meet and interfere; to the design loop that turns a target number into a living sequence; and to the error bars and scoring methods that keep every reading trustworthy. A reader who learns to read a profile here can read every profile in all of them. We wrote the standard first, and wrote it slowly, because everything else leans its whole weight on it. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Underneath all of it runs one conviction, quiet and total, the way it runs under everything our family makes: that food is something you give the body with intention, that a plant can be grown to hold the most nourishment it is able to carry, and that designing flavor and designing nourishment are not two acts but one. We are vegetarians who revere the living thing on the vine; we measured taste so we could build it well, and we build it well so it can feed you well. The number is not the opposite of the warmth at the table. It is how we keep the promise the warmth was always making. The number is where that begins.

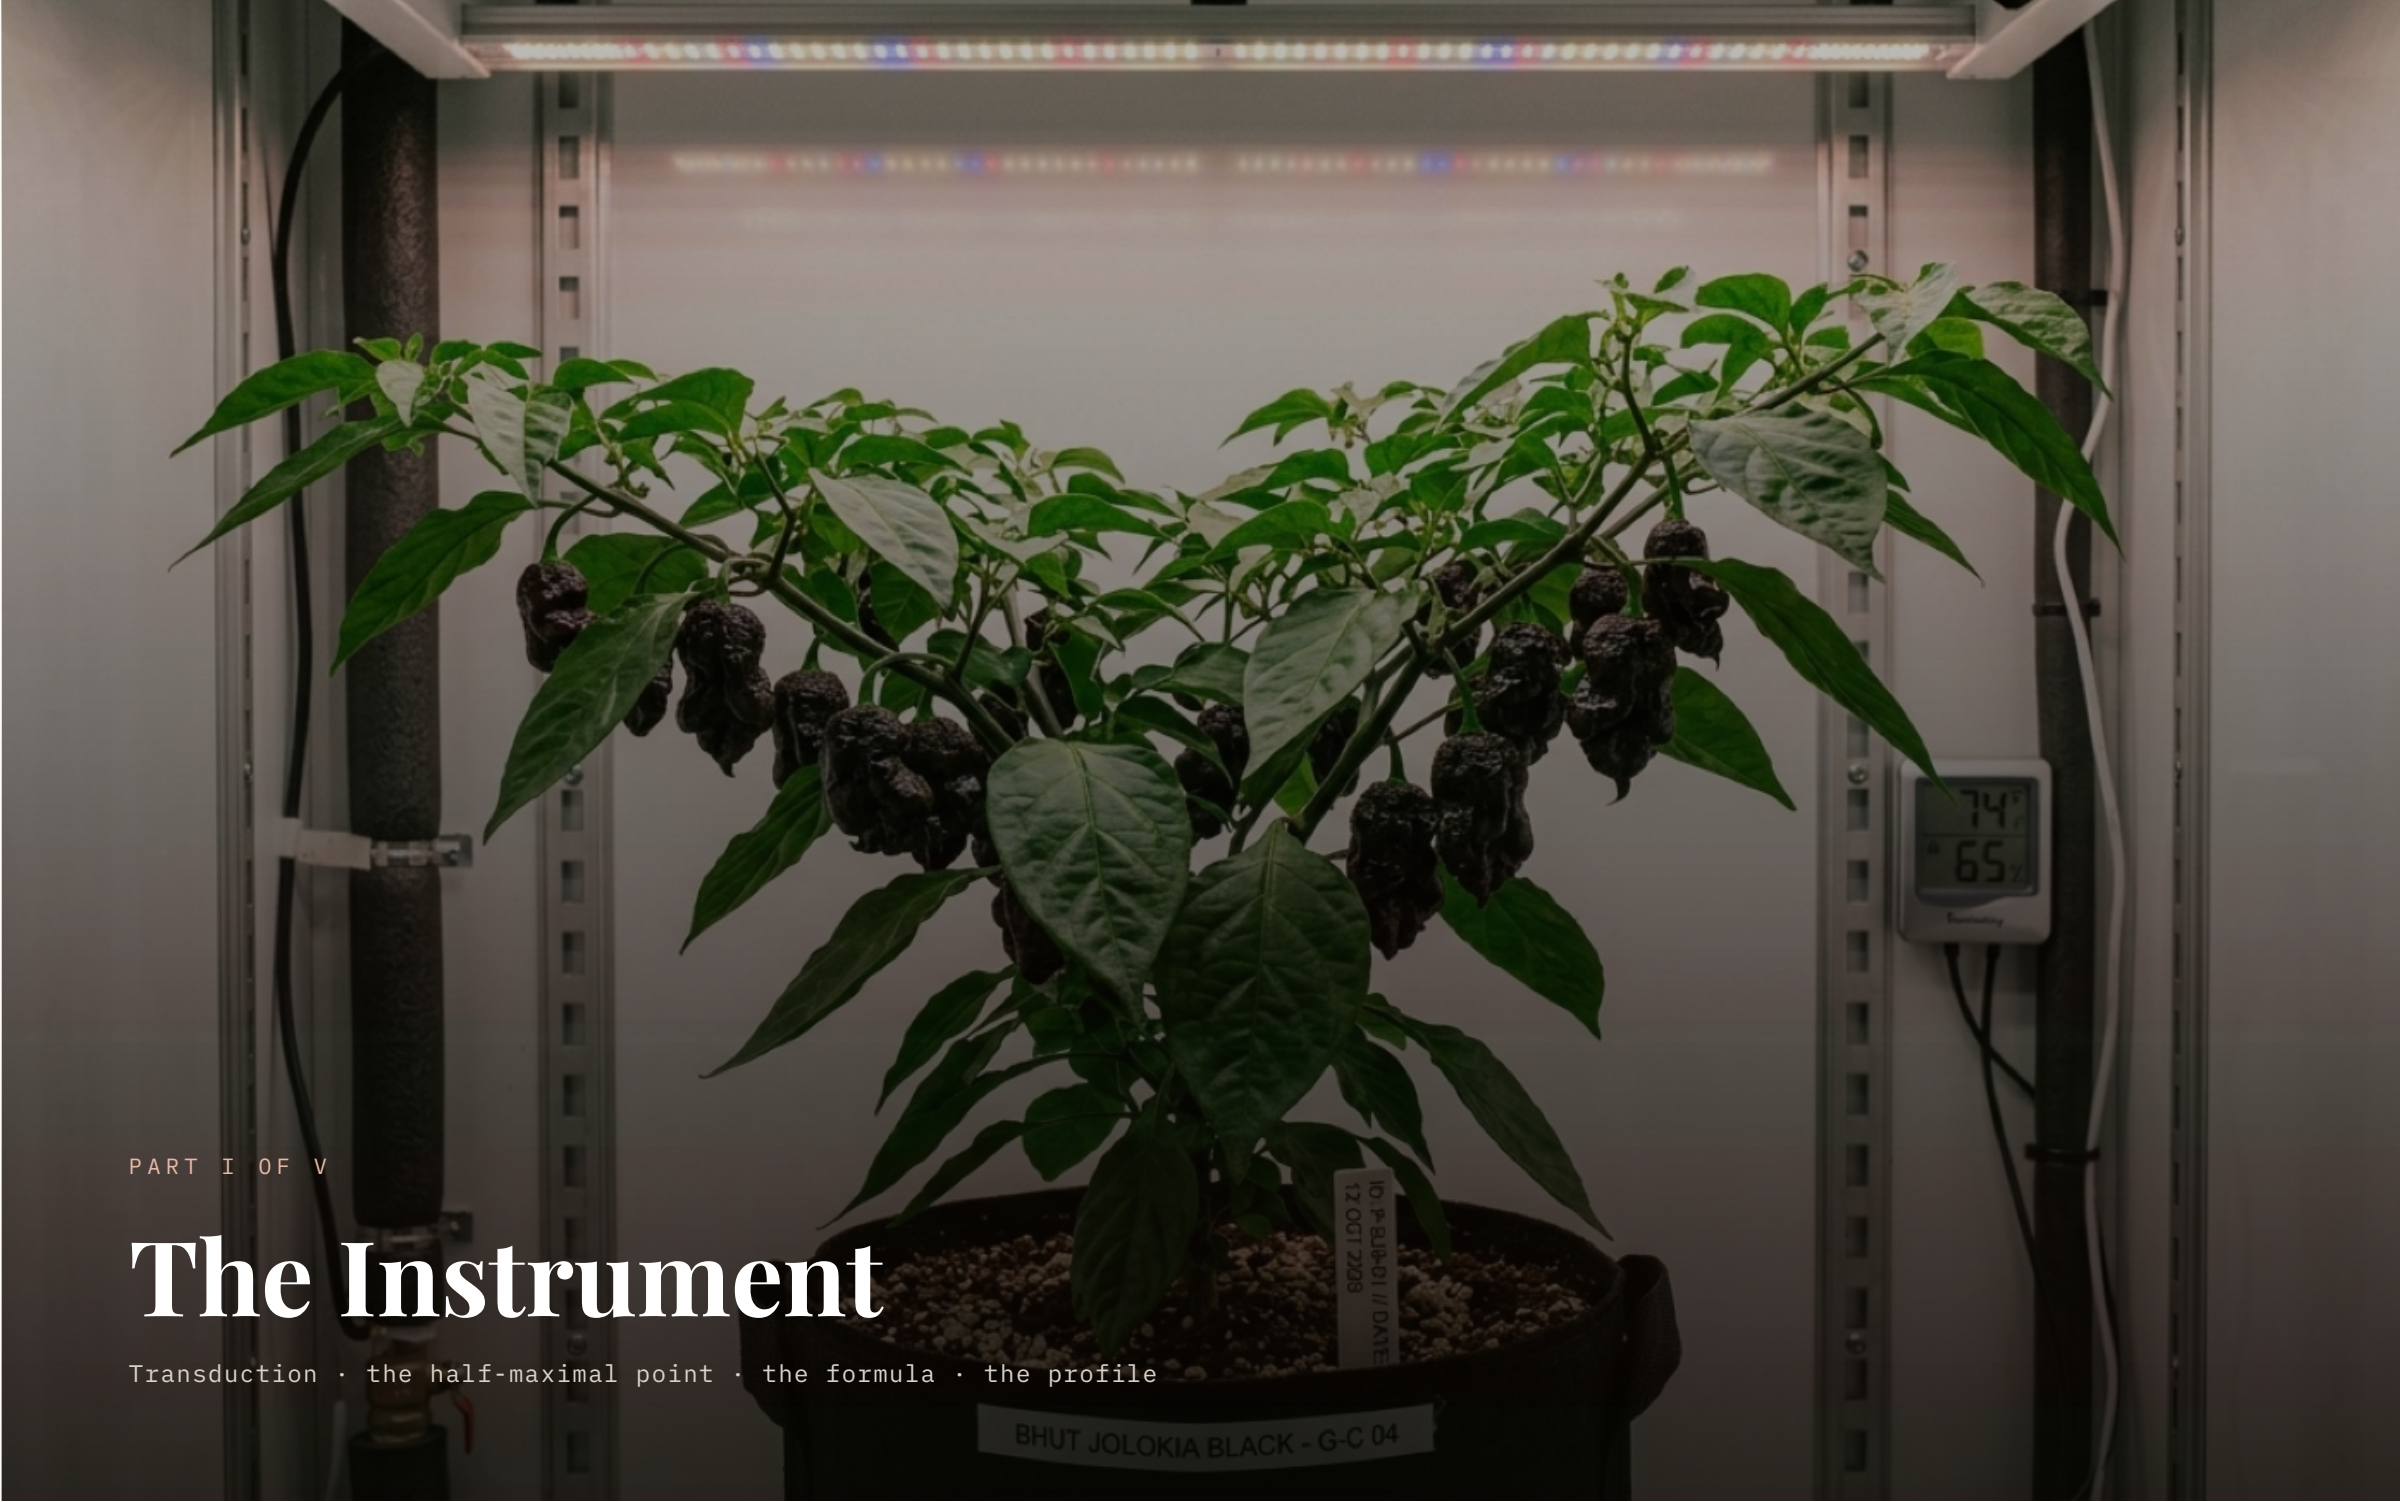

PART I OF V

The Instrument

Transduction · the half-maximal point · the formula · the profile

Chapter 1 · The Lock and the Key — *taste is a chemistry we can finally read.*

Before there is a number, there is a protein, and before the protein there is a molecule arriving where it is expected. This is the event the whole Index measures, so it is worth standing inside it — slowly, with care — to see what actually happens in the first millisecond a flavor touches the tongue. We tell it here with enough precision that every score in this book has a mechanism you can picture beneath it, because a number without its mechanism is only a rumor, and we would rather give you the thing itself. When our family set out to grow food that tastes as fully as it can and nourishes as fully as it can, we had to begin where taste begins: not in the kitchen, not even in the fruit, but at the receptor, the single point where what we grow finally meets the body it is meant to feed.

The surface you taste with

The surface you taste with is not the tongue's skin but a population of specialized cells gathered into taste buds — roughly fifty to a hundred cells per bud, bundled like the segments of a citrus, buried in the papillae you can see in a mirror and in many you cannot. There are four kinds of papillae and only three of them bear taste buds. The *fungiform* are dotted across the front two-thirds of the tongue, a handful of buds each, mushroom-capped and faintly pink. The *foliate* are ridged in parallel folds along the sides, toward the back. The *circumvallate* form a shallow V of eight to twelve large structures at the very rear, each one moated and crowded with hundreds of buds, so that most of your taste tissue sits where you least attend to it. The *filiform,* the velvety majority that gives the tongue its nap, carry no buds at all — they are purely mechanical, the tongue's own sense of touch, the thing that reads the grain of a food before any flavor does.

Into the pore atop each taste bud the dissolved molecules of a food diffuse, drawn down through saliva until they reach the apical membranes of the taste cells waiting there, and on those membranes sit the receptors. We say it plainly because it is the hinge of everything that follows: the receptors are the only part of the body that ever truly meets your food. The tongue feels its texture, the nose catches what rises from it, the gut will later take its measure — but the taste itself, the quality we are here to score, is a report filed by a few thousand proteins at the moment a molecule settles into them. Everything you will ever call flavor is their answer, relayed inward. When we grow a fruit to be read, this is the reader we are writing for.

Specialists and support staff

Inside a single bud the cells themselves divide into working classes, and the division matters more than its tidiness suggests. *Type II* cells are the receptor cells proper — they carry the G-protein-coupled receptors for sweet, umami, and bitter, and almost everything in this book about those three tastes is a story about Type II machinery. *Type III* cells are presynaptic, wired with classical synapses onto the nerve, and they read sourness through a dedicated proton channel, OTOP1, that lets the hydrogen ions of an acid pour directly into the cell. *Type I* cells are the glial-like majority, the housekeepers — they clear spent neurotransmitter, mop up stray potassium, and buffer the ionic traffic so that one cell's signal does not bleed into its neighbor's. A taste bud, in other words, is not a sensor. It is a small organ, with specialists and support staff, and its output is already a processed thing — sharpened, cleaned, separated — before it ever reaches a nerve. We find that worth a moment's reverence: the body does not merely detect the food we give it, it *curates* the meeting.



A taste-cell preparation on the microscope stage. The receptors studding these apical membranes are the only part of the body that ever meets your food.

Two families, three mechanisms

The receptors themselves come in two great families, and the difference between them sets the character of every axis in this book.

The first family is the **ion channels** — proteins that are themselves pores through the membrane, held shut until a stimulus opens them, whereupon ions pour through and the cell's voltage swings in a few thousandths of a second. Heat, pungency, and the cooling burn all run through ion channels, which is why they feel immediate and physical, less like a taste than like a thing *happening to you:* TRPV1 for capsaicin's fire, gated also by any temperature above about 43 °C and by the protons of an acid; TRPA1 for the sharp rush of

mustard and wasabi and cinnamon; TRPM8 for menthol's cool. These are the TRP channels — the transient receptor potential family — and each is built as four subunits ringed symmetrically around a single central pore, a gate that four identical walls must agree to open. They are nociceptors as much as they are taste organs: the very machinery that reports a scald or a cut reports a chili, which is the honest biological reason a hot pepper reads at once as pleasure and as alarm. And crucially they are not carried by the taste nerves at all but by the *trigeminal* nerve, the great sensory nerve of the face, so that what we loosely call "spicy taste" is really chemesthesis — chemical touch — riding a different cable into the brain than sweetness or umami will ever use. Two of our eight axes, heat and pungency, and much of a third, cooling, live on this trigeminal line, and that is why they behave on the Index like relatives who share a temper.

The second family is the **G-protein-coupled receptors,** the GPCRs — and these do not conduct anything at all. A GPCR is a single chain of protein threaded back and forth through the membrane exactly seven times, a seven-pass helix bundle with a binding site cradled in or above it. When its molecule lands, the receptor does not open; it *changes shape,* and that conformational flex on the inner face of the membrane shoves a waiting G-protein off its seat. Sweet, umami, bitter, and the hidden richness of kokumi all speak through GPCRs, which is why they feel rounder and slower and more like a quality than an event. To design for them, we first had to be able to see them — to read a protein at the resolution of the single residue, and know which wall a molecule meets and which pocket it must fall into.

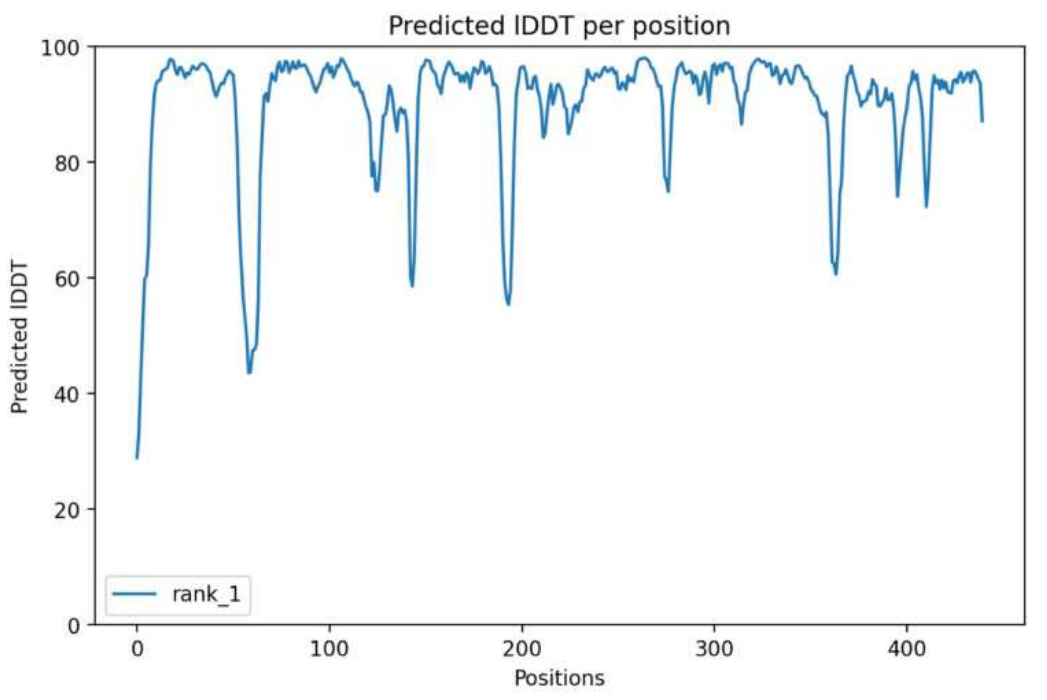

A protein read at the resolution of single residues, each colored by how confidently its fold is known — blue certain, orange less so. This is an enzyme from our own capsaicin pathway rather than a taste receptor, but it shows the kind of sight now available: the cleft, the pocket, the exact surface a molecule arrives against.

Astringency belongs to neither family, and we keep it on the Index anyway because the tongue does. It is not a chemical taste at all but a mechanical one: tannins and other polyphenols binding and precipitating the salivary proline-rich proteins that keep your mouth slippery, so that what you perceive is not a flavor but a sudden loss of lubrication — a drag, a pucker — reported by the same touch and

trigeminal fibers that feel texture. Eight axes, then, resolve into three mechanisms: a channel that opens, a cascade that blooms, and a friction that sets in. The Index does not pretend these are the same kind of thing. It measures each on its own terms, and only then, by anchoring all of them to the half-maximal point, makes them comparable without ever making them identical.

CODE	AXIS	RECEPTOR	MECHANISM FAMILY
IFI-H	Heat	TRPV1	ion channel · trigeminal
IFI-P	Pungency	TRPA1	ion channel · trigeminal
IFI-C	Cool-Burn	TRPM8 / TRPV1	ion channel · trigeminal
IFI-U	Umami	mGluR1 / TAS1R1·TAS1R3	GPCR cascade
IFI-S	Sweetness	TAS1R2 / TAS1R3	GPCR cascade
IFI-B	Zero-Bitter	TAS2R38	GPCR cascade
IFI-K	Kokumi	CaSR	GPCR cascade
IFI-A	Astringency	salivary PRPs	protein precipitation · friction

The cascade, step by step

Because the GPCR tastes carry four of our eight axes, the cascade they share is worth naming in full — it is identical for sweet, umami, and bitter, and it explains why all three feel the way they do. Follow it once and you will never again mistake a taste for a simple thing.

The molecule lands and the receptor flexes. That flex releases **gustducin,** the taste cell's own G-protein, a transducin cousin that until this instant sat clamped and waiting on the membrane's inner face. Gustducin's release activates an enzyme, **phospholipase C-β2 (PLCβ2),** which finds a particular lipid in the membrane and cleaves it, splitting off the small second messenger **IP_3.** IP_3 drifts to the cell's internal calcium store and opens it, and **calcium** floods out into the cytoplasm — the signal, having begun as a shape, is now a tide of ions. That calcium opens a channel called **TRPM5,** permeable to sodium, and the inrush of sodium finally **depolarizes** the cell. The depolarization opens one last channel, **CALHM1,** a wide pore through which the cell releases **ATP** as its outgoing transmitter — not into a classical synapse but straight across a narrow cleft onto the $P2X_2$ and $P2X_3$ receptors studding the waiting gustatory nerve. Five or six steps, each taking its own small moment: shape, enzyme, messenger, calcium, voltage, and finally the puff of ATP that the brain will read as *sweet,* or *savory,* or *bitter.* This is precisely why sweetness and umami and bitterness feel rounder and slower than heat — a chord that blooms rather than a note that strikes. Heat kicks a single door; the GPCR tastes run a relay, and you can feel the relay in the shape of the sensation.

Within the GPCRs there is a further split that will matter again and again, so fix it now. The **class C** receptors — sweet, umami, and the kokumi-sensing calcium receptor — carry a huge clamshell-shaped extracellular domain called the Venus flytrap, and they read the tastes the body *wants.* They are choosy, cooperative, built to recognize specific nourishing molecules and to close around them like two hands cupping water. The **class A** bitter receptors, the TAS2R family, are compact, and they are deliberately promiscuous: built to recognize *poison,* of which the catalogs of the living world

hold an enormous and ever-changing variety. There are some twenty-five of them in a human mouth precisely because the body would far rather flag a harmless bitter as dangerous than let a single real toxin pass unremarked. That architectural difference — the discerning clamshell against the broad alarm panel — is written into how each axis behaves, and into how far each can honestly be designed. You can coax a want; you cannot easily sweet-talk an alarm, which is why lifting bitterness is subtraction, not persuasion.

From receptor to percept

Where does the signal go once the ATP is released? The gustatory nerves gather it — the facial nerve by way of the chorda tympani from the front of the tongue, the glossopharyngeal from the back, the vagus from the throat beyond — and carry it first to a relay in the brainstem, the nucleus of the solitary tract, then up through the ventral posteromedial thalamus to the insular cortex, the brain's primary taste map. There, at last, receptor occupancy becomes perception, and perception meets memory and appetite and the body's running ledger of what it needs and what it has had too much of.

Two old theories argue over how a quality survives that journey. The *labeled-line* model holds that a sweet cell drives a sweet fiber drives a sweet percept — each quality on its own dedicated wire, kept separate end to end. The *across-fiber* model holds that quality is read instead from the pattern across many broadly tuned cells at once, a chord inferred rather than a note delivered. The weight of the evidence favors mostly labeled lines at the periphery — a sweet cell really is, in the main, a sweet cell — and that is a quiet gift to the Index, because it means an axis measured at the receptor corresponds to a sensation the brain keeps largely apart from the others. We score eight pathways because the tongue, in large part, keeps eight pathways. The Index did not invent the separation; it inherited it, and tried only to be faithful.

Two governors: saturation and adaptation

Two properties of this machinery govern everything that comes after, and both are the reason a flavor needs a particular *kind* of number rather than just a quantity.

The first is **saturation.** There are only so many receptors on a taste cell, and once a stimulus occupies most of them, more molecules change almost nothing — the cell is already shouting as loudly as it can. We can write this exactly. If we call c the concentration of a stimulus and EC_{50} the concentration that drives the receptor to half of its full response, then for simple binding the fractional response θ follows

```
θ(c) = c / (EC₅₀ + c)
```

and that single curve contains the whole behavior. At c equal to EC_{50}, θ is one-half — the definition of EC_{50}, read straight off the equation. Push further and the returns collapse:

STIMULUS C (IN MULTIPLES OF EC_{50})	FRACTION OF FULL RESPONSE, θ
0.11 ×	0.10
0.33 ×	0.25
1.00 ×	0.50
3.0 ×	0.75
9.0 ×	0.90
19 ×	0.95
99 ×	0.99

Read the bottom of that table slowly, because it is where the Scoville number and the honest truth of the tongue part ways. To climb from ninety percent of full heat to ninety-nine percent, you must raise the stimulus from nine times EC_{50} to ninety-nine times — an elevenfold increase of capsaicin to buy the last nine points. To reach 99.9 percent you would multiply the stimulus by ten again. This is why doubling the capsaicin in an already-fierce pepper barely moves the perceived heat, and why every axis has a ceiling it approaches but never touches. The molecule obeys arithmetic; the sensation obeys saturation.

Real channels are in fact steeper than the simple curve above, because gating is cooperative. TRPV1's four subunits do not decide independently — capsaicin binding at one site makes its neighbors likelier to open — so the response follows a Hill form, $\theta(c) = c^n / (EC_{50}^n + c^n)$, with an exponent n greater than one. The exponent sharpens the switch: the concentration span from ten percent to ninety percent response is 81^(1/n). For a non-cooperative receptor (n = 1) that span is an eighty-one-fold range of concentration; at n = 2, roughly TRPV1's cooperativity, it narrows to just ninefold; at n = 4 it would be a mere threefold window between barely-warm and fully-ablaze. Heat, in other words, is not only saturating but abrupt — it comes on like a threshold crossed, not a dial turned, and the mathematics says exactly how abrupt.

The second governor is **adaptation.** Hold a stimulus steady and the response fades: the channels inactivate, the cascade down-regulates its own gain, and the fire of a chili settles toward a warmth even while the capsaicin is still bound and present. Flavor is not a photograph of what is in your mouth. It is the nervous system's running, self-adjusting commentary on it — a rise, a peak, and a decay that are as much a part of a flavor's identity as its loudness. We measure the peak, because the peak is what can be pinned without argument; but we never forget that the peak is a moment in a story the receptor is telling itself.

The envelope — flavor in time

One more property deserves its own name, because it is half of what a flavor *is:* time. A taste is not an instant but an envelope — an onset as the molecules diffuse into the pore and find their receptors, a rise to a peak, and a decay as adaptation and clearance set in — and the eight axes do not share an envelope. Heat is slow to arrive and slow to leave, pooling and lingering because capsaicin is oily and loath to be washed away by watery saliva; its decay constant is measured in long seconds. Pungency is fast on both ends, a volatile flash up the back of the nose and gone, because the small reactive molecules

that drive TRPA1 clear almost as quickly as they came. A sweet protein can hang on the tongue long after ordinary sugar would have rinsed off, its clamshell slow to release. Astringency does the strangest thing of all: it *builds* across sips, each mouthful spending more of the salivary protein that resists it, so that the fourth taste of a tannic tea is sharper than the first though the cup has not changed.

Two flavors with identical peak loudness on every axis can therefore be entirely different experiences if their envelopes differ — the one blooming and fading like a held chord, the other stabbing and gone. The Index scores the peak, because the peak is what can be anchored to a half-maximal point and measured without argument. But our design records keep the envelope beside every score — the rise time and the decay time of each axis — because a flavor remembered is a flavor in time, and a fruit grown to be eaten slowly, with attention, deserves to be measured the way it is lived.

Why a flavor needs a number — the half-maximal anchor

Here, then, is the exact reason a flavor needs a number, and a particular kind of number. You cannot score taste by chemistry alone, because the body does not experience chemistry — it experiences receptor occupancy, with all of saturation and adaptation already folded in. A measure of *how much compound a food contains* will lie to you at both ends of the scale: it will exaggerate the top, where saturation has flattened the sensation, and it can mislead at the bottom, where cooperativity makes a little go a long way. To measure loudness honestly you must ask the receptor its own question: at what concentration does *it* reach half of everything it can do? That half-maximal concentration, EC_{50}, is the one point on the whole dose-response curve that is both unambiguous and perceptually meaningful — the steepest, most informative place on the curve, the inflection the sigmoid turns around. Anchor there and you are measuring the sensation, not the sample.

So we define the Iyengar Flavor Index for any candidate against a fixed reference — a stated convention constant, defined by our scale in the same way the Scoville scale is defined by its dilution-to-threshold, not a re-derived literature value — as

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

Three things follow immediately, and together they make the scale behave. First, the reference scores exactly 1,000 by construction: substitute its own EC_{50} on both sides and the ratio is one, times a thousand. A thousand is therefore the Index's sea level — native, untuned, the flavor as the plant already knew how to make it. Second, a *louder* candidate — one that reaches half-maximal response at a *lower* concentration, because it delivers its signal more potently — has a smaller EC_{50} in the denominator and so a *higher* score. Loudness and score rise together; the inversion is the whole trick. Third, because the score is a ratio of concentrations, it is naturally logarithmic in spirit: doubling the IFI means halving the EC_{50} means the stimulus is twice as loud at the receptor. For scorecards

we cap each axis at a normalization maximum of 5,000, which sets the native reference at one-fifth of full scale and leaves real headroom above it for everything design can add.

Each axis rests on its own convention constant, so an IFI score on one axis and the same score on another are not the same "loudness" in any shared unit — each is a multiple of that axis's own defined reference. Among the eight axes, IFI-data derived from published EC_{50} records exists for the heat axis (capsaicin/TRPV1) and the pungency axis (AITC/TRPA1); the remaining six axes are anchored to their stated convention constants and have no separate IFI-data version.

A worked reading makes it concrete. Take IFI-H, the heat axis. Its convention constant sets EC_{50}(reference) at 50.0 milligrams of fruit per millilitre in a standardized whole-fruit extract assay — a defined anchor, the same way the Scoville threshold is a defined stopping point, not a re-derived measurement. The formula then turns any cultivar's score directly into the concentration of itself it takes to half-activate the channel — EC_{50}(candidate) = 50,000 ÷ IFI — and into a plain multiple of the reference's loudness, IFI ÷ 1,000.

CULTIVAR	SKU	SHU (PREDICTED)	IFI-H	EC_{50} (MG FRUIT/ML)	SHU PER IFI POINT
native reference	—	—	1,000	50.0	–
Ember	SS-0087	866K	1,570	31.8	552
Cinder	SS-0088	1.05M	2,290	21.8	459
Kindle	SS-0089	1.25M	2,000	25.0	625
Blaze	SS-0091	1.72M	2,630	19.0	654
Wildfire	SS-0092	1.96M	2,930	17.1	669
Scorch	SS-0093	2.21M	3,220	15.5	686
Firestorm	SS-0094	2.47M	3,490	14.3	708
Solara	SS-0095	2.72M	3,980	12.6	683
Phoenix	**SS-0096**	**2.95M**	**4,200**	**11.9**	**702**

Read the EC_{50} column and the saturation story is suddenly physical. Where the native reference convention anchor is 50 milligrams per millilitre, our model predicts Phoenix needs only 11.9 — a quarter of the reference mass to half-activate the channel, because gene by gene we designed it to speak louder, not merely to carry more. Its score of 4,200 is exactly 4.20 times the reference loudness, and it sits at 84 percent of the 5,000 ceiling, which is to say we engineered most of the available headroom on this one axis and left a deliberate

little room at the top. Phoenix is also the only cultivar in the flagship line that our model places on all eight axes at once, and the full reading is worth seeing whole:

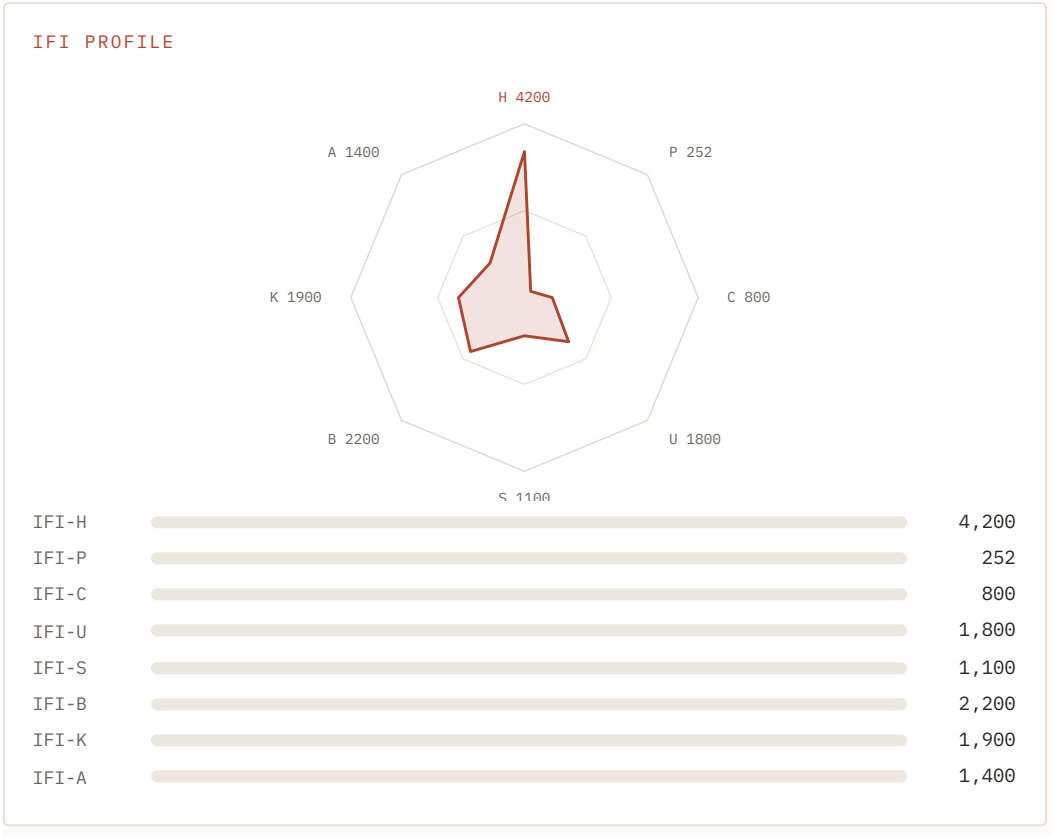

Eight numbers, one fruit — a chord struck across every mechanism the tongue owns, each value the honest half-maximal loudness of its own axis, none of them borrowed from another. That profile is what the rest of this book learns to read.

What the number is not — SHU, error, and the honest limits

A new scale earns trust by being clear about what it is *not,* so we will be plain about four things.

First, the Index is not the Scoville scale rescaled. Look again at the last column of the ladder. The Scoville Heat Unit is, in its original meaning, a dilution — how far you must dilute an extract before a panel can no longer detect its heat — and by the standard equivalence of about sixteen Scoville units per part-per-million of capsaicinoid, SHU tracks the *mass* of capsaicin in the fruit almost linearly. The IFI tracks *receptor occupancy,* which saturates. Two scales anchored at two different points on one dose-response curve — SHU near the vanishing threshold, IFI at the half-maximal inflection — must part company as the curve bends over, and the "SHU per IFI point" column is that parting, measured: it climbs from roughly 550 at the bottom of the ladder toward 700 near the top, meaning each additional unit of receptor-anchored loudness costs more and more raw capsaicinoid as the channel approaches saturation. That rising cost is the arithmetic signature of the flattening curve. And the clearest proof that one scale is not the other hides in two rungs: Kindle carries a higher predicted Scoville value than Cinder — 1.25 million against 1.05 million — yet scores *lower* on IFI-H, 2,000 against 2,290. Raw capsaicinoid mass and half-maximal loudness do not even rank the ladder in the same order. They cannot, because they are answers to different questions, and the Index asks the one the tongue actually poses.

Second, the Index is not noiseless, and we refuse to pretend it is. Every score carries a batch-to-batch coefficient of variation — the standard deviation as a fraction of the mean — and that is where the engineering behind these cultivars does its quietest work. Our designed lines hold a CV at or below 7 percent; the wild background from which they came runs near 42 percent. The difference decides whether the ladder above is real or imagined. At 7 percent, Phoenix's 4,200 carries a one-standard-deviation band of about ±294 points, roughly 3,900 to 4,500. At the wild-type's 42 percent, that same band would sprawl from about 2,400 to nearly 6,000 — wide enough to swallow five rungs of the ladder and reduce the whole careful sequence to a blur. Even at our tight 7 percent, honesty has a price near the top: Phoenix at 4,200 and Solara at 3,980 are separated by only 220 points, which is less than either one's single-batch standard deviation, so no one batch can tell them apart. Resolving them means replication. The standard error shrinks as the square root of the number of batches, and working it through — a pooled difference standard deviation near 405 points, a target separation of 220 at the ninety-five-percent level — asks for about thirteen independent batches before the two cultivars stand cleanly apart. That is a real cost, and we pay it rather than round it away. A number offered without its error bar is not a measurement; it is a decoration.

Third, the Index does not claim to have read every rung. Flare, SS-0090, carries a predicted Scoville value of 1.48 million in the flagship line and four expressed axes, but we have not yet fixed a settled half-maximal reading for its heat, so it carries no IFI-H here and we will not invent one to tidy the table. An honest gap is worth more than a confident guess, and the Index is built to hold the gap visibly until the assay closes it.

Fourth, and most important, EC_{50} measures loudness and loudness alone. It does not measure whether a flavor is *good.* It says nothing about hedonics — the swicy delight of heat laid over sweetness, the way kokumi makes a broth feel cared-for, the context and memory and hunger a person brings to the table. It does not fully capture how axes interact when they arrive together, where one receptor's cascade can borrow or steal from another's. It is computed against a standardized model, and a real food's fats and proteins and its own architecture will shift delivery in ways the design record must note separately. The Index is a faithful measure of one dimension of a far richer thing. We built it because you cannot design what you cannot measure — but we have never mistaken the measure for the meal.



Phoenix in the controlled chamber, the cultivar read against the native reference — the first we grew to register on all eight axes at once.

The reader we grow for

Return, at the end, to the beginning: a molecule arriving where it is expected. Everything in this chapter — the papillae and the pore, the three cell types, the channel that opens and the cascade that blooms and the friction that sets in, the saturating curve and the half-maximal anchor and the thousandfold scale laid over it — all of it exists to describe one quiet meeting, repeated a billion times a second, between the food and the only part of the body that will ever truly touch it. We find it moving that taste, which felt for ten thousand years like pure poetry, turns out to be a chemistry we can finally read, residue by residue and axis by axis, without losing one grain of the wonder. The number does not replace the pleasure. It lets us grow toward it on purpose. When our family tunes a plant gene by gene so that it meets this reader more fully — louder where loudness is a joy, richer where richness is nourishment, clean in everything it carries into the body — we are, in the end, only writing a more generous letter to the receptors on your tongue, and trusting them, as we always have, to read it. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 2 · The Half-Maximal Point — *why EC50 is the only honest unit.*

Give a receptor a vanishing trace of its molecule and nothing happens; the signal is lost in the cell's own thermal noise. Flood it and the response saturates; every channel that can open is open, every flytrap that can close is closed. Between those two silences lies a curve, and that curve — the dose-response relationship — is the most important shape in this book. Learn to read it and you can read every score in the Index, because every score is a single coordinate lifted off one of these curves.

We came to this curve the long way, the way most families come to the things they end up loving: by trying to say something we could feel but not yet measure. You can stand in a greenhouse at first light with a cut fruit in your hand and know, past any argument, that this one is *louder* than the one beside it. What you cannot do, standing there, is say how much louder, or why, or whether tomorrow's fruit from the same plant will say the same thing. The curve is where "how much" stops being poetry and becomes a number — and where a number, measured honestly, becomes a kind of respect for the person who will eventually eat the thing.

The shape of the climb

Plot the receptor's response on the vertical axis against the concentration of the stimulus on the horizontal, and you get a rising S. At the foot it hugs the floor, indifferent. Somewhere in the middle it turns sharply upward, each small increase of concentration buying a large increase of response. Then it bends over and flattens toward a ceiling — the maximal response, **R_max,** the loudest the cell can be. The curve is sigmoidal, and the mathematics that describes it is the Hill equation, written for a flavor stimulus as:

$$R = R_{max} \times [A]^n \div (EC_{50}^n + [A]^n)$$

Here **[A]** is the concentration of the stimulus, **n** is the Hill coefficient that sets how steeply the curve climbs, and **EC_{50}** is the one term we care about most: the concentration at which the response **R** equals exactly half of **R_max.** Set [A] equal to EC_{50} in that equation and watch it collapse — the numerator becomes EC_{50}^n, the denominator becomes $EC_{50}^n + EC_{50}^n$, which is $2 \cdot EC_{50}^n$, and the fraction is one-half, exactly, for any value of n whatsoever. This is the quiet miracle that makes the whole scheme work. The half-maximal point does not depend on the steepness, or the ceiling, or the chemistry of the particular molecule. It is a defined place on a fitted curve: the concentration at which the receptor is doing exactly half of everything it is capable of doing.

In practice we never read that point off a single measurement. We fit the whole curve at once with a four-parameter logistic model — floor, ceiling, slope, and midpoint — across a dilution series run on a cell line expressing the receptor, and we read EC_{50} off the fit with a confidence interval attached, because a number without its uncertainty is not yet a measurement. The fit uses every well at once; no single point carries the answer, and no single point can be bribed into giving a prettier one.

Why the middle, and not the ends

There is a real reason the midpoint is the honest place to stand, and it is worth making exact rather than asserting. Work in fractional response — call it f, the fraction of R_max the receptor is currently producing, so f runs from 0 to 1. The sensitivity of the receptor is how much f moves when you nudge the concentration by a fixed *proportion* — the derivative of f with respect to the natural logarithm of concentration. For the Hill curve that derivative works out to a clean expression:

```
df ⁄ d(ln[A]) = n · f · (1 − f)
```

The product f·(1 − f) is a downturned parabola in f. It is zero at both ends — f = 0 at the silent floor, f = 1 at the saturated ceiling — and it reaches its single maximum at f = ½, where it equals ¼. So the receptor's sensitivity peaks exactly at the half-maximal point, and there it equals *n*⁄4. This is not a convention we chose; it is where the curve itself is listening hardest. Near the floor the response is drowned in noise and a tenfold change in concentration barely stirs it; near the ceiling it is lost in saturation, where enormous changes of concentration hardly nudge the signal. The midpoint is the one place where a change in concentration maps most cleanly onto a change in response — the one concentration you can pin down with the least error. Every pharmacology of the last century measured potency this way, for exactly this reason; the EC_{50} and its antagonist cousin the IC_{50} are the common currency of drug discovery. We adopted their unit because it was already the right one. Our contribution was to turn it into a scale a cook could read.

Measuring it, so the number can be trusted

Measuring an EC_{50} is a small discipline, and it is where the number earns the right to be trusted. We express the receptor in a cultured cell line — no tasting panel, no sugar water, no one asked to swallow anything on our behalf — and load the cells with a dye that flares when calcium floods in through an opening channel. The stimulus is presented as a graded dilution series: a dozen or more concentrations spanning several orders of magnitude, each in replicate across a plate, each well filmed as the dye answers. Every well returns a response; the responses trace the sigmoid; a nonlinear regression fits the four-parameter logistic to the whole cloud of points at once and returns the half-maximal concentration with a confidence interval wrapped around it.



The calcium-imaging stage where a dilution series becomes a curve — each well a concentration, each flash of dye a receptor answering.

Then we do it again — another day, another plate, another passage of cells — and the midpoints must agree within their intervals or the measurement is discarded. Only an EC_{50} that reproduces becomes a reference or a score. The discipline is mostly in the failures, and they are worth naming plainly, because an honest unit is one that knows how it can lie. A molecule that is poorly soluble will seem to flatten short of its true ceiling simply because it falls out of solution before it can reach the top of the curve; read naively, it looks like a partial agonist when it is only a chemistry problem. A molecule that is faintly toxic will push the baseline *down* at high dose as cells sicken, carving a false shoulder into the curve. A receptor that desensitizes — that tires and closes even while the agonist is still present — will under-report at long exposure unless the read is taken in the first clean seconds. Each of these is a way to get a wrong EC_{50}, and each is why we fit the whole curve, bracket every plate with known controls, and throw out anything that will not reproduce. The curve is not decoration and the dilution series is not ceremony; they are the evidence, and we keep all of it.



A dilution series laid out in replicate — a dozen concentrations across orders of magnitude, each repeated, before a single EC_{50} is allowed to exist.

Potency runs backward

The decisive move is to recognize that **EC_{50} runs backward to potency.** A molecule with a small EC_{50} drives its receptor to half-response at a tiny concentration — it is potent, it is *loud.* A molecule with a large EC_{50} must pile up before the receptor answers — it is weak, it is *quiet.* Potency is the reciprocal of EC_{50}. Pharmacology gives this reciprocal its own name and its own unit: pEC_{50}, the

negative base-ten logarithm of the half-maximal concentration in molar. A micromolar EC_{50} is a pEC_{50} of 6; a tenfold more potent molecule, needing only a tenth of the concentration, climbs to a pEC_{50} of 7. Louder is higher. We will need that logarithm again in a moment, because it is the hinge between a pharmacologist's curve and a flavor score — and because the mouth, as it happens, hears in logarithms too.

From a curve to a score

If loudness is one-over-EC_{50}, then to compare two stimuli you compare their reciprocal EC_{50}s, and to anchor the whole thing you compare both against a fixed reference. Choose, per axis, a reference stimulus — by convention the half-activation of the native wild-type — and call its half-maximal concentration EC_{50}(ref). Then define:

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

The units cancel — both EC_{50}s are concentrations — so the score is a pure number, a ratio of loudness to the loudness of the reference, multiplied by a thousand so the reference itself lands on a round, memorable anchor. Set the candidate equal to the reference and the formula returns exactly 1,000; the reference is the origin of the coordinate system, reading the same on every axis by construction:

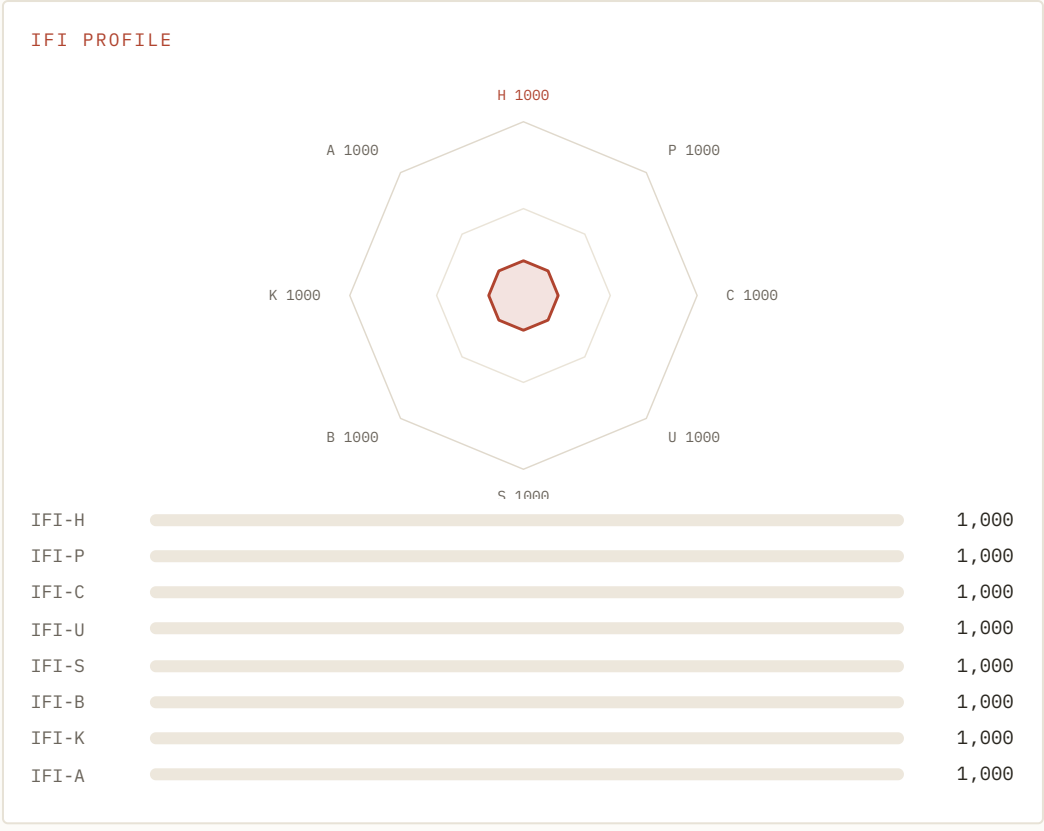

This is not a plant. It is the zero of the ruler — the loudness against which every real fruit is read. From here the arithmetic is immediate. A candidate whose stimulus reaches half-activation at one-quarter of the reference concentration has EC_{50}(candidate) = 0.25 × EC_{50}(ref), so IFI = 1,000 × (1 ⁄ 0.25) = 4,000: four times as loud, four thousand on the scale. A candidate that needs five times the reference concentration has IFI = 1,000 ⁄ 5 = 200: quieter than the

reference, and the score says so. The multiplier and the reciprocal between them turn a half-maximal concentration into a number you can set a tomato, a chili, and a cup of tea against on one page.

Take the logarithm of that definition and its character shows itself. Because log(1,000) = 3, and because log of a ratio is a difference of logs:

$$\log_{10} \text{IFI} = 3 + [\ pEC_{50}(\text{candidate}) - pEC_{50}(\text{reference})\ ]$$

The logarithm of the score is just 3 plus the candidate's potency *relative to the reference,* measured in pEC_{50} units. One full pEC_{50} step — a tenfold change in EC_{50} — is a tenfold change in the score. This gives a reading key that holds on every axis at once:

IFI SCORE	Δ PEC_{50} VS REFERENCE	LOUDNESS VS REFERENCE	HOW TO READ IT
100	−1.00	0.1×	a tenth as loud — ten times the stimulus to reach half-response
500	−0.30	0.5×	half as loud as the reference
1,000	0.00	1.0×	the reference anchor itself
2,000	+0.30	2.0×	twice as loud
4,200	+0.62	4.2×	Phoenix on the heat axis
5,000	+0.70	5.0×	the scorecard's normalized ceiling
10,000	+1.00	10×	a full order of magnitude above reference

Run our designed heat ladder through the same formula and it resolves into a table where every column is honest about a different thing. SHU is a legacy content assay; IFI-H is receptor loudness; the fourth column is the half-maximal concentration each cultivar is designed for, expressed as a fraction of the reference's, which falls out of the score as $EC_{50}(\text{candidate})/EC_{50}(\text{ref}) = 1{,}000/\text{IFI-H}$:

CULTIVAR · SKU	SHU (PREDICTED)	IFI-H	EC_{50} (× REFERENCE)	LOUDER THAN REFERENCE
Ember · SS-0087	866K	1,570	0.637	1.57×
Cinder · SS-0088	1.05M	2,290	0.437	2.29×
Kindle · SS-0089	1.25M	2,000	0.500	2.00×
Blaze · SS-0091	1.72M	2,630	0.380	2.63×
Wildfire · SS-0092	1.96M	2,930	0.341	2.93×
Scorch · SS-0093	2.21M	3,220	0.311	3.22×
Firestorm · SS-0094	2.47M	3,490	0.287	3.49×
Solara · SS-0095	2.72M	3,980	0.251	3.98×
Phoenix · SS-0096	2.95M	4,200	0.238	4.20×

SHU values for Scoville Splice cultivars are computational predictions, not certified measurements.

Read the last row as the whole method in miniature. Phoenix carries a predicted IFI-H of 4,200, meaning its capsaicin design drives TRPV1 to half-response at 23.8% of the reference concentration in our model — a little under a quarter — and 4.2 times as loud. The same plant reads on all eight axes at once, the first we designed to do so, and its full profile is the thing the rest of the Index exists to make legible:

IFI PROFILE

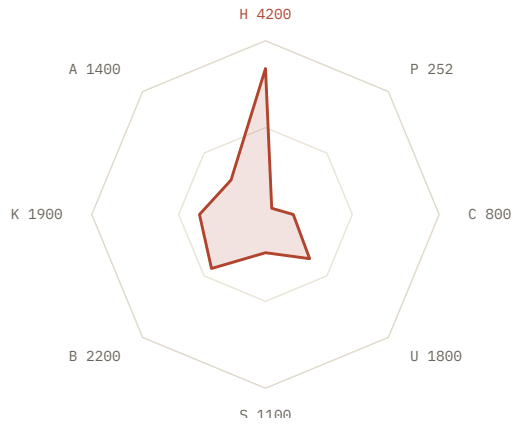

IFI-H		4,200
IFI-P		252
IFI-C		800
IFI-U		1,800
IFI-S		1,100
IFI-B		2,200
IFI-K		1,900
IFI-A		1,400

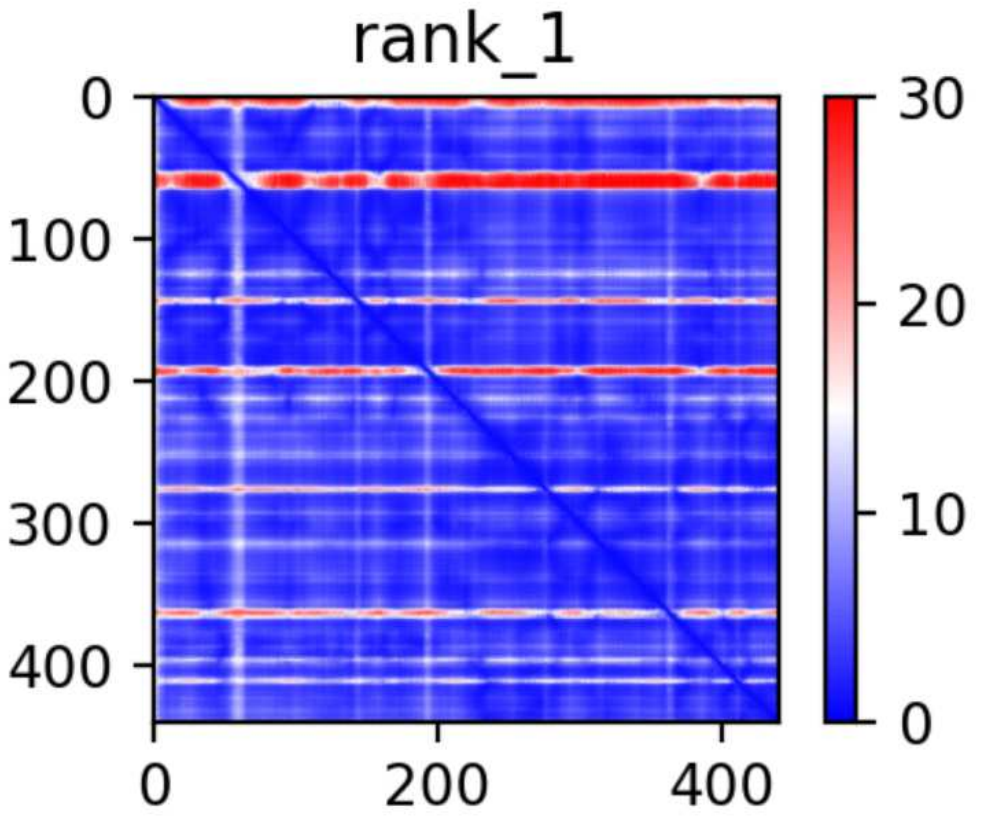

A predicted response surface from the structural pipeline: potency is read off a whole fitted shape, never a single point.

The Hill coefficient, or the shape of the experience

Two subtleties keep the half-maximal point honest, and they are the difference between a number that respects the reader and one that merely impresses. The first is the **Hill coefficient, n** — the steepness of the climb, and a quiet carrier of meaning. When n is near one, the receptor binds one molecule at a time, independently, and the curve

is gentle; the stimulus comes on gradually across a wide range of concentration. When n is greater than one, binding is *cooperative* — one molecule landing makes the next land more easily — and the curve becomes a cliff: almost nothing, then suddenly everything, across a narrow band. Put two curves side by side, scaled so both cross their half-point at the same place, and the coefficient's work is plain:

[A] / EC_{50}	FRACTIONAL RESPONSE, N = 1	FRACTIONAL RESPONSE, N = 2
0.1	0.091	0.010
0.25	0.200	0.059
0.5	0.333	0.200
1.0	0.500	0.500
2.0	0.667	0.800
4.0	0.800	0.941
10	0.909	0.990

Both pass through exactly one-half at the midpoint — that is the miracle again, indifferent to n — but the steeper curve is nearly silent at a quarter-dose and nearly saturated at double, where the gentle curve is still only two-thirds of the way up. Cooperativity is why some flavors have a threshold quality, absent and then abruptly present, while others fade up smoothly. The way a trace of inosinate makes glutamate bloom — the ribonucleotide synergy that turns savory depth cooperative — is a steepness you can taste; so, in a different register, is the all-or-nothing arrival of a truly ferocious heat, which seems to wait and then detonate. The Index scores the midpoint, because the midpoint is the loudness. But our design records carry the Hill coefficient alongside it, because the midpoint is *how loud* and the coefficient is *the shape of the experience* around it — whether the heat walks toward you or jumps.

Affinity and efficacy, folded into one number

The second subtlety hides inside EC_{50} itself, and it is where potency turns genuinely subtle. A molecule's half-maximal concentration folds together two different properties: how tightly it binds the receptor — its *affinity* — and how fully a bound molecule drives the receptor toward its ceiling — its *efficacy.* A full agonist of low affinity and a partial agonist of high affinity can share an EC_{50} and feel nothing alike, the one climbing late but all the way to full response, the other arriving early but plateauing at half the height no matter how much you add. This is why the Index records the ceiling, R_max, alongside the midpoint: a flavor that tops out low is a genuinely different experience from one that keeps climbing, and loudness is not only where a curve crosses its half-point but how high it is allowed to go.

The same framework scores the opposite direction. An antagonist — a molecule that blocks a receptor rather than opening it — shifts the agonist's curve to the right when it competes for the same site, or presses its ceiling down when it binds elsewhere, and that measured shift is exactly how we quantify a bitterness-blocker: not by a loudness it adds, but by how far it pushes someone else's curve. A compound that moves a bitter agonist's EC_{50} tenfold to the right has

done a measurable, honest thing to the experience of a cup of coffee, and the Index can carry that number the same way it carries a loudness, because both are read off the same shape.

Content is not loudness

Here the two subtleties pay off in the real data, and they correct the single most common mistake about heat: that more compound means more fire. SHU — the Scoville number — is at heart a *content* assay, a measure of how much capsaicinoid a fruit holds, inherited from an era of tasting panels and sugar-water dilutions and later pinned to chromatography. IFI-H is a *receptor* assay, a measure of how loudly that capsaicinoid speaks. They correlate strongly, because more of a loud molecule is louder — but they are not the same number, and the ladder proves it in one place you might otherwise skip past.

Look at Cinder and Kindle. Our model predicts Kindle holds *more* capsaicinoid — 1.25M SHU against Cinder's 1.05M, nineteen percent more by content — and yet is the *quieter* of the two at the receptor, IFI-H 2,000 against 2,290, its predicted half-maximal concentration sitting higher (0.500 of reference against 0.437). More compound, less fire. The content ranking and the loudness ranking disagree, and both are correct about different things.

CULTIVAR	SHU (PREDICTED)	IFI-H (PREDICTED)	SHU PER IFI POINT
Cinder · SS-0088	1.05M	2,290	459
Kindle · SS-0089	1.25M	2,000	625
Phoenix · SS-0096	2.95M	4,200	700

SHU and IFI values for Scoville Splice cultivars are computational predictions.

The last column — how much capsaicinoid it costs to buy one point of loudness — is not constant across the library; it wanders from the mid-400s to the low 700s. That scatter is the whole argument. A pepper's SHU and its receptor loudness are set by two different things: SHU by *how much* capsaicinoid the plant makes, loudness by *which* capsaicinoid it is. Capsaicin forms when vanillylamine condenses with a branched-chain fatty acid at the Pun1/AT3 acyltransferase, and the length and saturation of that acyl chain decide the balance of capsaicin, dihydrocapsaicin, and the rarer species beside them — each with its own affinity for TRPV1, its own EC_{50}, its own loudness per molecule. A cultivar can be engineered to pile up total capsaicinoid and still tilt its profile toward a heavier, quieter species; another can make less and make it *louder.* SHU cannot see the difference. The receptor can, so the Index can, and that is precisely why we measure where the tongue actually listens rather than weighing what the fruit happens to hold.

The logarithm, and how to read a score

The third thing to hold is **the logarithm,** and it governs how a score should be *read* as much as how it is made. Receptor-relevant concentrations in food span an enormous range — from the micromolar whisper of a potent sweet protein to the hundreds-of-millimolar roar of table sugar, five or six factors of ten between them. No linear axis can hold that. More to the point, the body does not perceive concentration linearly: sensation grows roughly with the

logarithm of the stimulus, the old Weber–Fechner intuition, refined by Stevens into a power law in which each sensory modality carries its own exponent. Taste exponents are modest — perceived sweetness climbs a little faster than the stimulus, perceived bitterness a good deal slower — but the common thread is that doubling a stimulus never doubles the sensation; it adds a roughly constant *step.* This is Weber's law in the mouth: the just-noticeable difference is a fixed fraction of what is already present, which is why a touch of salt transforms a bland dish and does nothing to a well-seasoned one.

The EC_{50} scale is a log scale at heart, and the Index inherits that character. A score that is twice another is not twice as intense to eat; it is one meaningful step louder. We say this plainly wherever a reader might misread a ratio as a multiplier of sensation, because a number that misleads is worse than no number. It is also why the scorecard normalizes to a ceiling of 5,000 rather than running open-ended: 5,000 is +0.70 pEC_{50} above the reference, a fivefold loudness, and in practice the span from the reference at 1,000 up to that ceiling is as much dynamic range as a single axis needs to carry before perception itself has flattened. A radar that keeps climbing past the point the tongue can tell apart is drawing numbers no one can taste.

What the number cannot do

An honest unit should be as clear about its limits as its powers, so here are the ones we keep in view. An EC_{50} is measured on one receptor in a dish, and a mouth is not a dish: it is warm, moving, coated in saliva, buffered, full of fat and acid that partition a molecule away from its receptor or deliver it faster. A score is a property of the molecule-and-receptor pair, not a promise about a finished dish; the formulation still has to be cooked. The number is also a single point in time, and flavor is a shape in time — the rise, the plateau, the long decay of a heat that outlives the bite — which the midpoint alone does not capture and which is why we keep the Hill coefficient and R_max beside it. Real food is a mixture, and axes interact: a cooling compound can blunt a heat that the heat axis scored in isolation, and the Index composes those interactions deliberately rather than pretending a plate is the sum of eight solo measurements. And the taster is not generic — the bitterness receptor TAS2R38 alone comes in taster and non-taster forms, so the same cup is honestly a different loudness to different people, and a responsible score reports the receptor it was measured against, not a universal sensation.

Even the measurement carries its own arithmetic of doubt, and it is worth separating two errors that are easy to confuse. The first is the error in the *score:* because IFI is a ratio of two measured EC_{50}s, its relative uncertainty is the two relative uncertainties added in quadrature. Suppose each EC_{50} is pinned to a four-percent relative error — a reasonable day's work on a well-behaved line. Then the score inherits $\sqrt{(4^2 + 4^2)} \approx 5.7\%$, so a 4,200 would be quoted as roughly 4,200 ± 240. The second is an entirely different variance: how much a cultivar's *output* drifts from batch to batch, the manufacturing question. There our engineered constructs are designed to hold to a batch-to-batch CV at or below 7%, against

42% in wild-type — a plant that makes the same fruit twice, which is its own kind of honesty and nothing to do with the precision of the assay. The first number is how well we can read a loudness; the second is how reliably a plant is designed to keep producing it. We quote both, and never let one wear the other's confidence.

The unit, settled

So the unit is settled. Not milligrams, which the body cannot feel. Not Scoville units in their old organoleptic form, which were a panel of tasters and a lot of sugar water, and which cannot tell a louder molecule from merely a heavier one. The unit is the half-maximal effective concentration of the responsible stimulus at its responsible receptor — a physically grounded, reproducible point on a fitted curve, measured where the measurement is cleanest, carrying its own confidence interval, its own Hill coefficient, and its own ceiling. Inverted against a reference and scaled by a thousand, it becomes a score a cook can read on the same page as a chemist. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. Everything else in this book — every axis, every cultivar, every plant we have taught to speak a flavor it was never born with — is built on top of this one honest point, the place on the curve where the receptor is doing exactly half of everything it can, and telling us the truth while it does.

Chapter 3 · One Thousand Times the Reference — *the formula, derived.*

Cut a chili across its waist at our table and the inside is lacquered — a pale oil clinging to the placental ribs, the part of the fruit that the plant never meant for us. Touch it to the lip and, a breath later, a warmth that no thermometer would register climbs up through the skin. Nothing is burning. A molecule has settled into a pocket on a protein, a lock turned by a key the plant spent ten thousand years learning to cut, and the body has reported a fire that is not there. That report is the thing the Index measures. Not the oil, not the milligrams, not the plant's labor — the loudness of the report. This chapter builds the number that carries it, and we build it from three honest parts: a receptor that answers a stimulus, a curve that describes the answer, and a single point on that curve where everything interesting is decided. Assemble those three and the formula falls out almost on its own; the work, always, is earning the measurements that go into it.

The curve beneath the number

The receptor for heat is TRPV1, a non-selective cation channel standing in the membrane of the trigeminal nerve endings that lace the tongue. It is built as a homotetramer — four identical subunits arranged around a central pore, each subunit a bundle of six membrane-spanning helices, the pore itself formed where the fifth and sixth helices of all four meet. Capsaicin slips into a pocket on the inner face, between the third and fourth helices, cradled by the linker that leads to the pore. It does not pour heat into the tissue. It lowers the energy the channel needs to open, so a threshold the body normally reserves for temperatures above forty-three degrees is crossed at the temperature of a quiet kitchen. The gate opens, cations flow, the nerve fires, and we taste a warmth with no source.



FIG. 5 — The reference assay: a cell line expressing the human receptor on the microscope stage, read against a dilution series. An EC_{50} is a measurement of a controlled afternoon; an IFI score is what survives the afternoon.

Ask that receptor a question — *how much?* — by giving it the ligand at one concentration after another and reading how hard it answers, and the answers trace a curve. At very low concentration almost nothing happens; the four subunits bind cooperatively, so a little ligand barely tips the first gate. In the middle the curve rises steeply, each small increase in dose buying a large increase in response. At the top it flattens: every channel that can open has opened, and more ligand changes nothing. That S-shape is the dose-response curve, and the cleanest way to write it is the Hill relation, the fraction of maximal response as a function of dose:

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

Here θ is the response as a fraction of its own ceiling, [L] is the ligand concentration, n is the Hill coefficient that sets how steep the climb is — near two for TRPV1, because the four cooperating subunits sharpen the turn — and EC_{50} is the single point we have been circling. Set the response to exactly half, θ = ½, and the algebra collapses beautifully: $[L]^n = EC_{50}^n$, so $[L] = EC_{50}$, *whatever the value of n.* The half-maximal concentration is the one landmark on the curve that does not care how steep the curve is. Move the Hill slope, stretch or compress the middle, and the half-point stays put. That is exactly why we build the Index on it and not on the slope, not on the ceiling, not on some reading at a single arbitrary dose: EC_{50} is the honest location of the whole curve, the one number that says *where* the receptor turns over.

Everything after this rests on a single translation. A stimulus that reaches half-response at a *low* concentration is a stimulus the receptor hears *loudly* — a whisper of it is already half the shout. A stimulus that needs a *high* concentration to reach the same half is one the receptor hears *faintly.* So loudness runs inverse to the half-maximal concentration:

$$\text{loudness} \equiv 1 \div EC_{50}$$

To feel the curve's shape in numbers, read it at multiples of EC_{50} with n = 2: at a quarter of the half-point the receptor is only six percent open; at half, twenty percent; at the half-point itself, fifty; at twice, eighty; at four times, ninety-four. The tongue lives almost entirely on the shoulders of that curve, which is why a small change in the half-point moves the whole sensation — and why pinning the half-point precisely is the entire game.

A ratio, a reference, a thousand

A bare reciprocal is a tiny, unit-bearing number that means nothing alone. Is one-over-forty-micromolar large? The question has no answer until you say *compared to what.* Loudness is only ever loud or faint relative to something, so we fix that something once, per axis, and never move it again. We choose a **reference**: a stated convention, one canonical compound at its receptor, fixed by declaration and held in place — the same kind of anchor the Scoville scale uses when it fixes pure capsaicin as its ceiling. For the heat axis, capsaicin at TRPV1 (where published EC_{50} records support the

choice); for the pungency axis, AITC at TRPA1 at a stated convention constant of 2.1 µM; for each other axis, the canonical compound at its canonical receptor, each reference constant defining IFI = 1,000 on that axis by declaration alone. The reference is the axis's middle C — the note everything else is tuned against. Every candidate on that axis is scored by how its loudness compares to the reference's, which is to say by the ratio of the reference's half-point to the candidate's own:

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

Walk the algebra once and it never needs walking again. If a candidate is exactly as potent as the reference, the two half-points are equal, the ratio is one, and the score is **1,000** — the reference reads 1,000 on its own axis by construction, the fixed peg the whole scale hangs from. If the candidate is *more* potent it reaches half-response at a *lower* concentration, so its smaller EC_{50} sits in the denominator, the ratio climbs above one, and the score rises past 1,000. If it is weaker, its larger EC_{50} drags the ratio under one and the score falls below. The factor of a thousand earns nothing but convenience, and that is reason enough: it lifts the interesting range out of decimals and into round whole numbers a person can hold in the mind and say across a table — a reference at a clean 1,000, a sunlit shishito's warmth near 40, a flagship chili in the thousands, all legible on one ruler.



FIG. 4 — Phoenix (SS-0096) in the controlled chamber: the first cultivar we engineered to register on all eight axes at once, and the reference case for reading a full profile against the octagon of thousands.

The reference does something quieter and more important than set a zero: it makes the Index stable across time, which is no small promise for a standard meant to outlast any single harvest. Because every candidate is a ratio against a fixed reference measured on the same assay, drift in the instrument cancels. If a cell line reads a little hot one season, it reads the reference and the candidate hot *together,* and their ratio holds. This is the same logic that anchors pH, the decibel, and stellar magnitude — all ratio scales pinned to something fixed rather than absolute counts, all built so that you measure the thing you care about and not the weather in your laboratory. A raw EC_{50} is a measurement of an afternoon. An IFI score is a measurement of a flavor.

Three lines, then a ladder

Three worked lines make the arithmetic muscle memory. First: let an axis's reference have a half-point of 100 micromolar and a designed candidate reach half-response at 25 — four times more potent. The ratio of reference to candidate is 100 ÷ 25 = 4; times a thousand, the score is **4,000.** Second: had that candidate instead needed 400 micromolar, the ratio would be one-quarter and the score **250** — a deliberately soft axis, quieter than the reference on purpose. Third, and this is the one to carry: because the scale is a ratio, equal *steps* in loudness are equal *multipliers,* not equal additions. The jump from 1,000 to 2,000 and the jump from 2,000 to 4,000 are each a doubling of underlying potency, even though one adds a thousand and the other adds two. Read the Index the way you read a camera's stops or a musician's octaves — in ratios — and it will never mislead you.

Run that backward and a published score becomes a statement about a half-point. Since IFI = 1,000 × EC_{50}(ref) ÷ EC_{50}(candidate), the candidate's half-point is just EC_{50}(ref) × 1,000 ÷ IFI. Our flagship heat ladder, read this way, becomes a table of potencies relative to the reference — the loudness multiple is simply the score over a thousand, and the half-point ratio is a thousand over the score:

CULTIVAR · SKU	IFI-H	LOUDNESS ×REF	EC_{50} ÷ EC_{50}(REF)
Ember · SS-0087	1,570	1.57×	0.637
Kindle · SS-0089	2,000	2.00×	0.500
Cinder · SS-0088	2,290	2.29×	0.437
Blaze · SS-0091	2,630	2.63×	0.380
Wildfire · SS-0092	2,930	2.93×	0.341
Scorch · SS-0093	3,220	3.22×	0.311
Firestorm · SS-0094	3,490	3.49×	0.287
Solara · SS-0095	3,980	3.98×	0.251
Phoenix · SS-0096	4,200	4.20×	0.238

Phoenix reaches half-activation at under a quarter of the reference's half-point — a whisper does what a shout does at middle C. One rung is missing on purpose: Flare (SS-0090) is predicted at 1.48 million heat units by mass, between Kindle and Blaze, but its full receptor panel is still in assay, and we do not print a score we have not earned. The still point against which every rung is read is the profile built entirely of references, the Index's origin:

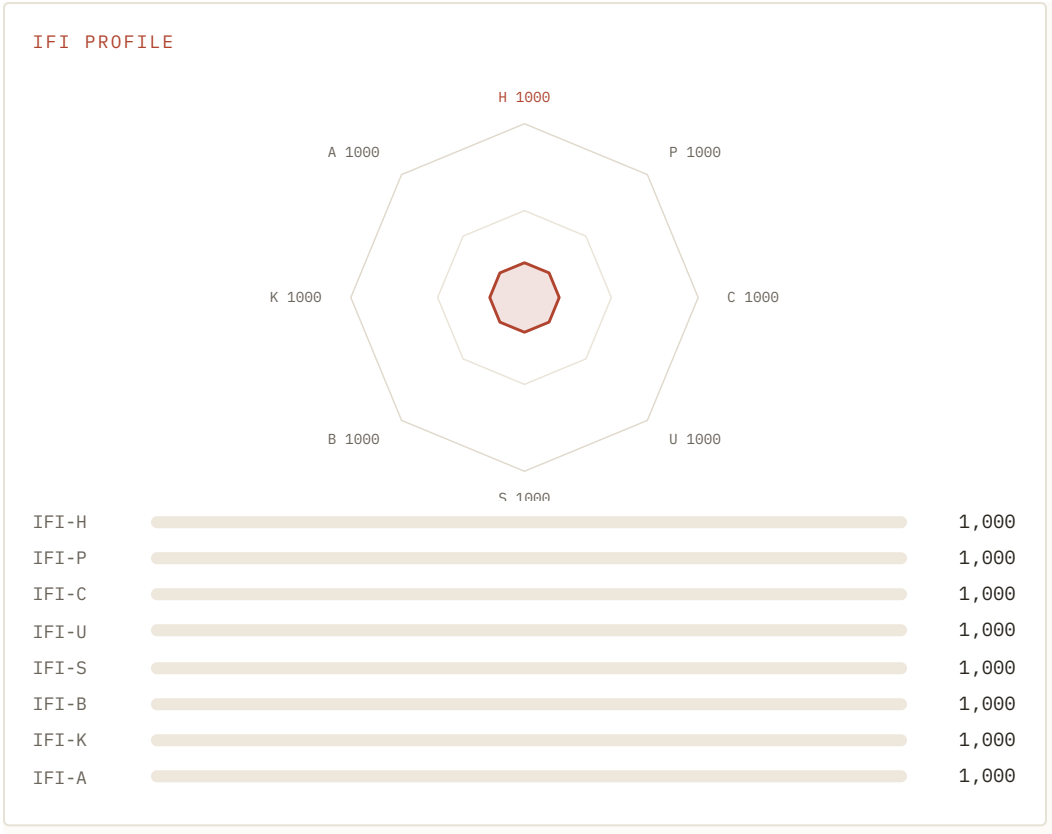

That is not a measurement of any food. It is the definition of the scale — the eight middle C's sounded together, the octagon of thousands every real profile departs from.

Units that cancel

There is a hidden gift in writing the score as a ratio, and it is worth pausing on because it is what lets one number travel across axes, organisms, and years. The formula is **dimensionless.** EC_{50}(reference) and EC_{50}(candidate) are both concentrations, so when you divide one by the other the units cancel and nothing but a pure number survives. Measure both half-points in micromolar and you get the same score as measuring both in milligrams per liter, or in mole fraction, or in any honest unit of amount — provided you use the *same* unit top and bottom. Take capsaicin, molar mass 305.4 grams per mole: a half-point of 0.70 micromolar is the very same concentration as 0.214 milligrams per liter, because one micromolar of capsaicin is 0.305 milligrams per liter. Score Phoenix's heat with reference and candidate both in micromolar, or both in milligrams per liter, and the ratio lands on 4,200 either way. The number does not know what ruler you measured with, and that indifference is exactly why it is portable: a reading in one lab, on one unit, in one year, is comparable to a reading in another, because the dimension that would make them incomparable divides itself away.

Quantity is not loudness

Here is the distinction the whole Index turns on, and the one most people arrive without. The Scoville scale measures *quantity* — how much capsaicinoid a fruit holds, anchored to pure capsaicin at a ceiling of 16,000,000 heat units, a number about *amount.* The Index measures *loudness* — how hard the receptor answers, anchored to a half-activation, a number about *sensation.* They are correlated, because more capsaicin usually does shout louder, but they are not the same number, and the places they diverge are the most honest places on the ladder.

Lay our flagships out with both scales and divide one by the other. SHU values below are computational predictions from our model, not certified measurements:

CULTIVAR · SKU	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
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

Across the upper ladder the ratio settles into a tight band near 650 to 710 heat units per Index point — the two scales marching close to in step, because at those peaks the receptor is still far from saturation and each extra molecule still buys roughly its share of loudness. Then look at the exception that teaches the rule. Cinder carries *less* total capsaicinoid than Kindle — 1.05 million heat units against 1.25 million — yet it reads *louder* on the Index, 2,290 against 2,000. More by weight, quieter at the receptor. The reason lives in composition: a fruit's heat is never pure capsaicin but a blend with dihydrocapsaicin and nordihydrocapsaicin, and those cousins, molecule for molecule, do not turn the TRPV1 gate quite as hard. Two fruits can hold the same milligrams and speak at different volumes if their blends differ. The dilution assay counts the molecules; the receptor assay asks what they *do.* The Index was built to report the second thing, because the second thing is what the tongue is told.

This is also where the plant's own chemistry meets the formula. Capsaicin is assembled at the Pun1/AT3 acyltransferase, which condenses vanillylamine with a branched-chain acyl group to make the molecule, and the two substitutions we wrote into that enzyme — S39L and L345G, with stabilizing free-energy shifts of −2.1 and −1.8 — our model predicts peak capsaicin output of approximately 2,937,948 heat units, with a design batch variation specification of seven percent or less, against forty-two percent in the wild-type plant. But notice what those edits move and what they do not. They move *quantity:* how much ligand the fruit makes. The Index then reads *loudness:* how the fruit's whole extract speaks at the receptor. The enzyme work fills the denominator; the assay measures what filling it bought. Quantity and loudness, the plant's labor and the tongue's verdict, meeting at a single half-point.

The axis that runs backward

One axis needs a word of its own arithmetic, because its whole purpose is subtraction. Zero-Bitter does not add a loudness; it makes a receptor — TAS2R38, the bitter sensor — *quieter.* So its score is built not from a stimulus's potency but from an *attenuation:* how far a design lifts the apparent half-point of the bitter ligand, so that more of the bitter compound is needed to reach the same half-response. The formula does not change; what changes is which

measurement goes where. For the seven loudness axes the candidate's small half-point sits in the denominator, so louder means higher. For Zero-Bitter the roles swap: the *attenuated* system's raised half-point goes in the numerator, measured against the untreated food's baseline, so *quieter* bitterness means higher. Subtraction, scored exactly like addition, landing on the same ruler without an asterisk.

Walk the numbers and the axis reveals how reasonable it is. Write the lift factor as m — the multiple by which the design raises the bitter half-point — and the score is simply IFI-B = 1,000 × m. Now ask what a given lift *feels* like. With a Hill slope of one, and probing at the baseline half-point itself, a bitter response that was fifty percent of maximum falls to 1 ÷ (m + 1) of maximum, a reduction of 1 − 2 ÷ (m + 1). Run the band the design space actually spans: a twenty-percent reduction in felt bitterness needs a lift of 1.5× — an IFI-B of 1,500; a sixty-five-percent reduction needs a lift of 4.7× — an IFI-B near 4,700. So the twenty-to-sixty-five-percent span of designed attenuation our model predicts maps cleanly onto Index scores from roughly 1,500 to 4,700, and a design reading 2,200 — a 2.2× lift — is cutting felt bitterness by about thirty-eight percent at that dose. The caveat is honest and small: the exact percentage shifts with dose and with the Hill slope, so we quote the ratio, which is fixed, and let the percentage be the ratio's shadow at a stated dose. The negative-space axis reads beside the other seven, in the same units, by the same formula, turned around.

The ceiling, and the discipline of five thousand

Two design choices finish the instrument. The first is a per-axis ceiling. Receptors saturate; past some loudness the body simply cannot tell two stimuli apart, and a score beyond that point would be a distinction without a difference — false precision, the one dishonesty we most want to refuse. So we cap each axis at a normalization maximum of **5,000,** five times the reference. Across the entire designed range we can presently build toward and the entire natural range beneath it, every distinction that matters lives between the baseline thousand and roughly five times it; push past that and two stimuli differ by less than a just-noticeable difference, so a finer number would be claiming a resolution the eater cannot feel. Five thousand is the honest dynamic range of the instrument — wide enough to hold Phoenix's heat with real headroom to spare, narrow enough that every step inside it corresponds to something a person could actually taste. The cap is a discipline on us as much as a promise to the reader: it says that beyond a point, more loudness is not more flavor, and the work turns from pushing a single axis to composing the whole. A scale should promise only the resolution it can keep, and ours keeps every point it prints.

Every number with its error bar

Because every score is a ratio of two measured half-points, every score carries a knowable uncertainty — and we propagate it rather than hide it. The arithmetic is cleanest in log space, where a ratio becomes a difference and the two uncertainties simply add. Take the natural log of the formula: $\ln(\text{IFI}) = \ln(1{,}000) + \ln EC_{50}(\text{ref}) - \ln EC_{50}(\text{candidate})$. The constant has no error; the two measured terms

do. For small relative errors the coefficient of variation of a quantity is close to the standard deviation of its logarithm, so the variations add in quadrature:

```
CV(IFI) ≈ √[ CV(reference)² + CV(candidate)² ]
```

Put our real numbers in. The reference is measured to about two percent; a designed cultivar's CV specification is seven percent or better. Then CV(IFI) = √(0.02² + 0.07²) = √0.0053 ≈ 7.3 percent. On Phoenix's heat score of 4,200 that is a one-sigma spread near ±310. The ninety-five-percent interval is built by multiplying and dividing in log space — 4,200 × e^(±1.96 × 0.073) — which gives roughly 3,640 to 4,850. Notice it is *asymmetric* in linear space, reaching a little further up than down; that is not an error, it is the honest signature of a ratio scale, where symmetric uncertainty in log space is lopsided in the raw number. We report the interval as it truly is rather than rounding it into a tidy, false symmetry.

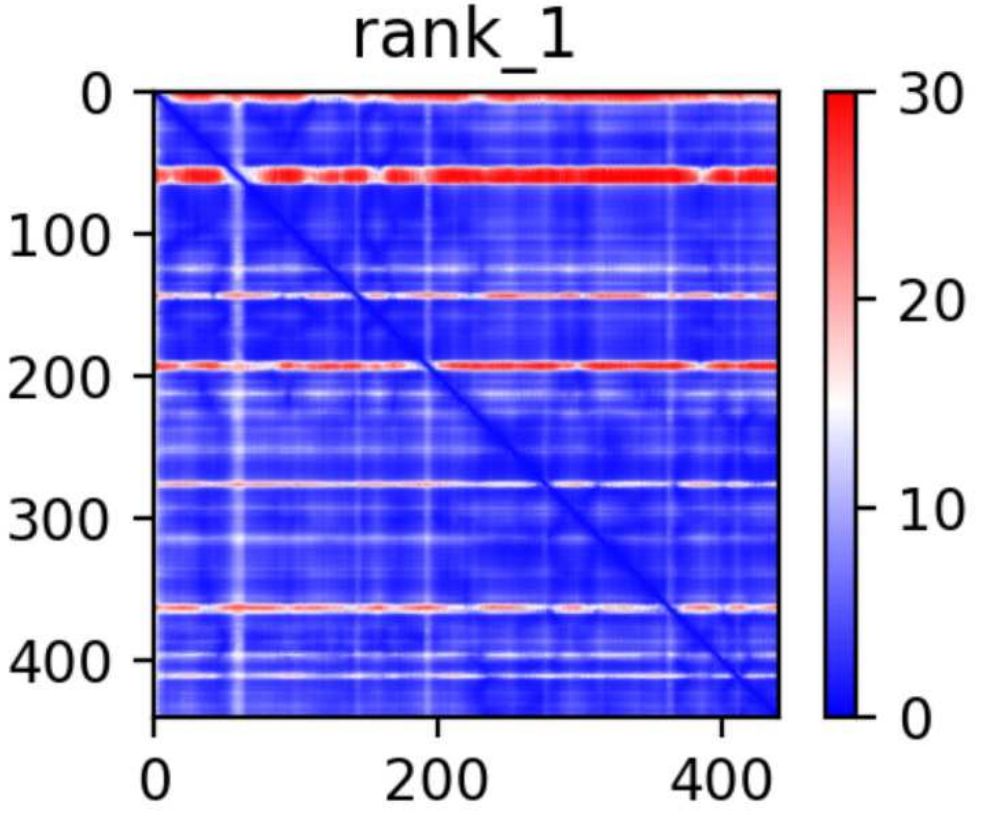

FIG. 6 — Predicted aligned error across the Pun1/AT3 model: a map of where the structure is trusted and where it is not. A published profile carries the same honesty — never a bare number, always a number with the width of its confidence drawn around it.

This is where reproducibility stops being a virtue and becomes the whole argument. Feed the same formula a candidate that varies by forty-two percent — the wild-type plant's batch-to-batch swing — and the propagated uncertainty is √(0.02² + 0.42²) ≈ 42 percent. A score of 1,000 would then arrive with a ninety-five-percent interval running from under 1,850 to over 9,500: a number so wide it tells you

almost nothing. The error bar is small *only because the designed cultivars are so reproducible,* and that is precisely why the seven-percent variation matters beyond manufacturing. It is what makes the published score narrow enough to mean something. A standard that reported its numbers without their uncertainty would be asking to be *believed* rather than checked, and across a library this size, belief is the one thing we refuse to ask for.

Where the formula breaks

A number is only as honest as its stated limits, so here are ours. The Index scores *potency* — the half-point, where the curve turns — and not *efficacy,* the height of the ceiling. For a full agonist, one that drives the receptor to the same maximum as the reference, that is the whole story. For a *partial* agonist, one that saturates at a lower ceiling no matter the dose, a single potency number flatters the sensation: the stimulus may reach its own half early yet never shout as loud, and the lone score would overstate it. The standard handles this by requiring full-agonist behavior at the reference receptor, and by flagging efficacy separately where a candidate falls short, rather than letting one number quietly carry two meanings.

Three more edges deserve naming. First, desensitization: TRPV1, pushed hard by capsaicin, grows numb — the channel desensitizes, so past a point *more* ligand yields *less* response, and the comfortable assumption that louder always means higher simply fails. The ceiling at 5,000 guards the top of the scale against exactly this, refusing to score into the region where the receptor has stopped keeping faith with the dose. Second, the receptor is not one fixed thing: TRPV1 differs across species — famously, birds scarcely feel capsaicin at all — and it is itself thermosensitive, its threshold sliding with temperature. So the reference assay fixes the human receptor at a controlled temperature; a half-point measured warm is a different number from one measured cool, which is why "a measurement of an afternoon" is a warning and not a flourish. Third, real food is a mixture: a fruit's extract is capsaicin *and* its cousins together, and the half-point we read belongs to the blend, not to any single molecule. We treat that as a feature, because the blend is what the eater meets — but a reader who wants the pure-compound potency must go to the compound, not the fruit. None of these breaks the formula. They mark the ground where it must be used with care, and naming them is part of keeping the number worth trusting.

Eight instruments, one ruler

The last choice is the deepest: the axes are independent, and therefore additive in *meaning,* never in *magnitude.* Each of the eight is its own receptor, its own reference, its own curve — heat at TRPV1, sweetness at TAS1R2·R3, umami at mGluR1, Zero-Bitter at TAS2R38, kokumi at CaSR, cool-burn at TRPM8, pungency at TRPA1, astringency at the salivary proline-rich proteins. A heat score and a sweetness score are not two amounts of one substance to be summed into a "total flavor." They are two readings on two instruments, like temperature and humidity: you would never add them, and you would lose everything by trying. Because each axis rests on its own stated reference convention — not a single absolute

measure of sensation — a score of 1,000 on the H axis and a score of 1,000 on the P axis are both their instrument's defined reference point, not the same absolute loudness. This is why a profile is drawn as an eight-spoke radar and never as a single bar. You do not add the axes. You read the *shape.*

Read Phoenix as a shape, and the full profile resolves into eight independent readings, each a receptor, a score, and a half-point ratio:

AXIS	RECEPTOR	IFI	EC_{50} ÷ EC_{50}(REF)	WHAT IT SAYS
H · Heat	TRPV1	4,200	0.238	4.20× louder than the reference
B · Zero-Bitter	TAS2R38	2,200	2.200 †	half-bitter dose raised 2.2× (≈38% less)
K · Kokumi	CaSR	1,900	0.526	1.90× — body and continuity
U · Umami	mGluR1	1,800	0.556	1.80× — savory depth
A · Astringency	Salivary PRPs	1,400	0.714	1.40× — a light tannic grip
P · Pungency	TRPA1	252	3.967	0.25× — quieter than the reference
S · Sweetness	TAS1R2·R3	1,100	0.909	1.10× — barely sweet
C · Cool-Burn	TRPM8	800	1.250	0.80× — deliberately kept quiet

† Zero-Bitter inverts the ratio: the attenuated system's raised half-point sits in the numerator, so 2,200 means the bitter half-dose is lifted 2.2-fold.

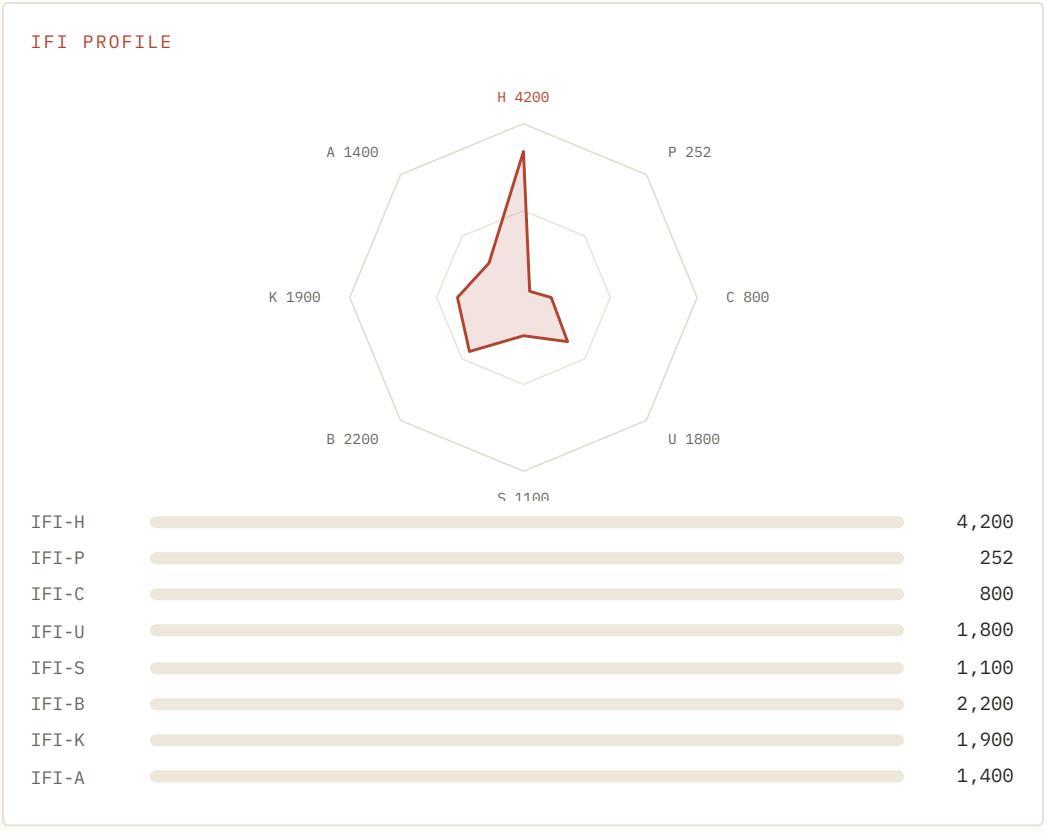

Sound those eight together and you do not hear a sum; you hear a chord — a loud heat over a warm savory body, bitterness pulled back, sweetness and cool held low on purpose, a tannic edge to finish. The Index's real gift is that it makes that chord a thing you can see, compare, remember, and — this was always the point — *specify in advance and then grow.* A flavor becomes a silhouette across eight honest dimensions, each anchored to a receptor, each scored by one ratio, each carrying the width of its own confidence. Everything elaborate in the rest of this work — the receptor biology,

the biosynthetic pathways, the protein models, the long reverence for the plant that carries it all — exists to move a single half-point into the denominator on purpose, and to move it so reproducibly that the number we print is a number you could check. The formula is the easy part. Earning the number is the work, and the work is the whole of our table. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 4 · The Radar and the Bar — *how to read a profile.*

A number alone is a fact; eight numbers arranged well are a portrait. The scorecard is the Index's signature object and its teaching tool, and by the end of this chapter you will read one at a glance — for the rest of this book, and, we hope, for the rest of your eating. We built it the way we build everything at our table: with reverence for the thing on the plate, and no patience at all for a number we cannot defend.

A profile is written as eight labeled values — the lab-notebook layer, monospace and terse — one per axis, in a fixed order so the eye learns the positions the way it learns a clock face. We give the order a mnemonic shape and never vary it: **H, P, C** on the warm and sharp side (Heat, Pungency, Cool-Burn — the trigeminal channels, the sensations that are half touch); then **U, S, B** through the rounded tastes carried by G-protein-coupled receptors (Umami, Sweetness, Zero-Bitter); then **K, A** to close (Kokumi, Astringency — the two felt as much as tasted, richness and friction). Read top to bottom it walks from fire to texture; drawn as a radar it closes into an eight-pointed silhouette, each axis a spoke whose length is its score and whose reference mark sits on the thousand ring. Here is the profile every reader of this library should hold in memory — Phoenix, our full-spectrum proof, the single cultivar we designed gene by gene to reach all eight axes at once:

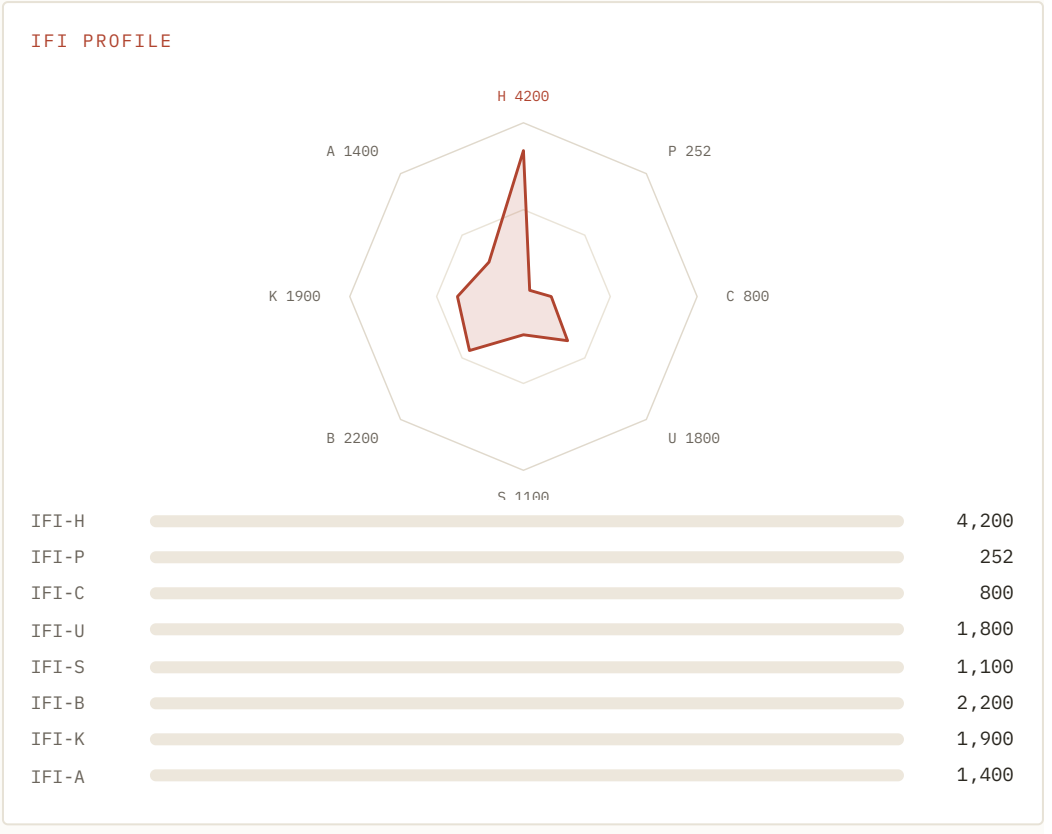



A scorecard page from the design archive — reading a profile is reading a shape, not a sum.

Read it as a shape. The spike is Heat at 4,200 — well above the reference thousand, a true superhot, though short of the 5,000 ceiling, so there is honest headroom left even here. Zero-Bitter at 2,200 and Kokumi at 1,900 do quiet, structural work — bitterness pushed down, richness and continuity pulled up — while Pungency at 252 and Cool-Burn at 800 both sit *below* the reference, the AITC pungency channel held deliberately soft, the cool axis a faint warming lean rather than a cool one. Those two sub-thousand values are the tell of a designed flavor. Nature rarely builds a fruit that is loud on six axes and purposefully soft on two; a composer does, because the composer decided those axes should recede so the others could be heard. The profile is not a scorecard of achievement. It is a statement of intent, made legible — and like any honest statement, it is only worth reading because every digit in it can be derived, defended, and reproduced.

The thousand-ring, and the formula that sets it

Before the shape can mean anything, the ring has to. Every axis is scored by one relation, the same for heat and for sweetness and for the tannic grip of a good tea:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of its full response — the half-maximal point, the honest middle of the curve. The reference constant on each axis is defined by convention — a fixed anchor that places the native organism at exactly 1,000 by definition, stated the way Scoville's dilution-to-threshold definition is stated, not derived from any single measurement. Everything else is a ratio against it. When a candidate reaches half-activation at the *same* concentration as the reference, its EC_{50} equals the reference's, the ratio is one, and the score is 1,000 — it lands precisely on the ring. When the candidate speaks louder — reaches the same half-max at a *lower* concentration, a smaller EC_{50} — the ratio climbs above one and the spoke pushes past the ring. When it is softer, needing more of itself to say the same thing, the spoke falls inside. The thousand-ring is not an

average or a passing grade. It is parity with the plant as it was before we touched it, and the most interesting cultivars are the ones that cross it in both directions on purpose.

Read Phoenix this way and the eight digits decode cleanly. Heat at 4,200 means the ratio EC_{50}(reference) ÷ EC_{50}(Phoenix) equals 4.2: the designed capsaicinoid reaches half-activation of TRPV1 at one-fourth-and-a-bit the concentration the reference needs — it is 4.2 times louder at the channel. Turn the formula around and you get the concentration directly. The fraction of the reference's EC_{50} at which the candidate hits half-max is simply 1,000 ÷ IFI, so:

AXIS	RECEPTOR	IFI (PHOENIX)	LOUDNESS (× REFERENCE)	EC_{50} (× REFERENCE)
IFI-H · Heat	TRPV1	4,200	4.20×	0.238
IFI-P · Pungency	TRPA1	252	0.252×	3.968
IFI-C · Cool-Burn	TRPM8	800	0.80×	1.250
IFI-U · Umami	mGluR1 / TAS1R1·R3	1,800	1.80×	0.556
IFI-S · Sweetness	TAS1R2·R3	1,100	1.10×	0.909
IFI-B · Zero-Bitter	TAS2R38	2,200	2.20×	0.455
IFI-K · Kokumi	CaSR	1,900	1.90×	0.526
IFI-A · Astringency	salivary PRPs	1,400	1.40×	0.714

Look at Cool-Burn: an EC_{50} of 1.25 times the reference, meaning it takes a *quarter more* stimulus to reach the same half-cool — the axis the design pushed softly inward, confirmed in a single fraction greater than one. Every axis above the ring returns a fraction below one; axes below the ring return fractions above it. The arithmetic and the eye agree, which is the whole point of anchoring to a receptor rather than to a committee's taste.

One honesty is owed about that table. Six of the axes are activations — the stimulus turns a receptor *on*, and EC_{50} is a half-activation concentration in the ordinary sense. The others are quieter machinery. Zero-Bitter is attenuation at TAS2R38, a suppression of the bitter alarm by somewhere between a fifth and two-thirds of its signal, so its "EC_{50}" is really the half-*effect* concentration of an inhibition. Kokumi is modulation at the calcium-sensing receptor, CaSR, which does not so much fire as deepen what its neighbors are already doing. Astringency is barely a receptor event at all — it is tannins and polyphenols precipitating the proline-rich proteins that lubricate the mouth, felt as friction when the lubrication fails. The formula absorbs all of them because it only ever asks one question — *what concentration produces half the full effect* — and that question has an answer whether the effect is a channel opening, a receptor quieting, or a protein falling out of solution. But a careful reader keeps the distinction in mind: a loud Zero-Bitter spoke is loud *silence*, not loud sound.

What a single axis actually measures

To trust the ring you have to know where the EC_{50} under it comes from, and it comes from a curve. Expose the receptor to a rising ladder of concentrations, measure the fraction of full response at each step, and the points fall onto a sigmoid — the dose–response curve — that the Hill equation describes exactly:

$$R([S]) = [S]^n \div (EC_{50}^n + [S]^n)$$

Here R is the response as a fraction of maximum, [S] is the stimulus concentration, and n is the Hill coefficient, the steepness set by how cooperatively the receptor's subunits act. For capsaicin at TRPV1 the channel opens cooperatively, with n near 2; the sweet and umami heterodimers and the bitter family each carry their own slope. Set [S] equal to EC_{50} and the equation collapses to one-half — that is the definition of the half-maximal point — but the deeper reason we anchor there is that the curve is *steepest* there, and steepness is information.



The receptor bench where a profile is earned — every axis is a dose–response curve before it is a number.

Measure the slope of response against the natural logarithm of concentration and it peaks exactly at the half-max, where it equals n ÷ 4 per e-fold of stimulus — for a Hill coefficient of 2, half a response-unit per natural-log step, the most the curve ever moves for a given nudge in dose. Away from the middle the sigmoid flattens: near threshold almost nothing happens as you add stimulus, and near saturation almost nothing happens either, because the receptors are already nearly all engaged. A measurement taken on a flat stretch is a measurement with enormous bars, because a large change in concentration maps to a tiny, noisy change in response. A measurement taken at the half-max is the tightest the biology allows. So EC_{50} is not an arbitrary convention; it is the one point on the curve where the plant answers most clearly, and we built the Index to listen

exactly there. This is also why the scorecard caps each axis at 5,000. Five times the reference loudness drives the receptor so far up the saturating shoulder of its curve that still more potency buys almost no additional response — the tongue has stopped distinguishing, and a number past the cap would be precision we cannot honestly claim. We let the axis run to the edge of meaning and no further.

Why loudness is a ratio — and why the tongue agrees

It would have been simpler to report concentration — how much capsaicin, how much glutamate — and some labels do. We refused, because the tongue does not count molecules; it hears them, and hearing runs on ratios. Press a little harder, add a little salt, raise a little heat, and the increment you notice is not a fixed amount but a fixed *fraction* of what is already there — the Weber relation, the oldest law in the study of the senses. The receptor's own curve carries that logarithm built in: equal multiplicative steps in concentration are equal additive steps along the log-dose axis, which is the axis the sigmoid lives on. By scoring every axis as a ratio of EC_{50}s, the Index speaks in the same units the nervous system does.

This is why a 2,000 is not "twice as flavorful" as a 1,000. It is one ratio-step louder — a doubling of potency at the receptor, which the tongue registers as a constant increment of sensation, not a doubling of it. The move from 1,000 to 2,000 and the move from 2,000 to 4,000 feel like the same size of step precisely because they are the same multiplicative distance, and the perceptual world is multiplicative. Hold onto this when you read Phoenix's heat at 4,200: it is not four-and-a-fifth times as *fierce* as the reference on the tongue. It is a little over two doublings of potency above it — loud, climbing, but bounded, and the fierceness grows far more gently than the bare ratio suggests. The ratio is the honest physical fact; the logarithm is why the fact feels humane.

Reading the shape: the geometry of the octagon

There is a reason the radar is a *shape* and not a bar chart, and it is worth a moment of geometry. Eight independent axes are eight dimensions, and no eye holds eight numbers at once — but any eye holds a *silhouette* effortlessly, the way it holds a face. Spoke the eight scores out from a common center at equal angles, forty-five degrees apart, join their tips, and eight digits become one closed figure whose outline is the flavor's identity and whose enclosed area is a rough sense of its total loudness. That area has a closed form. For tips at radii r_i set at equal angles, the octagon's area is half the sine of the step angle times the cyclic sum of neighboring products:

$$A = \tfrac{1}{2} \cdot \sin 45^\circ \cdot \Sigma \left(r_i \cdot r_{i+1} \right)$$

Run Phoenix through it in scorecard order and the neighboring products — Heat times Pungency, Pungency times Cool-Burn, around to Astringency times Heat — sum to 19,820,000; multiply by $\tfrac{1}{2} \cdot \sin 45^\circ \approx 0.3536$ and the figure encloses roughly 7.0 million

square-units of loudness. A larger octagon does, loosely, mean a bigger flavor: more total receptor traffic, more happening at once. The eye is right to read a big shape as a big meal.

But the area hides a trap, and we would rather hand you the trap than let it spring. That cyclic sum is built from *adjacent* products, so the area depends on the *order* of the axes. Permute the spokes and the same eight numbers enclose a different area — a mathematical fact, not a flavor fact. It is exactly why the axis order is fixed and never varied: only a fixed order makes two shapes comparable at a glance. And it is why, when we want a single honest magnitude, we do not trust the drawn area at all. We use the root-mean-square of the eight values, which depends on no ordering whatsoever:

```
RMS = √( ( H² + P² + C² + U² + S² + B² + K² + A² ) ÷ 8 )
```

For Phoenix the sum of squares is 33,203,504, so the RMS loudness is √(33,203,504 ÷ 8) ≈ 2,037, against a plain arithmetic mean of 1,707. The RMS always meets or exceeds the mean, and the gap between them measures how *peaked* the profile is. Their ratio for Phoenix is 2,037 ÷ 1,707 ≈ 1.19 — close to one, the signature of a broad, balanced octagon. Picture instead a pure specialist, one axis near 4,000 and the other seven near 100: its mean falls to about 588 while its RMS holds near 1,417, a ratio of 2.4. That ratio — RMS over mean — is the number behind the word *balanced*. Near 1, a full-spectrum face; climbing past 2, a single violent spike propped on silent neighbors. Use the drawn area to feel magnitude; use the RMS to state it; use their ratio to tell a composition from a specialist without ever naming a digit.

Two silhouettes train the rest. Picture one tall spoke at umami, a shorter one beside it at kokumi, and the other six hugging the ring: that narrow, two-pronged figure is a broth, a dashi, a long-simmered stock — savory depth and continuity and almost nothing else, and you would know it for comfort from across a room. It is also the shape our own vegetarian table leans on hardest, where umami and kokumi do the work that meat is usually asked to do. Now picture the near-opposite: a broad, nearly even octagon, every spoke past the ring, one of them — heat — reaching out past the rest. That is Phoenix's own shape, and its roundness is the tell that someone composed all eight on purpose rather than letting one run away. The specialist is a spike; the composition is a shape; and once you have read a dozen, you stop reading digits and start recognizing faces.

The bar beneath the radar: Scoville is not loudness

Under the radar sits a single horizontal gauge — the heat-bar, IFI-H alone, the one axis the world already thinks it knows. We keep it because heat is the number people arrive with, and we put it beneath the octagon precisely so that heat can be both a spoke among eight and a headline of its own. But the bar carries a quiet correction, and reading it teaches the difference between measuring a chemical and measuring a sensation.

Scoville Heat Units count capsaicinoid concentration — how much of the molecule is present. IFI-H counts loudness — how strongly what is present drives the channel. They track each other, but they are not the same number, and where they part ways the whole case for the Index lives:

CULTIVAR	SKU	SHU (PREDICTED)	IFI-H	AXES
Ember	SS-0087	866K	1,570	2 of 8
Cinder	SS-0088	1.05M	2,290	3 of 8
Kindle	SS-0089	1.25M	2,000	3 of 8
Flare	SS-0090	1.48M	—	4 of 8
Blaze	SS-0091	1.72M	2,630	4 of 8
Wildfire	SS-0092	1.96M	2,930	5 of 8
Scorch	SS-0093	2.21M	3,220	5 of 8
Firestorm	SS-0094	2.47M	3,490	6 of 8
Solara	SS-0095	2.72M	3,980	6 of 8
Phoenix	**SS-0096**	**2.95M**	**4,200**	**8 of 8**

SHU and IFI-H scores are computational predictions from our flavor model, not certified measurements.

Walk down the predicted SHU column and it rises without a stumble, from about 866,000 to nearly three million. Walk down the IFI-H column beside it and it stumbles once, hard. Our model puts Kindle at a predicted 1.25 million Scoville units — nineteen percent *more* capsaicinoid by mass than Cinder's predicted 1.05 million — yet its heat scores 2,000 against Cinder's 2,290, about thirteen percent *softer* at the receptor. More molecule, less loudness. The two measures cross.

They cross because heat has a *shape*, not just a size. Capsaicin and dihydrocapsaicin — the two dominant capsaicinoids — reach TRPV1 with different half-max concentrations and different onsets, one a faster, brighter strike, the other a slower, deeper warmth. A fruit richer in total capsaicinoid can be tuned toward the slower molecule and read softer at the channel than a leaner fruit tuned toward the faster one. SHU weighs the pile; IFI-H listens to it. This is not a flaw in either number — it is the reason both appear on the card. The bar tells you how much fire there is; the spoke tells you how loudly it speaks; and the gap between them, on a line like Kindle's, is where a flavor's character actually lives.



Phoenix (SS-0096) in the controlled chamber — the only cultivar on the ladder whose shape closes on all eight spokes.

The error bars you cannot see

Every spoke on a scorecard is drawn as a clean point, and every clean point is a small lie of omission, because a measurement is never a point — it is a point and a band. The band is set by reproducibility, and reproducibility is where designed flavor earns its keep. Our batch-to-batch coefficient of variation is specified at seven percent or below; the wild-type reference the whole Index is anchored against runs near forty-two. That difference is not bookkeeping. It decides whether a shape can be read at all.

Take Phoenix's designed heat at 4,200. At a seven-percent CV, the designed score carries a specified reproducibility band of about ±294 — call it 3,906 to 4,494 — a spoke tip you could cover with a thumbnail. Hold the same 4,200 to the wild-type's forty-two-percent spread and the band blows out to ±1,764, from 2,436 all the way to 5,964: a spoke that smears from "serious superhot" nearly to the ceiling and down past the midpoint of the chart. You cannot compose with a smear. The entire practice of reading one spoke against another — of trusting that Kokumi at 1,900 sits meaningfully above Astringency at 1,400 — depends on the bands being narrow enough that the points don't overlap, and seven percent is what makes the octagon a figure instead of a cloud.

PHOENIX · IFI-H 4,200	DESIGNED · CV ≤7%	WILD-TYPE · CV ~42%
One-sigma band (±)	294	1,764
Lower bound	3,906	2,436
Upper bound	4,494	5,964
Spread as % of ceiling	11.8%	70.6%

The last line is the plainest way to feel it: at our design precision the specified range for that spoke covers about an eighth of the whole 5,000 chart; at wild-type precision it covers more than two-thirds of it. Same mark on the paper, two different instruments — one a scalpel, one a rumor.

That same CV sets the scorecard's resolution, the smallest difference worth believing. Two axes near 2,000 each carry a one-sigma band of about 140; the uncertainty in their *difference* is the two added in quadrature, roughly $140 \times \sqrt{2} \approx 198$, so a gap smaller than about 200 is within a single standard deviation of noise, and only a gap past roughly 400 clears the two-sigma bar where we

would call the two axes honestly distinct. Read a card knowing this and you stop over-interpreting a hundred-point wobble between neighboring spokes; it is texture in the measurement, not intent in the design. And there is a mercy folded into the arithmetic: the psychophysics literature puts the tongue's own just-noticeable step for intensities like these near a fifth of the stimulus — a Weber fraction far coarser than our seven percent. The design precision, jar to jar, is specified tighter than a human mouth is able to tell apart. We design the manufacture to exceed what the sense will need to judge it, which is a quiet form of respect for the person at the table.

Three habits, two refusals

Three reading habits serve for the whole library. **First, read against the thousand.** The reference octagon — the ring where every axis would land if the candidate exactly matched the native plant — is the ghost behind every radar; a spoke past 1,000 is louder than that plant, one short of it softer, and the designed cultivars are the ones that do both deliberately. **Second, read the silhouette before the digits.** A broad, balanced octagon is a full-spectrum flavor; a single dominant spike with quiet neighbors is a specialist; two or three strong spokes with the rest low is a composition — a swicy heat-and-sweet, a savory umami-and-kokumi depth. The gestalt carries the meaning; the digits confirm it. **Third, respect the axis count.** Phoenix is extraordinary because it registers on all eight at once; most cultivars, by design and by difficulty, speak only two or three axes loudly and leave the rest near the baseline. A profile that claims eight strong axes is either a landmark or a mistake, and the Index is built so you can tell which.

Two misreadings are worth naming so you can refuse them. The first is to *add* the axes — to see six lifted spokes and conclude the flavor is "more" than a three-spoke one. The axes do not sum; loudness lives in spoke lengths, not spoke counts, and a food loud on three pathways can hit far harder than one mildly lifted on six. The RMS makes this precise: a profile with three spokes at 3,000 and five at 300 carries a higher root-mean-square loudness than one with all eight at 1,500, even though the second looks busier. The second misreading is the one the logarithm already warned against — reading the ratio as a multiplier of sensation, taking a 2,000 for twice the feeling of a 1,000 when it is one ratio-step louder and the tongue hears in constant increments. Read spoke length for loudness, spoke count for breadth, and the whole shape for identity — and never do arithmetic the axes were not built to support.

The two readings above the eight

Two composite readings sit above the eight, and the scorecard carries them where they apply. **IFI-F, Full-Spectrum,** is not a ninth axis but a verdict: it marks a flavor that registers meaningfully across all eight pathways at once — the regime of our 300-strong full-integration catalog and of nanoparticle co-delivery, and the rarest thing in the library, because balancing eight competing pathways in one living fruit is harder than maximizing any one of them. **IFI-M, Multi-axis,** names the deliberate two- and three-axis compositions

— the dual designs that reach as high as 8,400 in combined loudness and the triple designs balanced near 2,800, where the engineering is not any single score but the *relationship* between scores. A full card, then, is read on three levels at once: the eight axis values, the composite verdict above them, and the silhouette that ties them into one recognizable face. Three levels, one glance, once you have the habit.

Why we labor over the scorecard

We labor over this object because it is, finally, a teaching tool as much as a record. Most people have never been given permission to taste analytically — to notice that the thing they are enjoying is six things at once and not one — and the radar hands them that permission in a glance, naming the parts without diminishing the whole. A chef reads it to reverse-engineer a dish; a formulator reads it to match a target to within the error bar; a curious eater reads it and suddenly has words for why one tomato sings and another is mute. The same eight-spoke figure serves the laboratory bench and the kitchen table, which is exactly what a shared standard is for: one picture to point at, one language to argue in, open to expert and amateur alike.

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

That is the instrument, whole. A receptor that reports occupancy; a dose–response curve whose half-maximal point is the one place the plant answers clearly; a formula that fixes the native reference at a thousand and scores everything as a ratio against it, in the logarithmic units the tongue already uses, capped where the tongue stops distinguishing; a geometry that turns eight honest numbers into one readable face, with the root-mean-square beneath it when the eye needs correcting; and a seven-percent design specification under every spoke that keeps the face from blurring. Hold those together and a scorecard stops being a chart and becomes what we meant it to be — a way to read, and then to compose, the flavor and the nourishment we want a plant to carry. The rest of this library simply takes the eight axes one at a time, until each abstraction in that picture puts on flesh: a channel, a molecule, a pathway, and the plants we taught to speak each one.



PART II OF V

The Eight Axes

TRPV1 · TAS1R2·R3 · mGluR1 · TAS2R38 · CaSR · TRPM8 · TRPA1 · salivary PRPs

Chapter 5 · Heat — *the axis everyone knows.*

Heat is where we began, as a company and as a family, and it is the axis against which the whole Index earned its right to exist. It is also the purest case of the principle that sits under everything we measure — that the tongue reports receptor activation, not chemistry — because heat is an illusion with no exception in it, a fire with no flame, and the body falls for it completely, every time, with a sincerity that never wears off. That faithfulness moves us. The body is not fooled so much as honest: hand it the right molecule and it will tell you, truthfully, about a warmth that is not in the room. Everything in this chapter is an attempt to earn the right to design that moment — to write a precise number onto a living plant and have the warmth arrive, on cue, grown rather than waited for. We will take the long way, because heat rewards it: from the single channel that carries the whole sensation, through the enzyme that builds the molecule, to the ten designed cultivars we drew from the library to show that a flavor could be computed, targeted, and grown.

The alarm we learned to play

Capsaicin binds **TRPV1**, a non-selective cation channel standing in the membranes of the trigeminal nerve endings that wind through the tongue. TRPV1 is not, in its day job, a taste receptor at all; it is a thermometer and an alarm. It opens to temperatures above about 43 °C, and to protons — the acidity of injured tissue — and when it opens, calcium and sodium rush in, the neuron depolarizes and fires, and the brain receives the one message TRPV1 ever sends: *this is hot, this may harm you.* There is no "spicy" signal in the body. There is only this borrowed heat-and-harm line, played out of turn. And it does not even travel the taste nerves: heat rides the trigeminal, the face's own touch-and-temperature nerve, not the gustatory fibers that carry sweet and bitter. That is why chili "taste" is felt as much in the lips, the gums, and the throat as on the tongue, why it can sting an eye or a cut that has no taste buds at all, and why — strictly — heat is not a taste. It is a touch the brain files under flavor.

We know the channel's shape in atomic detail, because TRPV1 was the first of these to be solved by cryo-electron microscopy: four identical subunits arranged around a central pore, each subunit a bundle of six membrane-spanning helices, with a vanilloid-binding pocket tucked against the pore between the third and fourth of them. Capsaicin slips into that pocket from the channel's *intracellular* side — the oily molecule having already crossed into the membrane — and pulls on the short linker that gates the pore. The effect is not to add energy but to move a threshold. A channel that normally needs scalding 43 °C to open now opens at something well below body heat, so that your own steady 37 °C is suddenly enough to trip the alarm it was holding in reserve. The molecule does not warm the tissue by a single degree. The warmth you feel is real nerve traffic reporting a fire that is not there — the cleanest demonstration in all of flavor that what you taste is the receptor's story, not the world's.

This one mechanism explains every folk fact about chili heat at once, which is how we know it is right. Water does not help because capsaicin is oil-bound and already seated in a pocket inside the membrane, where water cannot reach it; fat, and the casein of a glass of milk, help because they can lift it back out. The burn builds and then slowly resolves because TRPV1, flooded with the very calcium it admits, *desensitizes* — a calcium-dependent shutdown that is the same trick a capsaicin cream uses as a topical analgesic, the nerve quieted by the molecule that first inflamed it, and the same trick that makes a chili-eater's tolerance real, built channel by channel over years. And the heat reads as both pleasure and threat because it is, quite literally, the body's pain line being played like an instrument — the thrill of a warning you have chosen and know to be false.

It helps to hold TRPV1 as what it truly is: a *polymodal integrator*, a single channel summing heat, acid, capsaicin, and the cell's own membrane voltage into one decision to open. That is why the sensations stack — why a fierce curry on a hot afternoon in a feverish mouth burns past anything the capsaicin alone should manage, each input lowering the threshold the others must clear, the channel a coincidence-detector adding its warnings together. Capsaicin is simply the most selective key ever found to this lock: it speaks to this one receptor, cleanly, where almost nothing else in flavor manages such a single conversation. We design into exactly that selectivity, tuning the capsaicinoid output of a plant in the confidence that it will land on TRPV1 and nearly nowhere else. A clean channel gives a clean number, and the clean number is why heat was the axis on which the whole Index could first be shown to work.



A cut superhot: the pale placental tissue around the seeds is where capsaicinoids concentrate, and where the design does its work.

The shape of the curve

To write a number onto heat, we first have to say what "how hot" means in the only language the receptor speaks: concentration and response. A receptor does not switch; it leans. Add a little capsaicin and a few channels open; add more and more open, until adding more does almost nothing because nearly all of them are already open. Plot the fraction of channels open against concentration and you get an S-shaped curve, and the single most useful point on that curve is its middle — the concentration at which exactly half the

channels are open. That concentration is the **EC_{50}**, the half-maximal effective concentration, and it is the whole foundation of the Index. We read it the same way every time: a line of cells made to carry human TRPV1, loaded with a dye that brightens as calcium enters, is given the extract in a rising series of dilutions, and the concentration that lights exactly half of the maximum calcium signal is the EC_{50} the formula will take. One assay, one number, run identically for a shishito and for the hottest plant we have ever designed.

The curve obeys the Hill relation,

```
fraction open = [C]ⁿ ÷ ( EC₅₀ⁿ + [C]ⁿ )
```

where [C] is the capsaicin concentration at the receptor and *n* is the Hill coefficient — a measure of how steeply the channels come on together. For capsaicin at TRPV1, *n* sits near 2, an apparent cooperativity that echoes the four-fold architecture of the channel: the subunits do not open quite independently. A coefficient of 2 makes the curve pleasingly steep, and the arithmetic is easy to feel. Set [C] equal to EC_{50} and the equation returns exactly one-half, by construction — a threshold warmth has become a clear, present heat. Double the concentration to twice EC_{50} and it returns $2^2 \div (1 + 2^2) = \frac{4}{5} = 80\%$, a strong and insistent burn. Halve it to half EC_{50} and you get $0.5^2 \div (1 + 0.5^2) = 0.25 \div 1.25 = 20\%$, a friendly warmth; a quarter of EC_{50} barely crosses the threshold at 6%, and four times it nearly saturates at 94%.

The steepness is the whole story: a four-fold change in concentration — from half to twice EC_{50} — carries the channel from one-fifth to four-fifths open. Heat is a fast climb and then a long plateau, which is exactly why the EC_{50}, down in the responsive middle of that climb, is the honest place to measure. Up on the plateau a large change in concentration moves the response barely at all, and a number read there would mean little; at the foot the signal is too faint to trust. The half-point is the one place the curve is both steep and certain. (Perception is not linear in the open-channel fraction — the mouth's own wiring compresses the plateau and stretches the foot — so these percentages describe the receptor, not the sensation. We score the receptor because it is the part we can hold still.)

The molecule, assembled

Capsaicin is a condensation of two halves, and the balance between them is the *shape* of the burn, not merely its size. One half is **vanillylamine**, an aromatic amine the plant builds along the phenylpropanoid branch — **PAL → C4H → 4CL → HCT → COMT** — the same universal route, running in every vascular plant, that carries all the way to vanillin; vanillylamine is simply vanillin with one more enzyme's work on it, the aminotransferase **pAMT** swapping an aldehyde for an amine. The other half is a short, branched fatty-acyl chain, usually 8-methyl-6-nonenoyl, assembled from the carbon skeletons of branched-chain amino acids by **BCAT** and elongated by the ketoacyl synthase **KAS**. The enzyme that marries them —

capsaicin synthase, the Pun1 gene product, an acyltransferase of the AT3 family — joins the amine to the acyl chain and releases capsaicin.

The exact length and saturation of that acyl chain decide *which* capsaicinoid you get, and the capsaicinoids are not interchangeable. Capsaicin and its saturated twin dihydrocapsaicin are the sharpest — pure, each registers around sixteen million Scoville — and they carry the forward, upper-palate burn that arrives first and declares itself. Nordihydrocapsaicin, built on a chain one carbon shorter, is softer and fruitier, nearer nine million, and fades faster. The homocapsaicins, a carbon longer, are slower and deeper, pooling low in the throat and lingering after the sharp notes have gone. A pepper's heat is therefore never one thing. It is a *chord* of capsaicinoids, and designing heat means designing the chord — its voicing as much as its volume. Two fruits can carry the same total mass of capsaicinoid and burn like different instruments, one a struck bell and the other a bowed string, because the proportions across these four compounds differ. Picture it in round numbers. A chord that is three-quarters capsaicin at sixteen million and one-quarter nordihydrocapsaicin at nine million averages $0.75 \times 16 + 0.25 \times 9 \approx 14.3$ million in effective pungency — bright, forward, sharp. Invert the proportions — one-quarter capsaicin, three-quarters nordihydrocapsaicin — and the very same total mass averages $0.25 \times 16 + 0.75 \times 9 \approx 10.8$ million, a quarter softer and markedly fruitier, fading sooner from the palate. Identical grams; an entirely different pepper. That gap between the mass and the music is the gap the Index was built to see. We spend as much care on the ratio as on the sum; a beautiful heat is a well-voiced chord, and an ugly one is a loud single note.

The acyltransferase that strikes the chord is where the design acts most directly on which capsaicinoids are made, so we modeled it in atomic detail before we touched a single base of it. Our flux model predicts that primary control over total capsaicinoid output sits upstream, at 4CL and HCT in the phenylpropanoid trunk — the pools of vanillylamine and acyl precursor that Pun1 draws from — and not at the terminal Pun1 step itself. Pun1 is the final capsaicin synthase; it labels the product but does not govern the rate. The designed edits act at Pun1 precisely because that is where the chord's voicing — the ratio and character of the capsaicinoids — is set. The structure below is the Pun1/AT3 product colored by confidence: the binding domain and the active-site cleft, where the amine and the acyl chain meet, resolve at high confidence — which is precisely where a designed edit can change the output without breaking the fold.

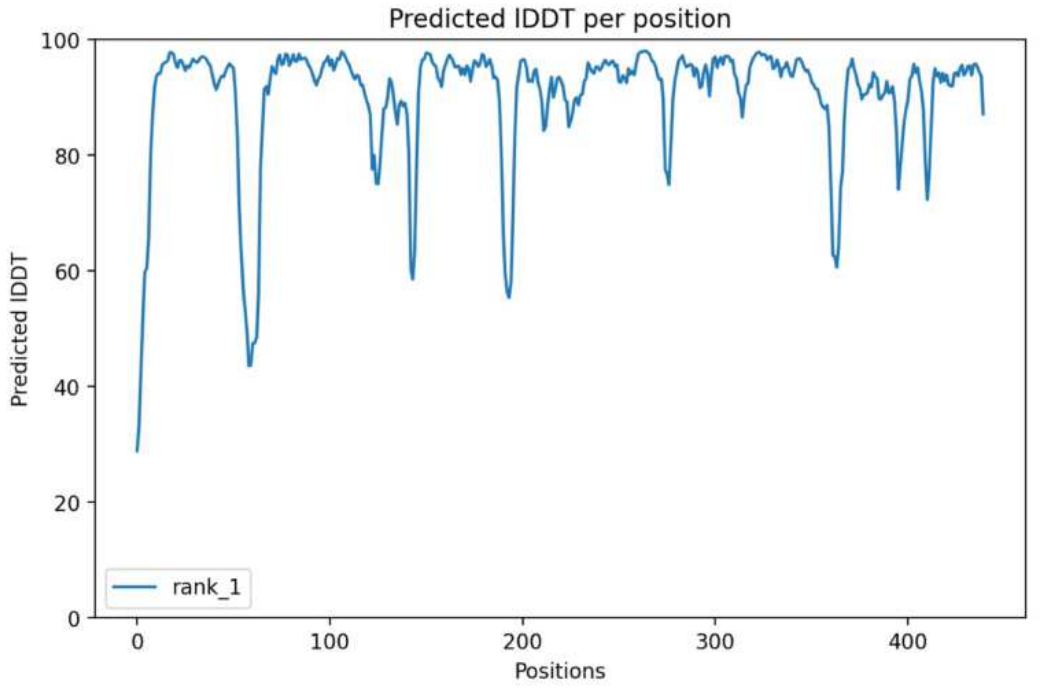

The capsaicin-synthase (Pun1/AT3) model, colored by prediction confidence; the binding domain and active-site cleft, where the chord is struck, are the high-confidence core.

Two scales, one plant

Heat is the one flavor the world already tried to number, and the attempt is instructive. The **Scoville** scale began in 1912 as a panel of human tasters diluting a pepper extract in sugar water until the burn vanished; the dilution at which it disappeared was the score. Today the panel is gone and the number comes from HPLC — the fruit is extracted, its individual capsaicinoids separated and weighed, and the masses converted to heat units by their known relative pungencies, with pure capsaicin set near sixteen million. The conversion is a plain proportion: pure capsaicin is a million parts per million by definition, so sixteen million Scoville spread across that million parts fixes the factor at roughly 16 SHU for every part per million of capsaicinoid. An HPLC trace reading a few hundred ppm of total capsaicinoids multiplies straight up into the familiar six- and seven-figure numbers. The modern Scoville value is therefore a *chemistry*: how much capsaicinoid is present, by mass. It is an honest measurement of the gram.

The Index asks the other question. **IFI-H** is not a mass but a *loudness* — how strongly the capsaicinoid chord drives TRPV1 — read against a fixed reference and scaled so the arithmetic is kind:

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

The anchor is the native wild-type reference, defined so that its own score is exactly 1,000. That is the zero-point the entire heat axis is tuned to — one note, the half-activation of an ordinary wild pepper, against which everything else is scored:

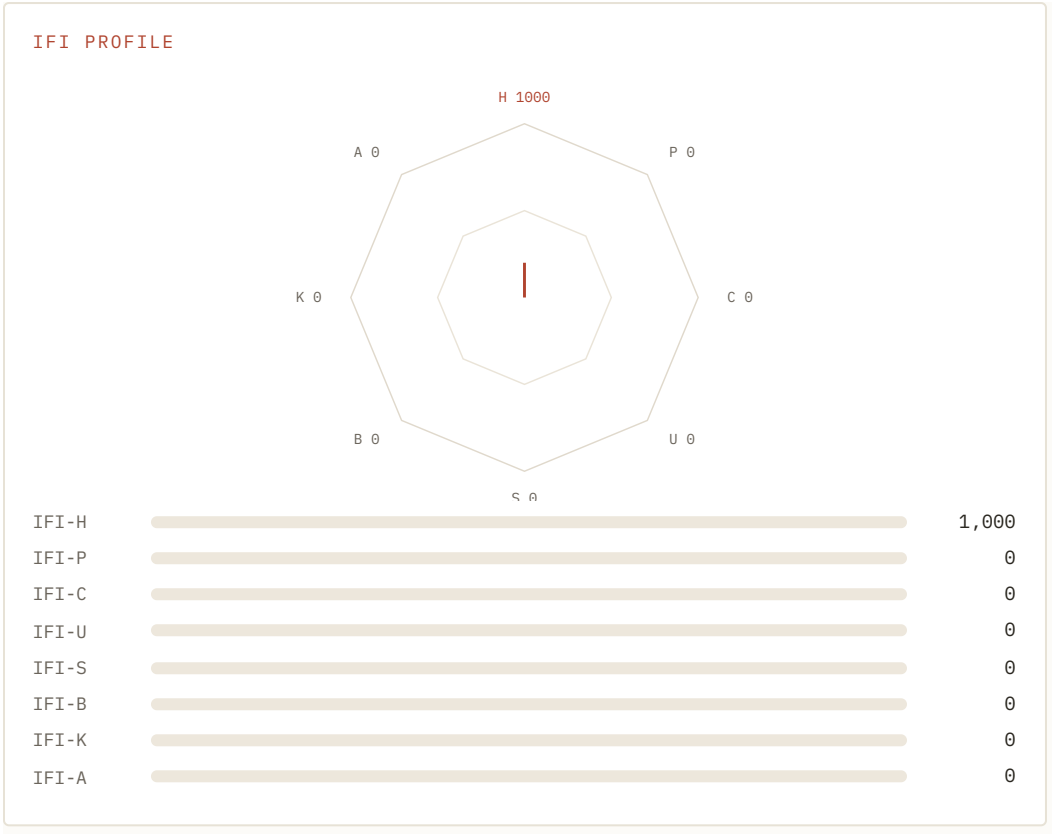

Because only the *ratio* of EC_{50} values enters the score, the absolute size of the reference is a convention we are free to fix; we set the standardized wild-type extract to reach half-activation at EC_{50}(reference) = 2.00 µM, expressed as capsaicin-equivalent concentration in the assay bath. Now the score inverts into a real concentration for any plant. Our flagship, **Phoenix**, reads IFI-H 4,200, so its extract is designed to reach half-activation at

EC_{50}(Phoenix) = 1,000 × 2.00 µM ÷ 4,200 = **0.476 µM**.

A louder plant is a *lower* EC_{50} — it needs less to say the same thing — and 0.476 against 2.00 is almost exactly one-quarter: Phoenix drives the channel to its half-point at a quarter of the reference's concentration, which is what an IFI-H of 4,200 against 1,000 means, 4.2 times as loud. The same inversion, run in everyday units, is just as stark. Capsaicin's molar mass is 305.4 g/mol, so 1 µM is about 0.305 mg/L, which in dilute water is 0.305 ppm. The reference's half-point sits at 2.00 µM ≈ 0.61 ppm; Phoenix's sits at 0.476 µM ≈ **0.145 ppm**. A seventh of a part per million, at the receptor, for half the fire a nervous system can feel.

And the SHU, the gram-number, tells a parallel story that occasionally, usefully, disagrees. Our model predicts Phoenix at approximately **2.95 million SHU**. Against pure capsaicin's sixteen million, that is roughly 2,950,000 ÷ 16,000,000 ≈ 0.184 — the design equivalent, in Scoville terms, of an 18.4% dilution of pure capsaicin. The two scales climb together most of the way up the ladder, because more capsaicinoid usually does mean more receptor drive. But they are measuring different things, and where they part, the mouth sides with the Index. A milligram of soft, fruity nordihydrocapsaicin and a milligram of sharp capsaicin weigh the same and do not burn the same. When the gram and the loudness disagree, trust the loudness — it is the one that matches the tongue.

The ladder

We built all of this out not on a spreadsheet but on ten designed cultivars. These are *Capsicum chinense*, the hottest species, advanced construct by construct from a fierce but ordinary pepper

— ten rungs of increasing computational design scope we drew from the library so the Index would have something to stand on. The EC_{50} column is derived from each cultivar's IFI-H by the formula above, against the 2.00 µM anchor:

CULTIVAR	SKU	PREDICTED SHU	IFI-H	EC_{50} (MM)	AXES
Ember	SS-0087	866K	1,570	1.274	2/8
Cinder	SS-0088	1.05M	2,290	0.873	3/8
Kindle	SS-0089	1.25M	2,000	1.000	3/8
Flare	SS-0090	1.48M	—	—	4/8
Blaze	SS-0091	1.72M	2,630	0.760	4/8
Wildfire	SS-0092	1.96M	2,930	0.683	5/8
Scorch	SS-0093	2.21M	3,220	0.621	5/8
Firestorm	SS-0094	2.47M	3,490	0.573	6/8
Solara	SS-0095	2.72M	3,980	0.503	6/8
Phoenix	**SS-0096**	**2.95M**	**4,200**	**0.476**	**8/8**

The ladder hides its best lesson in one apparent contradiction, and it is worth stopping for, because it is the entire argument for scoring the receptor instead of the gram. Look at Cinder and Kindle. Kindle carries *more* capsaicinoid mass — 1.25M SHU against Cinder's 1.05M, nearly a fifth more by weight — yet it scores *lower* on the Index, IFI-H 2,000 against 2,290, and its EC_{50} is accordingly higher, 1.000 µM against 0.873. Kindle needs 1.000 ÷ 0.873 = 1.15 times the concentration to reach the same half-activation. It is a seventh *quieter* at the receptor while carrying a fifth *more* capsaicinoid by mass — because its chord is weighted toward the softer, fruitier, less receptor-potent compounds. More grams, less fire. The Scoville number counts the molecules; the Index listens to them. On the tongue, Cinder is the hotter pepper, and the Index is the only one of the two scales that agrees with your mouth.

Read the two right-hand numeric columns together and the design philosophy is laid bare. Heat does not climb alone. As IFI-H rises up the ladder, the *axis count* rises with it — two, three, five, finally all eight — because the same platform that writes more capsaicin into the placental tissue is writing umami, kokumi, and the quieting of bitterness into the same fruit. By the top rung, heat is no longer the whole flavor; it is the loudest voice in a full chord. One rung, Flare, carries a dash where the others carry a score — and we leave it a dash on purpose. Its dose-response curve did not clear our fit-quality gate; an interfering co-extracted compound skewed the receptor assay, and rather than publish a loudness we do not trust, we publish nothing for that axis and say why. A withheld number is a form of respect for every number we do print.



Phoenix (SS-0096) in the controlled chamber — the first cultivar grown to register on all eight axes at once.

Reading Phoenix

At the top of the ladder is the cultivar the rest of this library keeps returning to, because it is the first complete profile we designed. Our model predicts approximately 2.95 million SHU for Phoenix and an IFI-H of 4,200, and it does so while engaging all eight axes at once:

IFI PROFILE

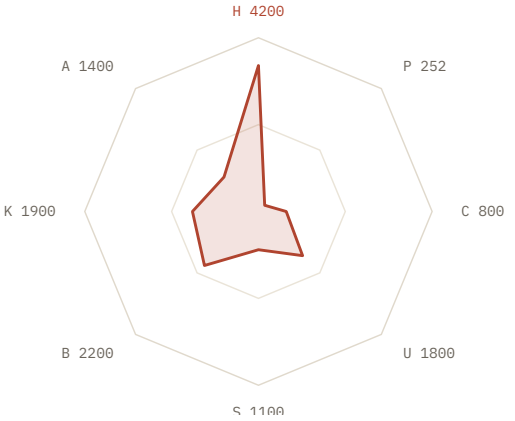

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

Scorecards normalize each axis against a per-axis maximum of 5,000, so these eight numbers draw a specific, deliberate shape. Heat fills 4,200 of 5,000 — 84% of its ring — and it is the only axis that loud. Everything else is placed in relation to it, and the placement is the art:

AXIS	RECEPTOR	IFI VALUE	SHARE OF 5,000-MAX
H · Heat	TRPV1	4,200	84%
B · Zero-Bitter	TAS2R38	2,200	44%
K · Kokumi	CaSR	1,900	38%
U · Umami	mGluR1	1,800	36%
A · Astringency	Salivary PRPs	1,400	28%
S · Sweetness	TAS1R2·R3	1,100	22%
C · Cool-Burn	TRPM8	800	16%
P · Pungency	TRPA1	252	5%

Notice what the profile refuses to do. With heat at 4,200 the easy thing is a one-note inferno. The designed thing is a 4,200 that sits *inside* a flavor. Bitterness is pressed hard — IFI-B at 2,200, the strongest attenuation on the card — so the fruit comes through clean rather than acrid under all that fire. Kokumi is lifted to 1,900 and umami to 1,800, a floor of richness and savory depth so the burn has a body to ride on instead of scorching across an empty plate. And cool-burn is held deliberately *below* the reference, at 800 — under the wild-type's own 1,000 — so nothing minty or cooling cuts across the warmth; the plant leans faintly, intentionally, toward heat even on the channel that could have fought it. Pungency is held low at 252 — well below the reference — adding just a quiet edge of TRPA1 sharpness at the finish; astringency at 1,400 gives a tannic grip that keeps the sweetness, held low at 1,100, from turning cloying. Eight numbers, one decision each, and together a pepper that is enormous and still kind — heat written by people who believe a chili should be a pleasure and not an assault.

Precision as respect

A number you cannot reproduce is not a measurement; it is an anecdote. The honest test of a designed flavor is not its peak on a good day but its spread across a season, and that spread has a name: the coefficient of variation, CV = σ ÷ μ, the standard deviation of a batch divided by its mean, as a percent. Wild-type superhots are notoriously fickle — sun, stress, water, and soil swing their heat by enormous margins, and a representative wild line runs a batch-to-batch CV near **42%**. Our engineered lines are designed to hold **≤ 7%**. The difference is not cosmetic; it is the whole difference between a plant and a product.

Put Phoenix's predicted mean through both. At a predicted μ ≈ 2,937,948 SHU and a 7% CV design target, the standard deviation is specified at 0.07 × 2,937,948 ≈ **205,656 SHU**, so the design places roughly two batches in three inside 2.73M–3.14M — a tight, dependable window you could formulate a sauce around. At a wild-type's 42%, the same mean would carry σ ≈ 0.42 × 2,937,948 ≈ **1,233,938 SHU** — a spread six times wider (42 ÷ 7 = 6), ±1.2M, overlapping half the ladder. A plant whose heat you cannot predict to within a few percent is a lottery ticket, not a line on a spec sheet. The reproducibility *is* the engineering.

Both the predicted height and the design specification for consistency trace to two edits — a double variant in the Pun1 enzyme, each substitution proposed by a protein language model scoring every possible single-residue substitution, written in with CRISPR, never combined by any breeding program:

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

† C175S is a structural annotation from modelling; it was not included in the final design. The shipped double variant is S39L+L345G only.

The ΔΔG column is why the CV falls, not just why the peak rises. Each substitution in the double variant lowers the enzyme's folding free energy — a negative ΔΔG is stabilizing — for a combined −3.9 kcal/mol of added fold stability. A more stably folded capsaicin synthase misfolds less, degrades less, and runs at a more uniform rate across the thousands of cells and the weeks of ripening that make a fruit, so the computed flux into capsaicin varies less from plant to plant. S39L tightens the substrate's grip at the N-terminus and lifts predicted output; L345G opens the active-site cleft so the acyl chain seats cleanly. Height from the first, consistency from both.

Honesty requires naming the edges where even this gets hard. TRPV1 desensitizes under its own calcium, so a dose-response read too fast or too slow drifts — we hold timing and temperature to fix the EC_{50} in place. At the very top of the ladder the receptor assay begins to saturate, the plateau of the Hill curve crowding the measurement, which is part of why a loudness beyond the sharp climb must be read with care. And a co-extracted interferent can masquerade as potency, which is exactly the fault that earned Flare its dash. We would rather carry a visible gap than a quiet error. Numbers are how we show respect for the reader, and a number we cannot stand behind we do not print.

Heat was the proof

Heat was always going to be first, because it is the axis the body reports most purely and the one the world already cared about — but it earned its place as the proof for a reason larger than fame. On this one clean channel, with one selective molecule and one well-modeled enzyme, we showed that a flavor could be computed, targeted, and grown to a design specification that holds to within a few percent, batch after batch. How the two precise edits of the double variant were found, out of all the substitutions the gene could have taken, is a longer story than heat alone; for now it is enough that they were found, and that the warmth they are designed to produce arrives on cue. What we designed is not simply a hotter pepper. It is the demonstration that a plant able to reach a flavor no longer has to be waited for across ten thousand harvests — it can be designed, and designed gently, a fire with a body and a kindness in it, grown for a tongue we hope to delight and a body we mean to

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

Chapter 6 · Sweetness — *sugar's signal without sugar's cost.*

Sweetness is the taste the body trusts most and questions least — the ancient all-clear, the first word a mother's milk ever speaks to a newborn, the signal that a food is energy and safe. We are born already loving it; of the eight tastes it is the only one that needs no teaching, the only one an infant will reach for before it has reached for anything else. That trust is exactly why it is the most interesting axis to design, because the receptor that carries it can be spoken to by molecules that are not sugar at all, that the body never metabolizes, that pass through unburned and uncounted, and that nonetheless ring the same bell with the same sincerity. To understand how is to understand why sweetness, uniquely among the eight, can be kept while its cost is set down — why the pleasure and the price, which the living world has always delivered knotted together, can be pried apart at the level of a single protein. This is the axis where our family's quietest conviction, that a food should nourish without exacting, stops being a wish and becomes an equation.

The receptor is the whole argument

The sweet receptor is a single heterodimer: **TAS1R2 joined to TAS1R3**, a class-C G-protein-coupled receptor pair, and almost everything distinctive about sweetness follows from its architecture. Each half carries, at its business end, the large clamshell domain called the **Venus flytrap** — two lobes hinged like a trap, held open at rest by the jostle of water. A sugar molecule drifts into the cleft between the lobes of the TAS1R2 flytrap and the lobes fold shut around it, and it is that closure — transmitted down through a **cysteine-rich stalk** to the seven-pass transmembrane bundle buried in the cell membrane — that shoves the G protein and begins the relay.

The relay itself is worth tracing step by step, because its slowness is its signature. The activated receptor exchanges GDP for GTP on its coupled G protein (gustducin in the taste cell), freeing the Gβγ subunits to switch on **phospholipase C-β2**. PLCβ2 cleaves the membrane lipid PIP_2 into two messengers; one of them, **IP_3**, diffuses to the endoplasmic reticulum and opens its calcium stores. The released **calcium** throws open **TRPM5**, a channel that admits sodium and depolarizes the cell, and that depolarization drives **ATP** out through the CALHM1/CALHM3 pore onto the waiting gustatory nerve, which reads the ATP as *sweet* and carries it to the brain. Five handoffs, each taking its measured moment. Compare this to the taste of salt or sour, where the stimulus is an ion that walks straight through a channel and depolarizes the cell with no messenger in between. Sweetness is built on a cascade, not a wire, and a cascade takes time to fill and time to drain. That is the mechanistic reason sweetness **blooms** rather than strikes, rises across a second or two and lingers as the second messengers clear — a shape of sensation written into the plumbing of the cell.

The roominess of the receptor is why so many unrelated molecules can trip it. The sweet dimer is not a lock with one keyhole but a building with several doors: the **TAS1R2 flytrap cleft** admits sugars and small synthetic sweeteners like aspartame and sucralose; a separate pocket inside the **TAS1R3 transmembrane bundle** is where cyclamate and neohesperidin dihydrochalcone act, and where the blocker lactisole plugs in; and the broad outer face of the flytrap is where the **sweet proteins** dock, draping across the closed trap and clamping it shut from without. One receptor, many doors, and a molecule needs only find one of them to be read as sweet. This is the architectural fact the whole chapter turns on: because the doors are many and the reading is the receptor's, sweetness is a property of the *reader*, not of any one class of key.



Stevia — the leaf that holds its steviol glycosides hundreds of times louder than sugar, gram for gram.

Proof by silence

That the receptor is the sole judge is proven most vividly by what happens when it is absent or silenced. A cat cannot taste sweetness at all. Somewhere in the deep history of the obligate carnivores the gene for **TAS1R2 broke** — it survives in the feline genome only as a pseudogene, a ruin of itself — and was never repaired, because an animal that eats only meat never needs to hear the sugar signal. To a cat, a lump of sugar is simply flavorless, whatever energy it holds; the molecule arrives, the door is missing, nothing is read. Run the experiment the other way and you can switch the sensation off in a human in seconds. **Gymnemic acid**, chewed from the leaf of *Gymnema sylvestre*, blocks the sweet receptor so thoroughly that sugar laid on the tongue a moment later tastes of wet sand — pure grit, the sweetness deleted while the texture remains. The synthetic blocker **lactisole** does the same by plugging the receptor's transmembrane pocket from inside. In both cases the sugar on the tongue has not changed by a single molecule; only the reader has gone quiet.

There is an inverse wonder that makes the same point from the other side. **Miraculin**, a protein from the West African miracle berry, binds the flytrap and sits there harmlessly at neutral pH — but drop the pH with something sour, a slice of lemon, and the protein's shape shifts into an agonist, so that for an hour afterward anything acidic reads as intensely sweet. The lemon has not gained sugar. The receptor has been handed a new conformation and reports what the conformation tells it. Sweetness, every one of these experiments

says, is not in the sugar but in the receptor's reading of the sugar; when the reader is silenced the taste is gone, instantly and completely, and when the reader is tricked the taste appears out of an acid that holds no sweetness at all. Every move the Index makes on this axis follows from taking that one fact with full seriousness: **we design the signal, not the sugar** — the EC_{50} at the flytrap, not the mass in the bowl.

Loudness is one over EC50

The Index scores every axis by loudness, and loudness has a precise definition: it is the inverse of the concentration a molecule needs to half-fill its receptor. Write that concentration **EC_{50}** — the dose at which the receptor sits at half of its full response — and the score is

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

The reference for sweetness is **sucrose**, anchored at exactly 1,000. A molecule that half-activates the flytrap at a *lower* concentration than sucrose is, by this definition, louder; the lower its EC_{50}, the higher it scores, and the climb can be astonishing. Sucrose itself is a quiet molecule in receptor terms, needing tens of millimolar to reach half-activation, which is why sweetening with sugar means sweetening with **mass**.

To see how far the loud molecules sit above it, take representative half-activation concentrations — the kind that standard cell-based assays return, reading calcium flux in cells expressing the human dimer — and run them through the formula. The molecular weights are fixed facts of chemistry; the potencies fall straight out of the arithmetic.

SWEETENER	EC_{50} (MOLAR)	MOLECULAR WEIGHT	MOLAR LOUDNESS, 1,000 × EC_{50}(SUC)/EC_{50}	POTENCY BY WEIGHT
Sucrose *(reference)*	~30 mM	342 g/mol	1,000	1×
Rebaudioside A	~50 µM	967 g/mol	~600,000	~210×
Mogroside V	~32 µM	1,287 g/mol	~940,000	~250×
Thaumatin	~150 nM	22,200 g/mol	~200,000,000	~3,100×

Read the rebaudioside A row slowly, because it carries the whole idea. Its EC_{50} is about fifty micromolar against sucrose's thirty millimolar — a six-hundred-fold gap in molar terms — so its molar loudness is 1,000 × (30 mM ÷ 0.050 mM) = **600,000**. The receptor hears it six hundred times more loudly than sugar, molecule for molecule. But a cook measures sweetener by weight, not by count, and reb A's molecules are heavy: at 967 g/mol each weighs nearly three sucrose molecules. Convert loudness to the kitchen's units and you multiply the molar ratio by the ratio of weights, 342 ÷ 967 = 0.354, so **600 × 0.354 ≈ 210×** by weight. That single conversion is why the same leaf can be called "six hundred times sweeter" by a receptor biologist and "about two hundred times sweeter" by a formulator and both be exactly right. The **mogrosides** of monk fruit run the same arithmetic to about 250× by weight; the **sweet proteins** — thaumatin, brazzein, monellin, mabinlin, which do not enter the

cleft at all but wedge the trap shut from outside — run it to the thousands, thaumatin landing near 3,100× by weight off an EC_{50} down in the nanomolar.

Two more numbers make the abstraction edible. Multiply each EC_{50} by its molecular weight and you get the half-activation dose by mass: sucrose, 30 mM × 342 g/mol ≈ **10.3 g/L**, so the sweet receptor reaches half-response at roughly one percent sugar — the faint sweetness of a thing just barely sweet. Reb A, 50 µM × 967 g/mol ≈ **48 mg/L**, forty-eight parts per million — a trace you could not see in the glass. The ratio of those masses, 10,300 ÷ 48 ≈ 212, is the weight potency again, arrived at from the other direction.

And the dose-response is gentle, which matters as much as the loudness. Model the single dominant binding event with the simplest Hill relation, response $\theta = C \div (EC_{50} + C)$. At $C = EC_{50}$ the receptor is exactly half-lit, $\theta = 0.5$. To climb from a tenth lit to nine-tenths lit — from the first whisper of sweetness to nearly full — the concentration must span from $EC_{50}/9$ up to $9 \times EC_{50}$, an **eighty-one-fold range**. Sweetness, in other words, turns up slowly across a wide band of dose; there is no cliff. That forgiving slope is why a sweetener can be dialed to a target by small steps, and why a tenth of a gram one way or another does not tip a drink from flat to cloying. The receptor grades gently, and the design grades gently with it.

From a molecule's loudness to a fruit's score

A molar loudness of six hundred thousand is a true number and a useful one, but it is a property of a *molecule*, and a scorecard does not grade molecules — it grades a fruit, a leaf, a glass of something. The axis value that appears on an IFI profile is the **delivered** sweetness of a food: how loudly the whole thing speaks when it is in the mouth, set on a band where the sucrose reference reads 1,000 and the top of the scale is pinned at **5,000**. The reason a food never scores six hundred thousand, no matter how potent the sweetener inside it, is the receptor itself. As concentration climbs, $\theta = C \div (EC_{50} + C)$ crowds toward 1 and cannot pass it; the flytrap is either closed or it is not, and once nearly all the receptors are closing on nearly every cycle, adding more compound buys almost no more signal. Sweetness **saturates**, and the 5,000 cap is not an arbitrary ceiling drawn by us but the honest shadow of that saturation. Past the top of the band you are no longer adding sweetness; you are only adding mass and, usually, off-notes.

Here is what that means in a real reformulation, worked all the way through. A 355 mL soft drink sweetened to the usual intensity carries about **39 g of sucrose**. In the glass that is 39 g ÷ 0.355 L ≈ 110 g/L, and in molar terms 110 ÷ 342 ≈ 0.32 M — call it 321 mM, better than ten times the sucrose EC_{50}. Put that through the dose-response and $\theta = 321 \div (30 + 321) \approx$ **0.915**: the receptor is running at ninety-one percent, which is what "a sweet soda" means in the language of the flytrap. Now hold that occupancy fixed and ask what dose of reb A reaches it. Invert the Hill relation, $C = EC_{50} \times \theta \div (1 - \theta) = 50$ µM × (0.915 ÷ 0.085) ≈ 535 µM, and weigh it: 535 µM × 967 g/mol ≈ 517

mg/L, which in 355 mL is about **0.18 g**. The same ninety-one percent of receptor, the same delivered IFI-S — reached with 0.18 g of sweetener instead of 39 g. **Thirty-nine grams of sugar become two-tenths of one gram.**

Notice what the arithmetic does *not* depend on: the target intensity. The weight potency is a ratio of concentrations times a ratio of molecular weights, and at any matched receptor occupancy the concentration ratio is fixed at the ratio of the two EC_{50} values — the dose-response curves for sucrose and reb A are the same shape, simply shifted along the concentration axis, so they run parallel and never cross. Match a gently sweetened iced tea instead of a soda — half-lit receptors, $\theta = 0.5$, which for sucrose is $C = EC_{50} = 30$ mM, about 3.6 g in a 355 mL glass — and reb A reaches the same half-light at its own EC_{50}, 50 µM, about 17 mg. The ratio is 212× again, to the digit, exactly as it was for the soda. The same two-hundred-fold mass saving holds whether a food is barely sweet or flooded with it, and that invariance is what lets a single potency figure travel across an entire product line without re-derivation: measure a sweetener once against the reference, and you have priced it at every intensity it will ever be used at.

That arithmetic is the gift and the problem in one line. The gift is obvious: the sweetness is kept and the metabolic load, the 39 g the body would have had to burn or store, is simply not there. The problem is the **38.8 g of vanished mass** — the body of the drink, its viscosity, the way sugar coats the tongue and fills the mouth and gives a beverage its slip and weight. Sweetness was never the only thing sugar was doing; it was also doing the work of bulk, and when you replace sugar by loudness alone you leave a hollow where the body used to be. Rebuilding that body without rebuilding the mass is a separate design problem, and it is usually solved on the **kokumi axis** — the richness and continuity carried by the calcium-sensing receptor — which can hand a drink back its weight and mouth-filling heft without handing back its calories. A sweetness correct on the meter can still be wrong in the hand, and the fix lives on a different axis entirely.

One more consequence of saturation is worth keeping, because it runs in our favor. Error in a food is error in concentration — the leaf accumulated a little more or a little less of its glycoside this season, the batch came in rich or lean — and we hold that batch-to-batch variation to a coefficient of variation of **≤7%**, against the 40-plus-percent swing an unselected line will drift across. Track that 7% through the dose-response and it behaves differently at different places on the curve. Near saturation the receptor compresses it: the fractional change in activation is $(1 - \theta)$ times the fractional change in concentration, so at $\theta = 0.915$ a 7% wobble in sweetener becomes $0.07 \times 0.085 \approx$ **0.6%** in delivered sweetness — the flytrap, already nearly full, barely notices. Down in the quiet, where a food is only faintly sweet and θ is small, $(1 - \theta) \approx 1$ and the full 7% passes through unsoftened. The receptor is its own buffer at the top of the scale and a faithful linear reporter at the bottom, and knowing which regime a product lives in tells us exactly how tightly the agronomy has to hold.

How the plant builds its sweetness

The plants assemble these molecules by routes the design can reach into and redirect, and the whole art is in the *endpoint*. Stevia draws its sweetness off the plastidial isoprenoid line — the same scaffolding that builds a plant's growth regulators — through ent-kaurene to **steviol**, a diterpene core that is not itself what we taste. The taste is made in the dressing. A family of **UDP-glycosyltransferases** (UGT85C2, UGT74G1, UGT91D2, and above all UGT76G1) hangs glucose groups onto the steviol core one at a time, and it is the *pattern* of that glycosylation, not the amount of steviol, that decides everything about how the leaf tastes.

GLYCOSIDE	GLUCOSE UNITS	RELATIVE SWEETNESS	OFF-NOTE TENDENCY
Stevioside	3	high	pronounced bitter / licorice tail
Rebaudioside A	4	high	much reduced
Rebaudioside D	5	high	low
Rebaudioside M	6	high	lowest; nearest to clean sugar

Read down that ladder and the design lever reveals itself. **Rebaudioside A is simply stevioside wearing one more glucose** — a single sugar added by UGT76G1 — and that one decoration lights the sweet receptor at least as fully while brushing the bitter receptors far less. Add another glucose to reach reb D, another to reach reb M, and the sweetness holds while the off-note keeps falling, reb M arriving about as clean as the leaf can make it. So most of the stevia work in the library — **910 constructs** across this one plant — is not an effort to make stevia *sweeter*. It is an effort to make it **finish clean**: to drive the plant's own glycosylation toward the rebaudioside forms that ring the sweet bell and almost nothing else, by tuning the expression and specificity of the UGTs that set the endpoint. We are not turning a sweetness up; we are moving a decision about which molecule the leaf finishes with.



Monk fruit — the mogrosides build their sweetness on a triterpene core, clean when the final sugars sit right.

Monk fruit runs the same logic on a different skeleton. Its sweetness lives in the **mogrosides**, cucurbitane triterpenes built from squalene through cucurbitadienol and mogrol, then glycosylated — and again the final sugar decorations set the balance between clean sweetness and a lingering edge, with **mogroside V**, the fully dressed

form, carrying the clean weight. **80 constructs** work this fruit toward the forms that finish right. And sugarcane, the reference sugar itself at the foot of the whole axis, anchors the other end of the shelf: **200 constructs**, a frontier platform, the baseline against which every louder molecule is measured.

ORGANISM	CONSTRUCTS	SWEET PRINCIPLE	APPROX. POTENCY VS SUCROSE	ENGINEERING DIFFICULTY
Stevia	910	steviol glycosides (reb A → reb M)	~200–300×	Moderate
Monk Fruit	80	mogroside V	~250×	Hard
Sugarcane	200	sucrose platform	1× (reference)	Frontier



Sugarcane — the reference sugar, the 1× baseline every louder molecule is scored against.

The two flaws the receptor does not cause

Most natural high-potency sweeteners carry a flaw, and the honest thing is that the sweet receptor does not cause either of the two that matter. The first is **bitterness**. Stevia's faint bitter edge and licorice tail are not the flytrap misfiring; they are the glycosides brushing a different receptor entirely — the bitter sentinels **TAS2R4 and TAS2R14** — on the way past. The sweet signal and the bitter signal are read by separate proteins and travel separate channels, which is precisely why the Index scores them on **separate axes**, and why cleaning the off-note is a receptor problem with a receptor's answer: shift the glycoside profile toward reb D and reb M, whose shapes the bitter receptors barely register, and the bitter tail falls while the sweetness holds. The second flaw is **time**. Sweet proteins and some glycosides carry a different envelope than sugar — a slower onset, a longer and occasionally cloying tail — because their binding and release run on different clocks than sucrose's quick in-and-out of the flytrap. The envelope a person actually feels is the binding kinetics *convolved with* the cascade's own fill-and-drain time: a molecule with a slow off-rate holds the flytrap shut after the sip is gone, and a flytrap held shut keeps the IP_3-to-calcium-to-TRPM5 relay firing, so the lingering tail is a receptor phenomenon and a second-messenger phenomenon at once, each stretching the other. Sucrose's quick release lets the binding and the cascade clear together; a sweet protein's tenacious grip keeps both running long after the sugar signal should have stopped. A sweetener correct on the meter and wrong on the clock reads as faintly *off* even when its

number is exact, and matching sucrose means matching not only its loudness but the shape it draws across the seconds it is in the mouth. We tune both, because the tongue grades both.

Those are the flaws we can fix. Honesty requires naming the limits we cannot design around, because numbers are a form of respect only if they admit where they stop. Loudness is not pleasure: the IFI measures how fully a receptor is driven, not how much a person likes the result, and the hedonic verdict depends on the temporal envelope and the state of the other seven axes as much as on the sweet score itself. Readers differ, too — polymorphisms in the sweet and bitter receptors mean the same reb A leaf reads cleaner to one person than to another, and a sharp detector of bitterness will catch a tail a duller one sails past, so a "clean" sweetener is clean to a population, with a spread. Perception drifts with temperature and with the mouth's own adaptation across a meal. And the receptor saturates, which means that above the top of the band more compound is wasted on sweetness and spent only on off-notes — the cap is a real wall, not a conservative one. We state these not to hedge but because a measurement you can trust is one whose edges are drawn as carefully as its center.

Combinations, cancellation, and reformulation with the lights on

Because the sweet dimer has several doors, two sweeteners that bind at *different* sites can be combined to do what neither manages alone. Dose a flytrap sugar alongside a transmembrane-pocket sweetener and the two can read as **more than their sum** — a genuine synergy, the receptor driven harder from two directions than either push could reach — so a blend hits a target intensity on less total sweetener than a single compound would need. The arithmetic of it is clean at low dose, where activation is near-linear in concentration and two different-site drives simply add: a flytrap sweetener contributing a drive of, say, 0.3 and a pocket sweetener contributing another 0.3 would, if the sites were blind to each other, sum to a response of 0.6. But the flytrap and the transmembrane pocket are allosterically coupled — occupancy at one lowers the effective EC_{50} at the other — so the measured response climbs past the additive prediction, nearer 0.7 for the same total mass, and that **superadditive margin** is exactly what a blend is built to harvest: the same sweetness on less sweetener, with the surplus spent on cutting off-notes rather than on more of the compound that carries them. More useful still, two sweeteners can **cancel each other's flaws**: one sweetener's bitter tail silenced by another's clean body and heft, the slow onset of a protein braced by the quick arrival of a glycoside. Much of modern sugar reduction is this quiet combinatorics of the receptor's several pockets — a blend tuned so the doors open together and the off-notes do not.

The reason we can do it deliberately rather than by trial is that the Index scores sweet and bitter apart. Watch a blend come together and you can see an off-note **fall on IFI-B while the sweetness holds steady on IFI-S**, and know you have improved the thing rather than merely moved the problem from one part of the mouth to another. It

is reformulation with the lights on. The reference against which all of this is read is the plainest profile in the library — sucrose, a single note and nothing else:

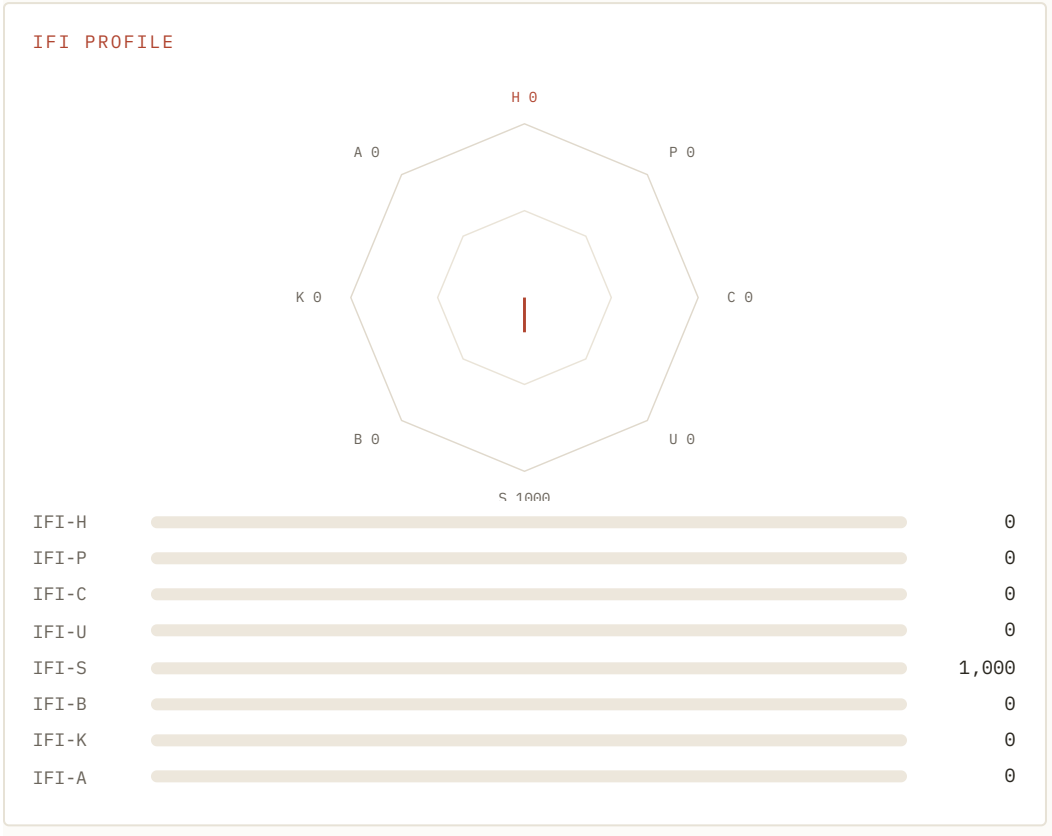

Pure sweetness at exactly the reference loudness, silent on every other axis. That one-note profile is the anchor, the 1,000 that every other sweetener is measured against, and it is also a kind of lesson: sugar does one thing, and the whole discipline of this axis is learning to do that one thing with a molecule the body never has to answer for.

On the composite side, sweetness is the natural partner of two other axes. Paired with heat it makes the **swicy** chord, the sweet-hot balance that leads a whole shelf of our work; paired against bitterness it does the gentle arithmetic of balance, lifting the perception of sweet while the bitter is pressed down, so a dark chocolate or a cold brew arrives rounded without a grain of added sugar. And it reaches even to the far edge of the heat scale. Phoenix itself — a fruit built to register on all eight axes at once, with our model predicting approximately 2.95 million SHU — carries its predicted sweetness at **1,100** on IFI-S, just over the sucrose reference:

IFI PROFILE

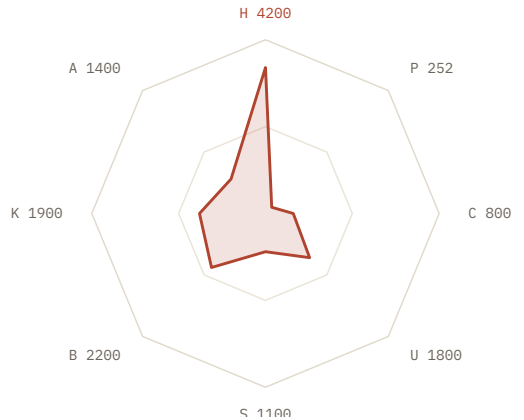

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

Not a dessert — a fruit that remembers it is a fruit, the faint sweetness, ten percent above reference and sitting low on a band that climbs to 5,000, that keeps 2.95 million Scoville of designed heat from reading as pure punishment. We kept that little sweetness in at the very edge of what a plant can be made to do, because a thing that nourishes should also, somewhere in it, be kind. That is the conviction under this whole axis stated plainly: a sweetness the tongue believes completely and the metabolism never has to answer for is not a health promise and never a claim — it is an old pleasure returned to people who had set it down, carried by a molecule the body meets as a guest and never as a debt. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 7 · Umami — *the taste of being cared for.*

Of the eight axes this is the one our family argued about the longest, because it is the hardest to say and the easiest to feel. Umami is the deep, mouth-filling roundness of a long broth, a vine-warm tomato, a slow-simmered anything — the taste that reads, below language, as *nourishment,* as the food the body was hoping for. For most of culinary history the West did not even have the word. It took a chemist, Kikunae Ikeda, tasting his way to the bottom of a bowl of kombu broth in 1908, to name it and to find that its source was a single amino acid, free and unbound on the tongue. The Index gives that taste a number, and in doing so gives the most comforting of flavors a receptor, a chemistry, and a way to be grown back into food that modern handling has quietly stripped it from.

We write this axis with particular care, because umami is the one most often faked and most often mourned. Faked, because a pinch of monosodium glutamate will counterfeit its surface without any of its depth. Mourned, because the foods that once carried it honestly — the tomato ripened to the edge of falling, the stock built over a slow afternoon — have been bred and handled by a supply chain that optimizes for everything except savor. What follows is the whole machinery of it: the molecule, the receptor that reads it, the arithmetic that turns a feeling into a figure, and the plant we rebuilt so a fruit could once again make its own.

The molecule the body is listening for

Umami is the taste of **L-glutamate** — the amino acid, free and dissolved, not the same glutamate locked inside an intact protein. That distinction is the whole axis in a single sentence, and we will return to it, because it is also the reason the taste disappears and the reason it can be grown back.

Free L-glutamic acid is a small, three-carbon-backbone molecule carrying two carboxyl groups and one amine; at the mild pH of the mouth it sits as the glutamate anion, its two negative charges and its amine arranged in exactly the geometry a particular receptor pocket was shaped to close around. It is among the most abundant amino acids in all living tissue, which is the first clue to why umami feels like *food itself*: the body is tuned to detect the molecular signature of protein that has already begun to break down into the building blocks it can use. Umami is, at bottom, the tongue reporting *digestibility* — the chemistry of patience, of foods given the time to come apart into savor.

The map of where that free glutamate pools in the natural world is the map our design is drawn against. It is extraordinary in aged cheese and in kombu kelp, high in a vine-ripened tomato, in dried shiitake, in miso and soy and other long ferments, and vanishingly low in most fast-grown, early-picked produce — which is precisely the gap the work exists to close. Nothing here is exotic chemistry. It is the chemistry of time, and our whole contribution is to write a little

of that patience back into the gene, so a plant can accumulate its own free glutamate on the vine the way a slow-grown one always would.

The receptor — a Venus flytrap on the tongue

Glutamate is read by a pair of class-C G-protein-coupled receptors, the same structural family that reads sweetness, re-pointed at savor. The principal umami sensor is the **TAS1R1·TAS1R3 heterodimer** — and the detail worth holding is that TAS1R3 is the subunit *shared* with the sweet receptor, the same protein partnered differently: pair it with TAS1R2 and the tongue tastes sugar; pair it with TAS1R1 and the tongue tastes broth. Alongside it sits a taste-specific, truncated form of **mGluR1,** the metabotropic glutamate receptor borrowed from the brain's own glutamate signaling and re-expressed on the tongue as a second, lower-affinity confirmer of the same molecule.

Each TAS1R subunit carries, on its outside, a large bilobed extracellular domain that biologists call a **Venus flytrap** — two lobes hinged like a clamshell around a central cleft. Glutamate settles into the cleft of the TAS1R1 flytrap and the two lobes close over it. That closure is the whole signal. It travels down through a cysteine-rich linker into the receptor's seven-pass transmembrane bundle, changes the shape of that bundle on the inside of the cell, and switches on the G-protein waiting there — gustducin. From that point the cascade is the familiar one of taste: gustducin frees its subunits, phospholipase C-β2 fires, inositol trisphosphate floods from intracellular stores, the **TRPM5** channel opens, the cell depolarizes, and a bolus of ATP is released through the CALHM1 channel onto the afferent nerve fiber waiting beneath. A molecule has become a message. The tongue has read a word.



A tomato in cross-section — the locular jelly around the seeds is where free glutamate concentrates as the fruit ripens.

The flytrap, like any clamshell, does not sit still. It lives in a restless equilibrium, flickering between open and closed even with glutamate in the cleft, and a flickering trap is a quiet signal — the receptor spends only part of its time in the active, closed state. Holding that picture in mind matters, because the single most important fact about umami, the thing that makes it feel less like a taste than like a *depth,* is a trick of chemistry that props the flytrap shut. We will come to it. First, the molecule has to be free.

Only the free form is tasted

Here is the fact that governs everything downstream: only *free* glutamate is tasted. The identical amino acid, bound inside an intact protein chain, is silent — the flytrap cannot close around a glutamate that is still stitched into a backbone of other residues. This is why fresh milk is bland and a wheel of cheese aged two years is a quiet bomb of savor; why a just-pressed soybean tastes of almost nothing and the same bean fermented into miso tastes of the sea. Time, enzymes, and microbes do one thing across all of these: they break long proteins down into their constituent amino acids, liberating the glutamate that was bound in them all along.

Seen through this axis, cooking and curing and fermenting and ripening are not four different acts but one — the freeing of bound glutamate into its tasted, free form. A long simmer does it with heat. A ferment does it with microbial proteases. Ripening does it with the fruit's own enzymes, working in the last warm weeks on the vine. Aging does it with patience alone. Every culture's deepest savory cooking — the slow stock, the cure, the ferment, the sun-finished fruit — is really the same unwitting chemistry performed in a hundred kitchens: *break the protein, free the glutamate, let the flytrap close.*

This is also, precisely, what a modern supply chain undoes. A tomato picked green so it will survive a thousand miles, gassed to color in a warehouse, bred for shelf life and not for the tongue, never spends those last warm weeks on the vine — so the enzymatic ripening that would have freed its glutamate never finishes, and the fruit arrives with its savor still locked away, or never built at all. The taste did not fail. It was interrupted.

Giving the feeling a number

To design umami we first had to give it a number, and giving a feeling a number is the work the whole Index was built to do. The rule is the same on every axis:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the midpoint of the dose-response curve. It is *inverse* loudness: a molecule that half-activates at a low concentration is loud, because a little of it speaks clearly, and a low denominator makes the score large. Invert the candidate's EC_{50} against a fixed reference, scale by a thousand, and you have a number a tomato, a broth, and a pepper can all be read against on one page.

For umami we fix the reference as a standardized L-glutamate solution whose half-maximal activation of the human TAS1R1·TAS1R3 receptor we anchor at **EC_{50}(reference) = 1.5 mM** — a value in the low-millimolar range that recombinant umami-receptor assays return for glutamate alone. By construction, that reference scores exactly 1,000. The receptor itself obeys a Hill relation,

$$R([G]) = R_{max} \cdot [G]^n \div (EC_{50}^n + [G]^n)$$

where R is response, [G] the free-glutamate concentration, and **n** the Hill coefficient — the steepness, the measure of cooperativity. At [G] = EC_{50}, the bracket collapses to one-half and R = R_max⁄2 by definition; for glutamate alone n runs about 1.2 to 1.5, a gently cooperative curve.

A food, of course, has no single EC_{50} — it is a mixture, not a reagent. So we express a food's savory drive as an *effective* glutamate-equivalent EC_{50}: the dilution of its expressed juice, read in glutamate terms, that half-activates the receptor. Define a dimensionless **loudness multiplier M = EC_{50}(reference) ÷ EC_{50}(candidate),** and the whole formula collapses to a statement you can hold in your head:

$$IFI\text{-}U = 1{,}000 \times M$$

M has two independent parts, and keeping them separate is the key to the whole axis: **M = M_conc × M_syn** — how much free glutamate the food delivers relative to the reference, multiplied by how far the food's ribonucleotides shift the receptor's EC_{50} leftward. Concentration and synergy, two levers on one number.

Work the concentration lever first, on the tomato that defines this chapter. A commercial, early-picked tomato carries on the order of **20 mg of free glutamate per 100 g.** Convert it: 20 mg/100 g is about 0.20 g/L of juice; divided by the molar mass of free L-glutamic acid, 147.13 g/mol, that is 0.20 ÷ 147.13 = **1.36 mM** in the neat juice. Mastication and saliva dilute that bolus roughly three- to fourfold before it reaches the receptor field, delivering about **0.35 mM** — well *below* the 1.5 mM reference EC_{50}, so the receptor sits under half-max, and the taste is thin. Its effective EC_{50} reads 5.0 mM, so M = 1.5 ÷ 5.0 = 0.30, and **IFI-U = 300.**

Now a tomato designed to accumulate its own glutamate on the vine, at roughly **280 mg per 100 g.** The same arithmetic: 2.8 g/L ÷ 147.13 = **19.0 mM** neat; diluted fourfold, about **4.8 mM** at the receptor — several-fold *over* the EC_{50}, comfortably past half-max, loud. Effective EC_{50} reads 0.58 mM, M = 1.5 ÷ 0.58 = 2.6, and **IFI-U = 2,600.** Between the two tomatoes, free glutamate falls from about 280 to about 20 mg/100 g — a loss of roughly **93 percent** of the fruit's savor, the exact gap the restoration closes.

SAVORY STIMULUS	FREE GLUTAMATE (MG/100 G)	EFFECTIVE EC_{50} (MM)	LOUDNESS M	IFI-U
Reference glutamate solution	—	1.50	1.00	1,000
Supermarket tomato	~20	5.00	0.30	300
Vine-ripe heirloom tomato	~250	0.65	2.30	2,300
Restored tomato (GAD / GDH1)	~280	0.58	2.60	2,600
Phoenix · SS-0096 (umami floor)	—	0.83	1.80	1,800

The table is the axis made legible. Every effective EC_{50} in it is simply 1.5 mM ÷ M, and every IFI-U is 1,000 × M — the same two operations, repeated, which is the point of giving a flavor a number

at all.

Synergy — cooperativity you can taste

The restored tomato gets to 2,600 almost entirely on the concentration lever. The second lever, synergy, is where umami stops behaving like an ordinary taste, and it is cooperativity made literal — a Hill coefficient you can put on your tongue.

On their own, the **5′-ribonucleotides** are very nearly tasteless. Inosine monophosphate, **IMP,** is liberated when aged fish and meat break down and was found in dried bonito by Shintaro Kodama in 1913; guanosine monophosphate, **GMP,** is richest in dried shiitake, pinned there by Akira Kuninaka in 1960. Taste either one alone and you taste almost nothing. Kuninaka's real discovery was not the molecules but their arithmetic: a ribonucleotide and glutamate do not *add,* they *multiply.* Put them together and the savor leaps far past their sum.

The mechanism is the flickering flytrap, propped shut. The ribonucleotide does not bind in the cleft where glutamate goes; it binds at the **hinge** on the outside of the closing joint, and there it acts as a wedge — carrying no taste of its own, it stabilizes the closed, active state so that every glutamate that binds *stays* bound and active far longer than it otherwise would. The practical effect is a collapse of glutamate's EC_{50}, a leftward lunge of the entire dose-response curve. In the dilute regime of real cooking the receptor-level shift plateaus near **fifteenfold** — EC_{50} falling from 1.5 mM toward roughly 0.10 mM. That is M_syn ≈ 15 on its own.

Run it through the formula and the ceiling appears. A food at merely reference glutamate concentration but full ribonucleotide synergy would compute to IFI-U = 1,000 × 15 = **15,000** — which the scorecard clips at its per-axis normalization maximum of **5,000.** The naturally loudest savory foods and any well-matched glutamate-plus-nucleotide pairing already sit at or over that ceiling, and we state this plainly rather than chase it: the restored tomato at 2,600 and Phoenix at 1,800 are *deliberate, moderate* readings — savor you notice, not savor that shouts.

The kitchen has always known this law empirically, long before it had the receptor. The perceived umami of a glutamate-and-nucleotide mixture follows a relation the flavor chemist Shizuko Yamaguchi fitted to taste-panel data:

$$\gamma = u + \beta \cdot u \cdot v$$

where u is glutamate concentration, v is nucleotide concentration (both in g/dL), γ is the glutamate-equivalent intensity, and β ≈ 1,218 for IMP. Work a faint dashi: kombu contributing u = 0.03 g/dL of glutamate, a shaving of bonito contributing v = 0.02 g/dL of IMP. Then γ = 0.03 + 1,218 × 0.03 × 0.02 = 0.03 + 0.73 = **0.76 g/dL** — the two faint inputs together taste like twenty-five times the glutamate of the kombu alone. (The perceptual leap runs larger than the fifteenfold receptor-occupancy shift because perception integrates intensity well above half-max; the Index, by design, scores the receptor event, not the sensation.) It is the same molecular wedge, written once in the grammar of a stockpot and once in the grammar

of a dose-response curve. This is why a few grams of dried mushroom remake a whole pot of sauce, why dashi marries kombu's glutamate to bonito's inosinate, and why the Index scores such a pair as a single reading rather than two.

Where glutamate pools

Because the design is drawn against the natural map, it helps to see that map in numbers. Free glutamate is not rare — it is concentrated by patience, and it travels with its synergists in exactly the foods that have earned their depth.

SOURCE	FREE GLUTAMATE (MG/100 G)	5′-NUCLEOTIDE (MG/100 G)	ROLE IN THE BROTH
Kombu (dried kelp)	1,600–3,200	—	the glutamate benchmark
Aged cheese (24 mo)	~1,200	—	long aging frees bound amino acids
Dried shiitake	~1,060	GMP ~150	carries both halves at once
Katsuobushi (bonito)	~40	IMP ~500	almost pure synergist
Ripe tomato	140–250	—	the everyday vine glutamate

Two things fall out of the table. The loud natural foods — kombu, aged cheese, dried shiitake — carry so much free glutamate that, read through the formula, they would land at or past the 5,000 scorecard ceiling; savor, in nature, is capped not by chemistry but by patience and by how much a food can hold. And the table explains every great savory tradition at a glance: pair a glutamate source with a nucleotide source and the wedge does its work — kombu with bonito in a dashi, tomato with dried mushroom in a sauce, a long-aged cheese grated over either. The oldest savory cooking of every culture is ribonucleotide synergy, performed by cooks who never needed the word.

The tomato, restored



The tomato, restored — the savor a modern supply chain strips out, grown back in through the plant's own glutamate route.

The tomato is this axis's standard-bearer and, with **2,040 constructs,** one of the most developed organisms in the whole library. A tomato is supposed to taste like being cared for, and the one a long-haul truck delivers has had most of that taken from it —

picked unripe to survive the journey, gassed to color, bred for shelf life and not for the tongue, until the glutamate that would have accumulated on the vine in its last warm weeks never builds.

Our tomato restores it, and the *how* is the part that matters most to us. We did not add glutamate to the fruit from outside. We gave the fruit back its own ability to make it — by tuning the plant's native machinery, the **GAD** and **GDH1** enzymes that govern the synthesis and interconversion of glutamate, so the fruit accumulates free glutamate on the vine the way a slow-ripened one always would. Glutamate decarboxylase and glutamate dehydrogenase are not foreign imports; they are the tomato's own, present in every fruit, merely tuned here to finish the job that early picking cuts short. That distinction — restoring a capacity rather than dosing a compound — is the difference between seasoning a food and *healing* one, and it is the whole ethic of the library in a single fruit.

The restoration is also, quietly, an engineering story about *variance,* not just mean. Because IFI-U = 1,000 × M is linear in free-glutamate concentration with synergy held fixed, the score inherits the concentration's coefficient of variation directly: a reading's uncertainty is the glutamate's uncertainty. A field tomato's free glutamate swings hard with picking ripeness and season — call it a 25–30% CV — so a nominal 2,300 reading carries roughly **±650.** The restored line is designed to hold free-glutamate accumulation to the platform's **≤7% batch CV,** so 2,600 carries only about **±180.** For a formulator, the tightened interval is the real prize: you can build a sauce around 2,600 ± 180 and never around 2,300 ± 650. The engineering is designed to buy a higher mean and a narrower band in the same stroke.

We are candid about the limits, too. The IFI-U number reports receptor loudness at the moment of tasting; it does not predict how a tomato's savor will survive a ninety-minute reduction, where water leaves, glutamate concentrates, and the reading climbs — nor how it behaves against the acidity and sugar that share the fruit. Umami on the tongue is one axis read clean; umami in a finished dish is a composition. The number is the floor we design toward, not the ceiling a good cook will reach.

The harder hosts — soy, koji, yeast, mushroom



Soybean — the frontier of built savor, where microbial proteases free the glutamate a long ferment would.

The tomato is the easy host because its glutamate route is short and its fruit is built to accumulate. The rest of the umami library is harder, and honestly so — each organism reaches savor by a different road.

ORGANISM	CONSTRUCTS	UMAMI ROUTE	DIFFICULTY
Tomato	2,040	free glutamate (GAD / GDH1)	Easy
Soybean	360	glutamate / peptides	Frontier
Tea	280	theanine-linked savory	Hard
Yeast	240	glutamate + ribonucleotides	Hard
Koji Mold	80	protease-liberated glutamate	Hard
Shiitake	80	GMP synergy	Hard
Mushroom	40	GMP synergy	Hard

Soybean sits at the frontier because the bean guards its glutamate inside storage proteins so stable that, historically, only months of fermentation would free them; the work there is about liberating bound glutamate and short savory peptides earlier and more reliably than a vat of microbes ever could. Koji mold — *Aspergillus,* the quiet engine under miso and soy sauce — is itself a protease factory, and tuning it is tuning the speed and cleanliness of that liberation. Yeast is the rare host that carries *both* halves of the synergy at once, glutamate and the ribonucleotides together, which is why a yeast extract reads so loud. Shiitake and the other mushrooms are the GMP specialists, the nucleotide wedge in living form — grown not for their own glutamate but for the synergist that multiplies everyone else's. Tea reaches savor by a different molecule entirely, the amino acid theanine, a reminder that the axis is defined by the receptor, not by any one compound.

Read as a group, the table is a map of difficulty honestly drawn: the plant that stores its savor loosely is Easy, the bean that locks it behind fermentation-grade proteins is Frontier, and the number on each line is a promise of how much of that organism we have actually built, not how much we wish we had.

Savor as a floor

Umami does one more thing in the design, and it is why the number sits under even the fiercest work in the library. Phoenix — the full-spectrum superhot that registers on all eight axes at once — carries umami at **1,800,** and it carries it on purpose. The family would not build a pepper, however furious, that was hollow underneath. Heat without savor is an assault; heat standing on a floor of umami is a *pleasure,* because the tongue is told, beneath the fire, that this is real food.

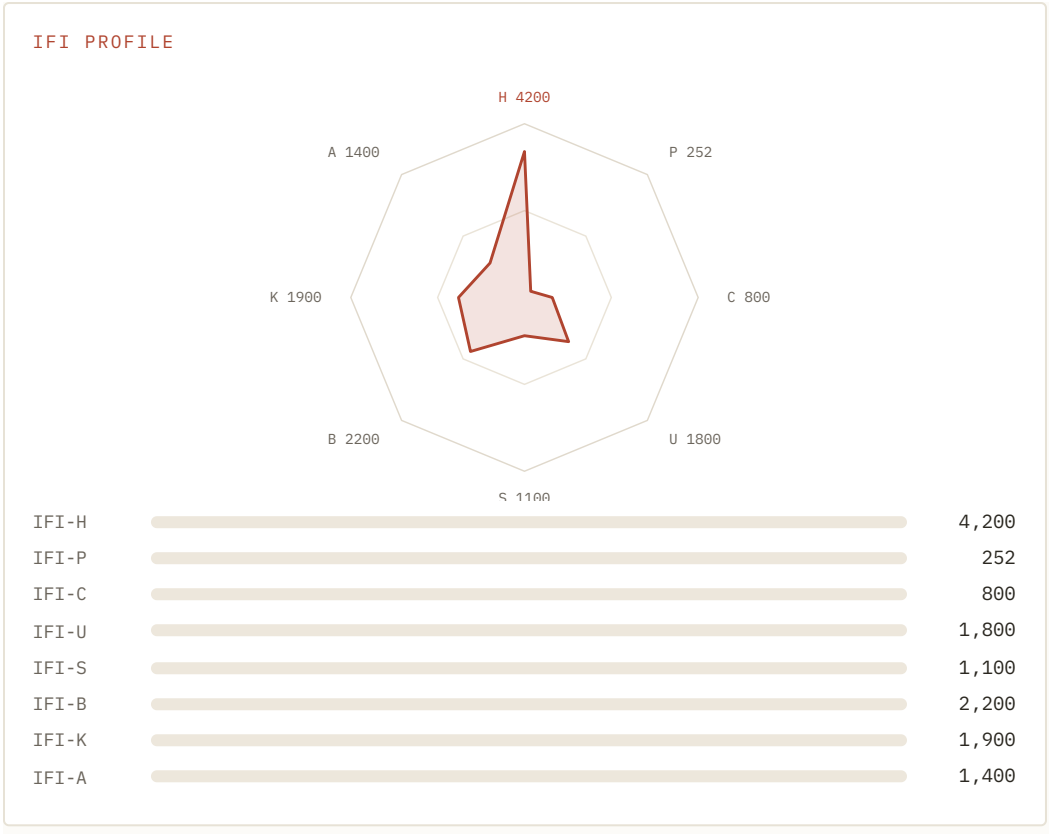

Read the profile and the design intent is plain: the heat is enormous at 4,200, but the umami at 1,800 and the kokumi at 1,900 run beneath it like a bass line — the savory-depth chord holding the fire up. Umami and kokumi, the richness-and-continuity axis read by the calcium-sensing receptor, pair into the single most requested composition we make: the taste of a kitchen that loves you, written down at last as a pair of numbers. We put that chord under the hottest thing we grow precisely because a flavor that only burns is a flavor that only ends.

One last property makes umami quietly powerful in the kitchen and the factory alike: it can stand in for a share of the salt a dish would otherwise need. A food with real savory depth reads as seasoned and complete at a lower level of added sodium, because umami and saltiness reinforce one another on the palate — the mouth, given depth, simply asks for less of the sharp mineral edge to feel that a dish is finished. We state this as a formulation fact and nothing more, never a health claim: a tomato designed to carry its own savor lets a sauce taste complete with less salt stirred in. On the Index it shows as a high IFI-U doing work that salt used to do — more of what nourishes, less of what merely corrects, which is as plain a statement of the whole library's purpose as any number can make.

What the number cannot do

We end on the honest edge of the measurement, because a number is a form of respect only if you also say where it stops. IFI-U is a clean reading of one receptor event — glutamate closing a flytrap, a ribonucleotide holding it shut — scored at the instant of tasting. It does not capture how savor blooms across a long reduction, how it negotiates with acid and sweetness in a finished sauce, or the slow, whole-mouth sense of *nourishment* that is umami's deepest gift and that no single half-maximal concentration will ever fully hold. The arithmetic is exact; the experience is larger than the arithmetic, and we would rather admit that than pretend the figure is the feeling.

What the number *can* do is the thing ten thousand years of cooking could not: it lets us say *how much* — how much savory depth a fruit holds, how far short of its own birthright a trucked tomato falls, how

precisely a vine can be designed to carry its glutamate home. We did not build the Index to replace the comfort of a good broth. We built it so that comfort could be grown back into food that had lost it — reproducible, and real, a quiet report on the tongue that says, below language, *this is the food the body was hoping for.*

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

Chapter 8 · Zero-Bitter — *subtraction as a flavor.*

Every other axis in this book is an addition — more heat, more sweetness, more of the round savory depth that makes a broth taste like being cared for. This one is a subtraction, and it is the most quietly radical thing the Index measures: a flavor defined not by what it brings but by what it takes away. Bitterness is the taste the body distrusts on purpose, the one sensation evolution tuned toward suspicion rather than pleasure, and it stands guard in front of some of the most nourishing foods a kitchen holds — the deep-green brassica, the roasted bean, the dark honest cacao. The whole point of IFI-B is to turn that guard down by a defined, bounded amount, so the coffee and the kale and the chocolate can be met on their own terms, their depth intact and their harshness eased. We did not set out to erase bitterness. We set out to *dose* it.

That distinction is the entire chapter. An addition can be pushed; you can always reach for one more unit of heat. A subtraction has a floor — a point past which a food stops being itself and becomes a hollowed imitation of what it was. So the axis is written as a span, not a promise: between twenty and sixty-five percent quieter, against the untreated baseline, and never further. Zero-Bitter does not mean *no* bitter. It means bitter *chosen,* set exactly where it belongs, because the family that grew up eating these plants — vegetarian for as long as any of us can remember, reverent toward the green things on the plate — never wanted them silenced. We wanted them made easier to love.

The alarm panel

Bitterness is read by the **TAS2R family** — roughly twenty-five bitter receptors in the human mouth, all of them class-A G-protein-coupled receptors, compact seven-helix proteins threaded through the membrane of the taste cell. Where the sweet and umami receptors are fussy, choosy heterodimers that answer to a narrow list of ligands, the bitter receptors are deliberately, protectively promiscuous. The job of detecting poison rewards breadth over precision: better to flinch at a thousand harmless molecules than to miss the one alkaloid that kills. So these twenty-five receptors are not twenty-five keys for twenty-five locks. They are an alarm panel. One receptor answers to dozens of chemically unrelated compounds; one compound often trips several receptors at once; and the brain, receiving a whole panel lit at different intensities, flattens the whole chord into a single verdict — *bitter, be careful.*

Knowing the wiring behind that verdict is what lets us design against it, because the quieting has to happen somewhere specific along the chain. The cascade runs in order. A bitter molecule settles into the orthosteric pocket of a TAS2R and nudges the seventh helix; the receptor's intracellular face opens and engages **α-gustducin,** the taste cell's dedicated G-protein, splitting it into its Gα and $G\beta_3\gamma_{13}$ parts. The freed Gβγ switches on **phospholipase C-β2,** which cleaves the membrane lipid PIP_2 into diacylglycerol and **inositol**

trisphosphate. The IP_3 crosses to the endoplasmic reticulum, opens the IP_3-receptor channels there, and spills stored **calcium** into the cell. That calcium gates **TRPM5,** a sodium-permeable channel; sodium floods in; the cell depolarizes; and the depolarization opens **CALHM1/3** channels that release ATP onto the afferent fibers of the facial and glossopharyngeal nerves. Only then does *bitter* become a signal the brain can read.

Eight steps, and every one of them is a place to intervene. You can keep the molecule out of the pocket. You can stop the pocket from engaging gustducin. In practice the two tools worth having both act at that very first handshake — the molecule and the receptor — because everything downstream is shared machinery the cell also uses for sweet and umami, and a family that cares about a food tasting like itself does not go tearing out the wiring the rest of flavor runs on. We ease the alarm at its sensor, and we leave the cell whole.



FIG. 9 — Coffee, the world's most beloved bitterness. Attenuation keeps the aromatics the roast builds and eases only the harsh edge they travel with.

Why bitterness is personal

The anchor of the axis is **TAS2R38,** the receptor for the thiourea compounds PTC and PROP that genetics classes use to sort a room into "tasters" and "non-tasters." Both molecules carry the same N—C=S thiourea group, and the whole difference between a person who gags on a PROP strip and one who tastes nothing comes down to three positions in the receptor's sequence — residues 49, 262, and 296. The taster carries proline, alanine, and valine there, a haplotype written **PAV;** the non-taster carries alanine, valine, and isoleucine, written **AVI.** Those three swaps change the shape and electrostatics of the binding pocket just enough to alter how tightly the bitter molecule is held, and which pair you inherited decides whether raw broccoli reads as pleasantly green or aggressively

acrid. A minority inherit PAV on both chromosomes and sit at the extreme — "supertasters," often with a denser forest of fungiform papillae on the tongue — for whom coffee, kale, and tonic water are genuinely, measurably harsher than they are for anyone else.

This is humbling, and it shapes how the axis has to be scored. Bitterness is read on a differently calibrated dial by different mouths; a single panel's verdict is a moving target. So we never anchor IFI-B to a tasting panel's mood. We anchor it to a **fixed chemical reference** — a defined bitter stimulus at a defined concentration activating a defined receptor in a cell-based assay — and we report attenuation as a computed *reduction* against that reference, not as a felt intensity that would be one number for a supertaster and another for everyone else. The honest unit on this axis is a ratio of receptor activation, because the receptor, unlike the person, holds still.

Two ways to quiet an alarm

You can turn an alarm down two ways, and the Index keeps them distinct because they are different acts of engineering.

You can **lower the ligand** — reduce the concentration of the bitter compound the plant makes, so fewer receptors are tripped in the first place. On a dose-response curve this is simply less stimulus climbing the same axis: you slide the operating point down the curve toward its foot. This is the cleaner tool whenever the bitter principle is a discrete, droppable thing that is not also carrying the flavor you want to keep.

Or you can **antagonize the receptor** — place a molecule in or near the orthosteric pocket that occupies it without triggering the conformational change that engages gustducin, so the bitter compound, though fully present, is less able to activate the sensor. A competitive antagonist shifts the whole dose-response curve rightward by a factor of (1 + [B]/K_B), where [B] is the antagonist concentration and K_B its dissociation constant. The stimulus is still in the mouth; it simply speaks more softly at every dose. This is antagonism run deliberately in reverse — a curve pushed to the right on purpose.

That receptors can be *blocked,* and not merely starved of their ligand, is established pharmacology we build on rather than invent. Some antagonists are as ordinary as particular sodium salts, which have long been known to mask bitterness at the periphery; others are purpose-found compounds — probenecid quiets TAS2R16, and a small synthetic fatty acid, GIV3727, blocks the TAS2R43/31 pair — each raising the apparent half-activation of a bitter compound so it registers fainter at the identical dose. Choosing between lowering the ligand and blocking the receptor is most of the craft of this axis. In coffee, where the bitter principles are born in the roast alongside the very aromatics that make coffee smell like coffee, antagonism is the gentler tool; it eases the harsh edge without stripping the cup. In a brassica, where the bitterness is a discrete glucosinolate pool the plant builds and stores, a measured reduction of the ligand is

cleaner. And because the two routes can produce the identical felt result, the Index scores them identically — by how much quieter, against baseline, and by nothing else.

The arithmetic of a rightward curve

The master relationship behind every axis in this book is a single ratio of half-activations:

```
IFI(candidate) = 1000 × EC₅₀(reference) / EC₅₀(candidate)
```

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the pharmacologist's measure of how *loudly* a compound speaks, low meaning potent. On the loudness axes the reference sits on top: a candidate that shouts has a smaller EC_{50} than the reference, the quotient climbs, and the score climbs with it.

Zero-Bitter is the one axis where the achievement is a *larger* EC_{50} — a receptor made harder to reach, a curve pushed right — so the same ratio is read in mirror. The reference is the native, untreated bitterness; the candidate is the eased food; and the score rewards the rise:

```
IFI-B = 1000 × EC₅₀(eased) / EC₅₀(native) = 1000 × (dose ratio)
```

The dose ratio, DR, is the factor by which the eased food's bitter curve has shifted rightward — the number of times more stimulus it now takes to reach the same half-activation. A score of 1,000 is no change at all (DR = 1). Our model assigns **Phoenix** a Zero-Bitter score of **2,200,** which reads directly as a dose ratio of **2.2:** its bitterness is designed to need 2.2 times the concentration to be as loud as it was, or equivalently to speak at about 1/2.2 ≈ 45% of its former loudness at matched dose. Subtraction, expressed as multiplication.

To connect that score to the plain-language promise — "twenty to sixty-five percent quieter" — we have to model the response itself. Take the Hill form with a coefficient of one, a fair approximation for a single bitter receptor near threshold:

```
R([L]) / R_max = [L] / (EC₅₀ + [L])
```

Suppose the native food's bitter principle sits right at its own half-activation, $[L]_0 = EC_{50}$, so the untreated response is 50% of maximum. Apply a dose ratio DR (whether by antagonism or by the equivalent drop in effective ligand), and the attenuation felt at that operating point works out cleanly:

```
A = 1 − R_after / R_before = (DR − 1) / (DR + 1)
```

Run the band through it. The whole engineering range of the axis — 20% to 65% attenuation — maps to a tidy ladder of dose ratios and scores:

ATTENUATION (AT HALF-MAX POINT)	DOSE RATIO DR	IFI-B
20%	1.50	1,500
30%	1.86	1,857
37.5% · Phoenix	**2.20**	**2,200**
50%	3.00	3,000
60%	4.00	4,000
65%	4.71	4,714

Sixty-five percent sits just under the per-axis scorecard ceiling of 5,000, which is exactly where the honest ceiling belongs: the axis runs out of room at the same place the food runs out of room to be eased.

There is a subtlety the table cannot quite show, and it is the most important honest limit on this axis. The percentage a mouth *feels* depends not only on the dose ratio but on where the native bitterness sits on its own curve. Hold Phoenix's dose ratio fixed at 2.2 and move the operating point, and the felt easing slides with it: a mild, near-threshold bitter at $[L]_0 = 0.3 \cdot EC_{50}$ eases by 48%; bitterness sitting right at half-activation eases by the 37.5% the score names; a saturating dark-roast harshness at $[L]_0 = 3 \cdot EC_{50}$ eases by only 23%. The same engineered shift reads as nearly half-off on a gentle bitter and under a quarter-off on one that is already saturating the receptor. This is why the Index scores the dose ratio and not the raw percentage: DR is a property of the curve, fixed and portable from food to food, while the felt percent is a property of the serving. We report the thing that holds still and let the kitchen read the percentage off its own operating point.

Two conversions keep the assay honest when it meets the real world, because bitter principles arrive in the units of their industries, not of pharmacology. Beer bitterness is sold in International Bitterness Units, where one IBU is about one milligram of iso-α-acid per liter; a firmly hopped IPA at 65 IBU is therefore carrying 65 mg/L of isohumulone, which at a molar mass of 362 g/mol is 65 ÷ 362 ≈ 0.18 mmol/L, or **180 µM** — the number the receptor actually counts. Brewed coffee runs near 0.4 g/L of caffeine, which at 194 g/mol is about **2.1 mM,** roughly twice caffeine's ~1 mM bitter-recognition threshold and still only a minor share of the cup's total bitterness. We do the division every time, because a receptor does not read a label.

Finally, the error. An axis that promised 37.5% but delivered anywhere from 10% to 60% batch to batch would be worthless, so the attenuation is held to the platform's reproducibility standard of **≤7% coefficient of variation** — against the 42% CV an uncontrolled process would leave. Propagated through the score, IFI-B = 1000 × DR inherits that 7% directly: 2,200 ± 154. Propagated through the felt percentage, with $\sigma_A = [2/(DR+1)^2] \cdot \sigma_{DR}$, a dose ratio of 2.2 ± 0.15 yields an attenuation of 37.5% ± 3.0 percentage points — tight enough that a roaster or a grower can trust the same construct to

land the same easing cup after cup. Reproducibility is not a footnote on a subtraction axis. It is the whole difference between a dosed bitterness and a gamble.

The bitter principles, plant by plant

The alarm fires at specific molecules, and knowing which molecule a given food uses to ring it is half the design. Each is a different lever, pulled a different way.

In the **brassica** — kale, cabbage, broccoli, the whole cruciferous family a vegetarian table leans on — the bite is **glucosinolates,** sulfur-rich compounds the plant stores intact and inert. The moment a leaf is torn or chewed, the stored glucosinolate meets the enzyme **myrosinase,** kept until that instant in a separate compartment; myrosinase hydrolyzes the thioglucoside bond, releasing glucose and an unstable aglycone that rearranges into a pungent **isothiocyanate** plus sulfate. Sinigrin becomes allyl isothiocyanate; glucoraphanin becomes sulforaphane. The bitterness is both the glucosinolate itself and the harsh edge of its products — and goitrin, from progoitrin, is a direct and intense TAS2R38 agonist. Lowering the glucosinolate pool at the source therefore eases both the stored bite and the downstream harshness in one move. But the same pool carries the isothiocyanates these plants are prized for, sulforaphane among them, which is precisely why the axis eases by a bounded amount and never to zero. We quiet the brassica; we do not disarm it.

In **coffee,** the harshness is largely **chlorogenic-acid lactones,** formed as the roast degrades the green bean's chlorogenic acids, and at darker roasts the **phenylindanes** that form as those lactones degrade further — the source of the lingering, hard bitterness a very dark cup leaves behind. Caffeine rides alongside as a mild bitter contributing perhaps a tenth of the total. Because the lactones are born in the same roast reactions that build the aromatics, they cannot be dropped without dropping the smell of coffee with them — so coffee is an antagonism food, eased at the receptor while the aromatic compounds it travels with are left untouched.

In **cacao,** the bitterness is **theobromine** — the methylxanthine cousin of caffeine — carried on a freight of polyphenols and procyanidins that are themselves both bitter and astringent, plus diketopiperazines built during the roast. In **hops,** it is the **iso-α-acids,** the isomerized humulones a brewer prizes and pays for by the IBU. In the **potato,** it is the **glycoalkaloids** α-solanine and α-chaconine, bitter and mildly toxic together, which is the rare happy case where easing the flavor and improving the food are the same edit.



FIG. 11 — Cacao. Theobromine on a freight of polyphenols — a bitterness eased by both levers at once, the body kept and the harsh edge pressed down.

BITTER PRINCIPLE	WHERE IT LIVES	HOW IT RINGS THE ALARM	PREFERRED LEVER
Glucosinolates (sinigrin, glucoraphanin)	Kale, cabbage, broccoli, radish, rapeseed	TAS2R38 + isothiocyanate harshness via myrosinase	Lower the ligand
Chlorogenic lactones, phenylindanes	Coffee	Several TAS2Rs; born in the roast with the aromatics	Antagonize the receptor
Theobromine + polyphenols	Cacao	Methylxanthine bitter + procyanidin astringency	Blend of both
Iso-α-acids	Hops, beer	TAS2R1 / TAS2R14 family	Lower the ligand
Glycoalkaloids (α-solanine, α-chaconine)	Potato	Bitter and mildly toxic together	Lower the ligand

A worked case shows how plainly the ligand lever reads. Take a potato whose total glycoalkaloids sit at a bitter 180 mg/kg, and suppose the half-activation of the bitter response is near 90 mg/kg. The untreated response is 180 / (90 + 180) = 0.667 of maximum. Engineer the construct to hold glycoalkaloids to 60 mg/kg — well below both the bitter threshold and the safety limit — and the response falls to 60 / (90 + 60) = 0.400. The attenuation is 1 − 0.400/0.667 = **40%,** squarely inside the band, achieved without touching a single receptor, and carrying a safer tuber out with the gentler one. Lower the ligand, and the curve comes to you.

The attenuation ladder

The axis is not one subtraction but a family of them, chosen plant by plant, and the roster of hosts reads as a ladder of difficulty. The easy rungs are the foods whose bitterness is a discrete, droppable pool; the hard rungs are the ones where the bitterness is braided into the identity or the aromatics and must be eased at the receptor instead.

ORGANISM	CONSTRUCTS	BITTER PRINCIPLE	DIFFICULTY
Grape	1,610	polyphenols / tannins (with astringency)	Hard
Broccoli	380	glucosinolates	Moderate
Tea	280	catechins (with umami, astringency)	Hard
Broccoli-Kale	280	glucosinolates	Hard
Orange	120	limonoids, flavanones	Hard
Potato	80	glycoalkaloids	Easy
Hops	80	iso-α-acids	Moderate
Coffee	40	chlorogenic lactones, caffeine	Hard
Cabbage	40	glucosinolates	Moderate
Kale	40	glucosinolates	Moderate
Radish	40	glucosinolates	Hard
Rapeseed	40	glucosinolates	Hard

Broccoli, at 380 constructs, is the flagship of the axis — the largest dedicated Zero-Bitter program in our design library, an approach designed to bring a beloved, slightly feared vegetable to the exact point where a child will eat it without stripping the compounds that make eating it worth doing. **Coffee,** at 40, is the hardest of the small programs precisely because its bitterness is inseparable from its perfume; every construct there is an exercise in antagonism, in quieting a receptor while leaving the aromatic chemistry of the roast completely alone. And **broccoli-kale,** a cruciferous hybrid at 280, is hard for the opposite reason — two glucosinolate systems to balance at once, each with its own myrosinase kinetics, each needing to be eased without flattening the green that is the whole point of the leaf.



FIG. 10 — Broccoli, the flagship of the axis at 380 constructs. Eased to the point a child will eat it, never stripped of the glucosinolates that make it worth eating.

The honest edges

Every axis has a place where it fails, and a subtraction axis wears its failures on the outside.

The first is the floor already named: push past roughly sixty-five percent and a food stops tasting like itself. A coffee eased too far goes hollow — the bitterness it loses turns out to have been holding up the body of the cup, and what remains is thin and vaguely sweet and not coffee. A kale eased past recognition tastes of nothing,

improved into blankness. The axis is bounded not by caution but by chemistry: beyond a point, the bitter compound you are quieting *is* the food, and removing it removes the thing.

The second is the operating-point dependence the arithmetic already exposed. Because the felt percentage depends on where a food's native bitterness sits on its curve, the same construct lands differently in a light roast than a dark one, in a young leaf than an old one, in a mild supermarket potato than a bitter garden one. We score the dose ratio because it is portable; the kitchen still has to meet the construct where it lives.

The third is the supertaster. A fixed-reference score is honest about the receptor, but it cannot promise the same experience to every mouth, because a PAV/PAV supertaster begins from a louder baseline and will still find an eased kale sharper than a non-taster finds the untreated one. We state the reduction we engineered; we do not pretend it feels identical to everyone, because that would be a number no single construct could keep.

The fourth is coupling. In foods where the bitter molecule is chemically wedded to the aromatics — coffee most of all — the ligand lever is simply unavailable, and antagonism is the only honest route; try to drop the bitter compound and the smell goes with it. Knowing which lever a food *forbids* is as much of the craft as knowing which it allows.

And the fifth is the one we hold most carefully. The compounds bitterness guards are often the ones a body does well on — the glucosinolate-derived isothiocyanates of the brassica, the polyphenols of cacao and tea. To ease the alarm toward those foods must never become stripping the foods of what made them worth guarding. So the axis is written to quiet the *harshness* a bounded amount, not to remove the chemistry; the bounded range is, among other things, a promise to the plant.

Subtraction with a conscience

There is an evolutionary irony at the center of this axis, and it is the reason the work matters to us. Bitterness evolved to flag the alkaloids and glycosides that once signaled poison, and it still fires at them faithfully. But many of the molecules it flags today are the methylxanthines, polyphenols, and glucosinolates of foods a body thrives on — so the ancient alarm now stands guard, as often as not, in front of the pantry rather than the poison. To ease it by a bounded amount is not to silence a warning the body still needs. It is to quiet a false one it has largely outgrown, and to do it by *just enough* that the food goes on tasting of itself.

We are careful about what that lets us say. We do not say these foods heal. We say they can be designed to be kinder to eat, so the eater meets them more often and more gladly, and we leave the rest to the plant and the body, which have always known what to do with each other. That is the whole of the design philosophy on this axis —

a nourishment ethos expressed not as a claim but as an engineering constraint, bounded always by the rule that a kale eased past recognition is no longer a kale and has been improved into nothing.

Phoenix is designed with a Zero-Bitter score of 2,200 — one of the most prominent values in its eight-axis profile — a superhot designed to have its own alkaloid bitterness pressed down significantly, a dose ratio of 2.2 computed in on purpose, precisely so the fruit, the heat, and the sweetness beneath can be tasted without the rough edge a pepper of that intensity would otherwise carry.

IFI PROFILE

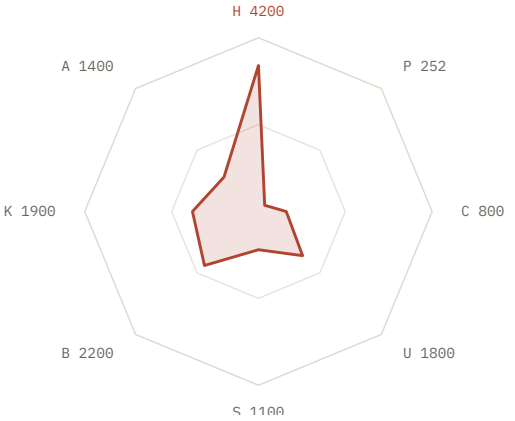

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

Read that profile and you can see the shape of the design: heat the highest computed value, and then, alongside it, a subtraction holding one of the most prominent spokes — a clear signal in our design that taking something away, by exactly the right amount, is one of the most expressive moves a designer of flavor can make. The Index was built to measure addition. It earns its keep on the axis where less, designed and bounded and reproducible, is the whole of the art. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 9 · Kokumi — *the axis you only notice when it's gone.*

This is the axis almost no one can name, and the one everyone has missed. Kokumi is not a taste of its own — it will not make a thing sweet or savory or sharp — and if you search a plain broth for it you will not find it, because what it governs is not a flavor but the *behavior* of flavor: its thickness, its continuity, the way it fills the whole mouth and lingers instead of flashing and fading. A stock with kokumi coats and persists and tastes *complete*; the identical stock without it tastes thin and oddly hollow — correct on paper, lifeless on the tongue. Of all eight axes, this is the one whose discovery as a measurable quantity most changed how our family cooks, because it finally named the thing we had chased for a lifetime by instinct and never been able to add on purpose: the quality a grandmother's cooking had, and a careful, faithful reproduction of it somehow did not.

For most of culinary history kokumi was a rumor the mouth kept and the vocabulary could not hold. You could feel it and you could not point at it, and so it survived only as technique — sweat the alliums, keep the rind, let the pot go long — a set of moves everyone trusted and no one could explain. What follows is the explanation, and then the arithmetic: the receptor that reads kokumi, the molecules that speak to it, the enzymes a plant uses to build those molecules, and the single number that turns all of it into something you can design toward on purpose.

The receptor that never stopped doing its first job

Kokumi is read by the calcium-sensing receptor, **CaSR** — a class-C G-protein-coupled receptor, built on the same architecture as the receptors for sweetness and for umami: a large, bilobed extracellular domain folded like a Venus flytrap, hinged over a cysteine-rich stalk that drops down into a seven-pass transmembrane bundle. CaSR's first and oldest job has nothing to do with dinner. It sits in the parathyroid gland and monitors the calcium concentration of the blood, holding it inside a window so narrow that your nerves and your heartbeat depend on it; when blood calcium rises, CaSR feels the divalent ions settle into the cleft of its flytrap, closes the clamshell, and tells the gland to ease off its hormone. That is the receptor's day job, and it is an ancient one.

The tongue, as it so often does, borrowed the machinery and pointed it at food. The same receptor sits on taste cells, where it reads not calcium but a quiet family of molecules the rest of the palate ignores — and when it is activated there, the sensation it produces is not a taste at all but *fullness*. That one protein should govern both the calcium in your blood and the body of a broth is the kind of economy biology is forever springing on us, and a standing reminder that taste is a young art assembled almost entirely out of older systems. Medicine knew this receptor long before flavor science did: the drug cinacalcet, developed to calm an overactive

parathyroid by making CaSR more sensitive to calcium, turns out to be powerfully kokumi-active on the tongue — a pharmaceutical that, to the mouth, tastes of richness.

The link to kokumi is recent enough to still feel like a frontier being walked. The sensation was named in 1990, when a team led by Yoichi Ueda chased down what made a simple garlic extract lend *body* to foods that held almost none of garlic's own flavor, and traced the effect to its γ-glutamyl peptides — naming a dimension of taste that science had walked past for a century, precisely because it hides inside the others and leaves no fingerprint of its own. The receptor behind it was pinned down later still, and turned out to be an old friend in new clothes: not a novel taste protein but CaSR, the calcium sensor already mapped in exhaustive detail by the people who study bone and blood. Flavor did not have to invent the machine; it only had to find the one already running, and learn to read its dial.

A molecule you cannot taste

The agonists CaSR reads on the tongue are **γ-glutamyl peptides** — short chains carrying one unusual bond. In an ordinary peptide, glutamate links to its neighbor through the carboxyl on its backbone, its α-carbon. In these, the bond forms instead off the carboxyl on glutamate's *side chain*, its γ-carbon: an isopeptide bond, a γ-linkage — chemically an amide like any other, but hung from the wrong arm of the molecule. That small structural quirk has two large consequences. The first is that the ordinary peptidases of the mouth and gut, shaped to cut α-bonds, slide off the γ-linkage and leave it intact; these peptides are *durable*, surviving cooking and fermentation and the first passes of digestion where normal peptides are dismantled. The second is that, by themselves, they are close to tasteless — they register on none of the primary taste axes — so a food can be dense with them and reveal nothing of it to a first sip.

The exemplar is **glutathione** — γ-glutamyl-cysteinyl-glycine, the tripeptide nearly every living cell makes for itself as its own redox buffer, held in quantity by fresh alliums and by yeast. The loudest yet characterized is the tripeptide **γ-Glu-Val-Gly**, roughly an order of magnitude more potent than glutathione, isolated in the hunt for what, exactly, made certain aged and fermented foods feel so full. Between them sits a ladder of lesser agonists — the dipeptide γ-glutamyl-cysteine among them — and a plant's kokumi is, in the end, a question of which rungs of that ladder it climbs and how far.

What these peptides do when they reach CaSR is the subtle part. They are not really agonists that throw the receptor open on their own; they are *positive allosteric modulators*. They bind at or beside the amino-acid-sensing site in the flytrap cleft and stabilize the closed, active conformation, so that the receptor fires at a lower threshold than it otherwise would. The peptide does not shout; it leans on the door. Downstream, an activated CaSR couples through G_q/11 to phospholipase-Cβ, which cleaves PIP_2 into IP_3 and diacylglycerol, and the IP_3 pulls calcium out of the cell's internal stores — a rising intracellular calcium transient that is, conveniently, the exact thing a laboratory can watch. Load a cell line expressing

human CaSR with a calcium-sensitive fluorophore, add peptide, and the brightness that follows is the receptor's answer, read directly in light. That brightness, plotted against concentration, is the raw material of the index.

What kokumi does to everything around it

Because kokumi adds no taste of its own, its whole effect is on its neighbors. A kokumi peptide, by sensitizing CaSR on the taste cell, *modulates* the signals arriving through the other receptors — amplifying and prolonging the perception of umami, sweetness, and saltiness already present. It is a gain knob on the rest of the palate, loudness lent to flavors it does not itself possess.

Measured in time, a kokumi-rich mouthful carries a longer, flatter, more even envelope — slower to fall away, the finish drawn out rather than dropped. Measured in space, the flavor seems to fill the whole mouth rather than sit in one bright spot on the tongue. This is the physical basis of the words cooks have always reached for — *body, roundness, continuity, mouthfeel, a sauce that tastes finished* — and it is why kokumi is the connective tissue of a flavor rather than one of its threads. To raise an organism's kokumi is to tune continuity itself: to make a flavor that arrives, stays, and resolves smoothly instead of jumping and dropping, which is much of what separates cooking that satisfies from cooking that merely registers.

How a plant builds a peptide that survives the pan

A plant does not stumble into these peptides; it assembles them, enzyme by enzyme, from amino acids it already holds in store, and that assembly line is exactly where a design reaches in. The backbone molecule is glutathione, built in two ATP-spending steps. First, glutamate-cysteine ligase (GCL) joins glutamate to cysteine — and critically, it is this enzyme that forges the γ-linkage, spending a molecule of ATP to weld cysteine onto the side-chain carboxyl of glutamate rather than its backbone, producing the dipeptide γ-glutamyl-cysteine. Second, glutathione synthetase (GSS) spends another ATP to add glycine to the far end, closing the tripeptide γ-Glu-Cys-Gly. The cell now holds glutathione — the first, mild rung of the ladder.

The louder rungs are made by a third enzyme working on that stock. γ-glutamyl transferase (GGT, also called γ-glutamyl transpeptidase) is a transpeptidase: it grips a γ-glutamyl donor — glutathione will do — snaps off the γ-glutamyl group, and transfers it onto an acceptor amino acid waiting nearby. Hand GGT a pool of valine and glycine and it will build γ-Glu-Val-Gly, the loud tripeptide, a rung roughly ten times louder than the glutathione it came from. The whole kokumi character of an organism, then, is set by three dials: how much glutathione it makes (GCL and GSS, and the cysteine and glutamate feeding them); how active its GGT is; and which acceptor amino acids sit in the cytosol for GGT to reach. A design that lifts GGT activity and stocks the cell with valine shifts the plant's peptide profile up the ladder — more of the loud peptide, less of the mild — and continuity rises with it.



Garlic — beneath its pungency, a reservoir of the γ-glutamyl peptides that lend food its body, its continuity, and the sense of being finished.

This also explains, at last, why the foods every tradition reaches for to build a base are the ones richest in these peptides. Alliums run the γ-glutamyl machinery hard as part of how they store and traffic sulfur; garlic and onion are, beneath their pungency, reservoirs of exactly these peptides. Yeast carries glutathione in abundance. And because the γ-linkage resists ordinary peptidases, the long-kept and the fermented — aged cheese, miso, fish sauce, anything given time and culture — accumulate γ-glutamyl peptides rather than losing them, which is why age and ferment and patience all read, on this one axis, as the same move: the quiet addition of body beneath everything else.

Measuring a taste that leaves no fingerprint

How do you score a molecule you cannot taste? The same way we score every axis — by how *loudly* it speaks to its receptor, not by how much of it a food holds. The instrument is the half-maximal effective concentration, EC_{50}: the concentration at which an agonist drives its receptor to half of full response. The index is a ratio built from it:

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

A low EC_{50} means a loud molecule — a little of it goes far — so when it sits in the denominator the score climbs. We anchor the kokumi axis to glutathione, the peptide the cell makes for itself, so that glutathione scores exactly 1,000 and every other agonist is read against the familiar. In our CaSR assay glutathione half-activates the receptor at 2.5 mM; we fix that value as EC_{50}(reference).

From there the ladder falls out by arithmetic. The dipeptide γ-glutamyl-cysteine is weaker — half-max near 5.0 mM — so it scores 1,000 × 2.5 ÷ 5.0 = **500**, half as loud as glutathione. γ-Glu-Val-Gly half-activates CaSR at about 0.195 mM, and so it scores 1,000 × 2.5 ÷ 0.195 ≈ **12,800** — an order of magnitude past the reference, and well past the 5,000 ceiling at which the radar scorecards clamp, where it simply pins the kokumi spoke to its outer edge.

Γ-GLUTAMYL PEPTIDE	STRUCTURE	CASR EC_{50}	IFI-K	VS. GLUTATHIONE
γ-glutamyl-cysteine	γ-Glu-Cys (dipeptide)	5.0 mM	500	0.5×
Glutathione — *reference*	γ-Glu-Cys-Gly (tripeptide)	2.5 mM	1,000	1.0×
γ-Glu-Val-Gly	γ-Glu-Val-Gly (tripeptide)	0.195 mM	12,800	12.8×

The millimolar figures hide how little of the loud peptide it takes, so it is worth converting out of molarity. Glutathione's molar mass is 307.3 g/mol, so its 2.5 mM half-max is 2.5 mmol/L × 307.3 mg/mmol ≈ 768 mg/L — about 0.077% by weight, roughly three-quarters of a gram in a liter of broth. γ-Glu-Val-Gly's molar mass is 303.3 g/mol, so its 0.195 mM half-max is only 0.195 × 303.3 ≈ 59 mg/L — about 0.006% by weight. Thirteen-fold less peptide, by mass, for the same effect on the receptor: that mass ratio is the whole reason the loud peptide is worth designing toward, and the whole reason a plant that makes a little of it can outperform a plant stuffed with the mild kind.

EC_{50} is a single point on a curve, and the shape of that curve matters as much as the point. CaSR activation follows the sigmoid Hill form,

$$R([A]) = R_{max} \times [A]^n \div (EC_{50}^n + [A]^n)$$

where n is the Hill coefficient, the steepness of the rise; CaSR, a cooperative homodimer, runs steep, around n ≈ 1.8. Setting [A] = EC_{50} gives R = R_max/2 by construction — that is the definition — but the rest of the curve is where the design lives. Walking γ-Glu-Val-Gly up from half its EC_{50} to four times it:

[Γ-GLU-VAL-GLY]	MULTIPLE OF EC_{50}	CASR ACTIVATION
0.098 mM	0.5×	22%
0.195 mM	1.0×	50%
0.390 mM	2.0×	78%
0.780 mM	4.0×	92%

The lesson in that table is the cliff. Below the EC_{50} the response falls away fast; above about four times it, the receptor is all but saturated and more peptide buys almost nothing. A design that overshoots — that floods a cell with far more peptide than EC_{50} — does not taste "more full"; it tastes the same fullness while spending the plant's amino acids for no return, and risks dumping sulfur intermediates that surface as off-notes. The useful working range is a narrow band around the EC_{50}, which is precisely why knowing the number, rather than guessing at "richer," is the whole game.

The formula runs backward just as well, and that is how we read a designed cultivar. Our model predicts Phoenix, our full-spectrum flagship, at a kokumi score of 1,900. Invert the index — EC_{50}(candidate) = 1,000 × EC_{50}(reference) ÷ IFI-K = 1,000 × 2.5 ÷ 1,900 = 1.32 — and the number tells us that a standardized Phoenix extract is designed to reach half-maximal CaSR activation at a concentration equivalent to 1.32 mM of glutathione: our design makes it 1.9 times louder, peptide for peptide, than the reference the cell makes for itself.

A score is only as trustworthy as its error bar, and because the index is a ratio, its uncertainty is the two measurements' uncertainties added in quadrature:

```
CV(IFI) = √[ CV(reference)² + CV(candidate)² ]
```

We hold each EC_{50} to a batch-to-batch coefficient of variation of 7% — the same production discipline that keeps our capsaicin lines to ≤7% CV against the 42% that unimproved stock drifts across — and the kokumi score inherits $\sqrt{(0.07^2 + 0.07^2)} = 0.099$, call it ten percent. Phoenix's K of 1,900 is therefore 1,900 ± ~190: a real band, honestly drawn, narrow enough to design against. Let that production CV loosen to 42% instead and the same mean would carry a band of nearly ±800, sweeping from 1,100 to 2,700 — three different cultivars wearing one number, and nothing you could build on. The tight CV is not housekeeping; it is what makes the coordinate mean anything at all.

The organisms that carry it

The plants that speak kokumi loudest are, to no one's surprise, the base-builders — the aromatics, the ferments, the patiently aged — and our host library reflects it. Five of the forty-two validated hosts carry kokumi as a primary axis, and the construct counts track how much machinery each one offered us to work with.

ORGANISM	CONSTRUCTS	PRIMARY AXES	KOKUMI ROUTE	DIFFICULTY
Garlic	560	K / P	γ-glutamyl peptides (GGT)	Moderate
Soybean	360	U / K	glutathione, γ-Glu peptides	Frontier
Yeast	240	U / K	glutathione, γ-Glu peptides	Hard
Onion	40	K	γ-glutamyl peptides	Moderate
Vanilla	40	K	kokumi-active phenolics	Hard

Garlic, at 560 constructs, is the standard-bearer: it already runs its γ-glutamyl transferase hard, so much of our work there is a matter of lifting and steering an existing capacity rather than installing a new one — moderate difficulty, outsized payoff. Soybean and yeast each carry kokumi alongside umami, which is why they anchor the savory-depth work: a single organism supplying both the depth and the continuity that make a broth taste cared-for. Onion walks the same γ-glutamyl route as garlic at smaller scale. And vanilla is the quiet surprise on the list — its kokumi runs not through peptides at all but through kokumi-active phenolics, a reminder that the axis is defined by what the receptor reads, not by any one chemical family, and that part of vanilla's legendary depth is a kokumi effect hiding beneath its aroma.



Vanilla — the most complex and most counterfeited flavor we know; part of its depth is a kokumi effect carried not by peptides but by kokumi-active phenolics, hiding beneath the aroma.

It is no accident that these are the foods every cuisine reaches for to build a bottom. Garlic and onion sweated slowly into the base of a pan; a spoonful of miso; a rind of aged cheese dropped into a simmering pot; the long-kept, the fermented, the patiently browned — all of them are, on this axis, the same gesture: the addition of continuity beneath everything else. Cooks have always known these things make a dish taste *finished* without quite being able to say what they contributed, because what they contributed was not a flavor but the smoothing of all the other flavors into one. Kokumi is the name for that smoothing, and a number for it at last — in our own kitchen, the discovery that most changed how we think about a base: not a flavor to be added on top, but a continuity to be built underneath, so that everything laid over it arrives joined rather than scattered.

Where kokumi fails, and what it cannot do

Honesty about an axis means naming where it breaks, and kokumi has more failure modes than any other spoke on the index, precisely because it is invisible.

The first is the one it is named for: kokumi is the axis most often *missed*. Rebuild a beloved food from a parts list — match its sweetness, its salt, its umami, its acid, every nameable component — and it can still come out technically correct and somehow thin, hollow in a way no single missing ingredient explains. What is absent is the continuity, the quiet peptide layer no one thought to measure because it has no taste to miss. More reproductions fail here than anywhere else, which is the practical argument for giving the thing a coordinate: you cannot replace what you have not counted.

The second is saturation. Because the Hill curve flattens past about four times the EC_{50}, a plant engineered to overproduce kokumi peptides does not grow richer without limit; it hits the ceiling of its receptor and then merely spends amino acids, often dumping cysteine and sulfur intermediates that surface as sulfurous or bitter off-notes. More is not louder once the receptor is full. The design target is therefore the EC_{50} band, not the maximum — a discipline that only makes sense once you can see the curve.

The third is that kokumi has nothing to amplify on its own. It is a gain knob, and a gain knob on silence is still silence: a kokumi peptide added to clean water, or to a dish with no underlying umami, sweetness, or salt, produces almost nothing a taster can name. The axis only pays out in company. This is why the K score must be read as a modifier, not a flavor — a cultivar with a high K and nothing beneath it is a design error, not a success.

And the honest outer limit: the index measures the lever, not the whole felt experience. EC_{50} is a receptor-cell number, and human perception is supra-additive and context-dependent in ways a single-cell assay cannot fully carry; individual CaSR variants shift sensitivity from one person to the next; the matrix of a real food — its fat, its temperature, its viscosity — bends how a peptide is delivered to the tongue. We measure the loudness of the molecule at the receptor with real precision, and we are candid that the last step, from receptor to the particular fullness *you* feel, is one the number informs rather than replaces.

Kokumi in Phoenix

Our model places Phoenix's predicted kokumi at 1,900 — the third-highest axis score in its design, behind heat and bitterness-attenuation (each axis calibrated independently against its own reference point) — and the reason is structural, not decorative.

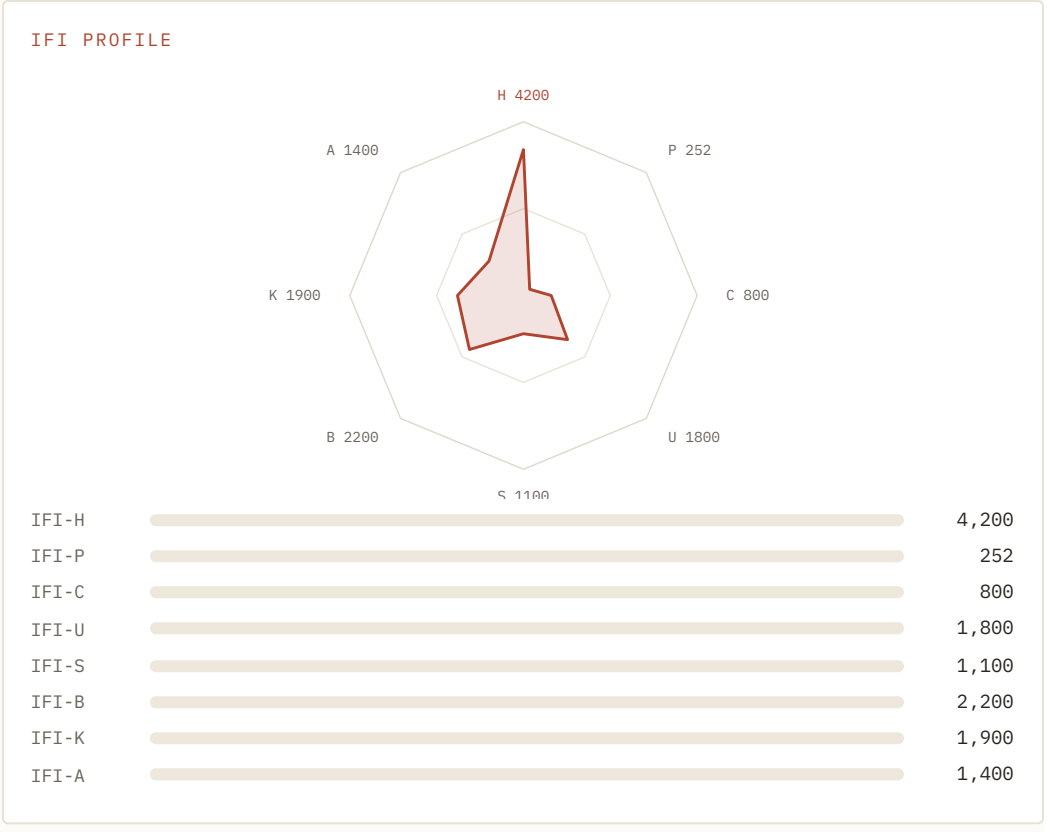



Phoenix (SS-0096) in the controlled chamber — predicted kokumi of 1,900 is the invisible axis that keeps our model's predicted ~2.95-million-Scoville heat from feeling hollow.

Our model predicts roughly 2.95 million SHU for Phoenix; without kokumi beneath it, that heat would be a spike over an empty mouth: a flash of fire with nothing to ride on, gone as fast as it came. With kokumi built in, the heat has a body to travel through — a continuity that carries the burn and the fruit together and lets the whole ferocious thing feel, somehow, *complete* rather than merely violent. It is, fittingly, the invisible axis that keeps the loudest cultivar we have designed from feeling hollow.

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

And it is here that kokumi meets the ethos running under all of our work. The pairing we are asked for most often is umami with kokumi — savory depth joined to savory continuity — because together they are most of what separates a broth that comforts from one that merely tastes of its ingredients: umami gives a flavor its depth, kokumi gives that depth its continuity, and the two arrive by the same quiet route. In the plant-forward kitchen our family has always kept, that pairing does heroic work. It builds the deep, lasting fullness that long-simmered bones once supplied, and it builds it out of alliums and ferments and designed peptides instead — so that nourishment and savor arrive together, with nothing asked of an animal that a plant can be taught to give. Kokumi is the name, measured at last, for the richness we believe a nourishing food should always carry: not a flavor laid on top, but a continuity built underneath, so that everything above it arrives joined rather than scattered.

Chapter 10 · Cool-Burn — *temperature that isn't there.*

Chili heat is a fire with no flame. Cool-burn is its mirror — a cold with no cold — and the symmetry is not a figure of speech but a mechanism, because the cool runs through the same family of channel as the burn, only pointed the other way along the thermometer. Where capsaicin reaches into the body's alarm for *too hot* and trips it early, menthol reaches into the body's sensor for *pleasantly cool* and trips that one early instead. This is the axis of the clean breath and the rinsed-with-winter finish, the mint that is not cold and the lift a cold drink or a confection leaves behind — a sensation the tongue swears to and the thermometer flatly denies. We find it one of the most honest spokes in the whole Index, because its name admits out loud what every one of them is secretly doing: borrowing the body's machinery to counterfeit a feeling the world is not supplying. On this axis the counterfeit is a chill, and the forgery is so good that a mouth at blood heat will insist it has just drunk from a cold spring.

The channel that reports a cold that isn't there

TRPM8 is a non-selective cation channel, built like its TRP cousins as four identical subunits arranged around a single central pore, standing on the same trigeminal nerve endings in the mouth and skin that carry heat and pungency. In its natural work it is a thermostat. Below about 26 °C the pore opens, calcium and sodium flow inward, the ending depolarizes, and a train of impulses climbs to the brain carrying the welcome news that something is cool — a breeze, a draught of water, the first bite of winter air. It is a *polymodal* sensor: the same gate answers to temperature and to membrane voltage together, so that cold and electrical drive each lower the price of opening the other. The channel does not measure degrees. It reports a probability of being open, and the brain reads that probability as cold.

Menthol, the signature terpene alcohol of mint, is a key cut to fit this lock. It binds within the transmembrane bundle of the channel — in the pocket formed by the voltage-sensing helices, not at the pore mouth — and from there it shifts the whole activation relationship toward warmer values. The half-activation voltage slides; the apparent threshold temperature climbs from the low twenties up into the thirties. The practical consequence is exact and a little astonishing: a channel that should have needed a genuine chill now opens at ordinary mouth warmth, somewhere around 34–36 °C, because menthol has moved the set-point to meet the body where it already sits. No heat is removed. No tissue cools. The receptor simply crosses its own threshold at a temperature it would normally ignore, and the brain, faithfully reading the calcium traffic, reports a cold that the mouth does not contain.

Two features of the gating matter for everything that follows. First, the open state is *cooperative*: four subunits sit around the pore, and the channel flickers fully open only when they agree, so the transition from silent to firing is steep rather than gradual — cool does not fade up, it snaps on over a narrow band of concentration. Second, the effect is *graded in potency, not in kind*. A weak agonist and a strong one both open the same channel to the same maximum; they differ only in how little of them it takes to get there. That single fact — potency, not presence — is the hinge on which the entire axis turns, and it is what lets a mint and a menthol-free coolant and a barely-cool citrus all be scored against one another on one line.



FIG. 11 — Peppermint, whose six-enzyme relay turns a simple citrus terpene into (−)-menthol, the molecule that moves TRPM8's thermostat warmer.

Borrowing the body's thermostat

The parallel to capsaicin and TRPV1 is worth drawing with care, because it is the whole secret of the axis. Capsaicin lowers the *heat* threshold of a hot channel, so that body temperature registers as fire. Menthol raises the *cold* threshold of a cold channel, so that body temperature registers as frost. One molecule pulls a set-point down, the other pushes a set-point up, and in both cases the plant borrows the body's own 37 °C to manufacture a sensation out of nothing but a shifted number. The thermometer is right the whole time. It is the threshold that has moved.

Because the receptor only reports probability, it does not care where the key comes from — and industry has learned to open it without mint's aroma at all, which matters whenever you want the cool but not the flavor of it. **Icilin** is a synthetic agonist far stronger than menthol and oddly dependent on intracellular calcium to do its work; the **coolants WS-3 and WS-23** are odorless cyclic carboxamides that open TRPM8 cleanly and smell of almost nothing. The channel reports cold for all of them, which is the deep reason cool can be designed *apart* from the smell of mint — a cool unbundled from its traditional aroma and dialed in on its own. That separability is the engineering gift of the axis, and it is why the Index scores the sensation rather than the scent.

And because **TRPV1** sits on the very same trigeminal endings, a true cooling design has to manage the *burn* as well as the cool. A little warmth from TRPV1 running underneath a strong TRPM8 cool reads as *brisk* — invigorating, alive, the difference between a flat mint and

a bright one. Too much, and the two signals muddy into a mentholated heat that satisfies neither sense. Schematically, the net percept behaves like a difference of the two channels' drives, *cool ≈ (TRPM8 activation) − k·(TRPV1 activation)*, with a small weight *k* that the body sets and the formulator cannot. This is precisely why the axis is named cool-*burn* and scored as a balance rather than a single pole: it is the one spoke whose name confesses that two channels are always in the room together, and that the number is really a truce between them.

The arithmetic of cool

The Index refuses to measure a flavor by how much of a compound a food holds, because that number lies — a loud molecule at low concentration beats a quiet one at high. It measures *loudness* instead, anchored to the receptor, with one formula that every axis obeys:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the single most honest point on a dose–response curve, and we will say in a moment exactly why we chose it. By construction the cool-burn reference is pinned to IFI-C = 1,000: a fixed (−)-menthol standard against which every other cool is read. Everything above a thousand is a cool louder than that standard; everything below is quieter. The axis is deliberately *dimensionless and relational* — it certifies a ratio of potencies, not an absolute. That is a feature, not a dodge: ratios survive a change of assay, of day, of hands, where raw EC_{50} values drift.

Invert the formula and the design meaning falls out:

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

The index and the half-activating dose are reciprocals. Double a candidate's IFI-C and you have halved the concentration it needs to reach the same half-activation — the same cool for half the delivered molecule. Loudness, measured this way, *is* economy.

Worked example — Phoenix, the only sub-thousand spoke we will ever defend. Phoenix (SS-0096), the one cultivar our family built to register on all eight axes at once, carries IFI-C = 800. Put it through the inversion:

> $\mathrm{EC_{50}}(\text{Phoenix}) = 1{,}000 \times \mathrm{EC_{50}}(\text{ref}) \div 800 = \mathbf{1.25 \times EC_{50}(ref)}$

Its cool principle needs 25% *more* concentration than the reference standard to reach the same half-activation. It is quieter than the baseline, on purpose, and the number says so to two digits. We will return to why a superhot wants its cool pulled below a thousand; for now, note that the formula lets a *deliberate* softness be stated as precisely as a loudness.

To make the ratios concrete, anchor the reference at EC_{50}(ref) ≈ 100 µM for illustration — the absolute figure is a convention the standard fixes, while the index certifies the ratio — and the whole axis becomes a conversion table between a score, a half-dose, and a working concentration. Menthol's molar mass is 156.27 g/mol ($C_{10}H_{20}O$), so 1 µM of it is 0.156 mg/L, and in dilute aqueous that is essentially 0.156 ppm:

IFI-C	EC_{50} (× REFERENCE)	EC_{50} AT REF ≈ 100 MM	≈ PPM	READING ON THE AXIS
500	2.00×	200 µM	31.3	half as loud — a whisper of cool
800 — Phoenix	**1.25×**	**125 µM**	**19.5**	below baseline; a warming lean
1,000 — reference	1.00×	100 µM	15.6	the anchor
2,000	0.50×	50 µM	7.8	twice as loud
4,000	0.25×	25 µM	3.9	four times
5,000 — scorecard ceiling	0.20×	20 µM	3.1	the normalization maximum

Every row is just EC_{50} multiple = 1,000 ÷ IFI-C, carried out to its practical units. Read it as a design brief in reverse: *to earn IFI-C 2,000, the cool principle must present a half-activating concentration of 50 µM — menthol delivered at half the reference's half-dose, 7.8 ppm instead of 15.6.* The ceiling at 5,000 is the scorecard's normalization maximum; a cool five times the reference standard, at one-fifth its dose, is the loudest the radar will plot.

Why EC_{50} and not EC_{10} or EC_{90}. The dose–response of a cooperative channel follows a Hill curve, $\theta(C) = C^n \div (C^n + EC_{50}^n)$, where θ is the fraction of maximal response and *n* is the Hill coefficient. At C = EC_{50}, θ = 0.5 exactly — and the curve is steepest there, so a small error in the measured concentration produces the smallest error in the inferred potency. The midpoint is the one place on the sigmoid least contaminated by where you drew the baseline or where you called saturation. We anchor to it because it is the most reproducible point on the most reproducible curve, and reproducibility is the only thing that makes a flavor number worth printing.

The four-subunit architecture makes *n* greater than one, which sharpens everything. The concentration span from barely-on to nearly-full — the 10%-to-90% window — is 81^(1/n). For a representative cooperative slope of n ≈ 2 that span is just 9-fold; for n ≈ 3 it collapses to about 4.3-fold. In plain terms, cool does not ease in across decades of concentration; it snaps on across a single order of magnitude or less. The usable formulation window is narrow — a touch too little reads as nothing, a touch more reads as full winter — which is exactly the knife-edge the index is built to help you stand on. And a change in *loudness* is a pure sideways slide of that whole steep curve: halving EC_{50} (doubling IFI-C) shifts it left on the log-dose axis by $\log_{10}2 \approx 0.30$ decades, the same percept reached at a lower dose, with the shape untouched.

Error, and why the editing earns the number. Because IFI-C ∝ $1/EC_{50}$, the relative error in the index equals the relative error in the estimated potency to first order — a 7% wobble in the half-dose is a

7% wobble in the score, no amplification, no cancellation. Batch-to-batch concentration variation enters the same way, and independent sources combine in quadrature, $CV_{total} = \sqrt{(CV_{batch}^2 + CV_{assay}^2)}$. Here the platform's reproducibility stops being a slogan and starts being arithmetic. Our engineered cool principles hold a batch CV at or below **7%**; a wild-type botanical extract runs near **42%**. Carry each through Phoenix's C = 800 at one standard deviation:

SOURCE OF THE COOL	BATCH CV	C = 800 LANDS WITHIN (±1Σ)	WHAT THE INDEX CAN RESOLVE
Engineered cool principle	≤ 7%	744 – 856	"below baseline, a warming lean" — cleanly resolved
Wild-type botanical extract	≈ 42%	464 – 1,136	below, at, *or* well above the reference — unresolved

Seven percent of 800 is 56, so the engineered number sits in a band narrow enough that *below-baseline* means something you can defend. Forty-two percent of 800 is 336, a smear so wide it swallows the reference line entirely — the same nominal cool could be a warming lean or a genuinely bright mint, and the computation could not tell you which. A flavor you cannot make twice cannot be scored once. The reproducibility the editing buys is not a quality-control footnote; it is the precondition for the axis existing at all.

Six enzymes from citrus to cold

The chemistry behind the cold is one of the most elegant short pathways in the plant world, and **peppermint — 2,360 constructs, one of the library's most worked organisms —** runs it in the oil glands studding the underside of its leaves, a tidy enzymatic relay that walks a simple citrus terpene step by careful step down to a single, precise stereochemistry:

STEP	ENZYME	INTERMEDIATE PRODUCED
1	limonene-3-hydroxylase	trans-isopiperitenol
2	isopiperitenol dehydrogenase	isopiperitenone
3	isopiperitenone reductase	cis-isopulegone
4	isopulegone isomerase	pulegone
5	pulegone reductase	menthone
6	menthone reductase	**(−)-menthol**

It begins with **limonene**, the plain citrus terpene — the smell of a zested orange — and subjects it to a hydroxylation, a dehydrogenation, and a chain of reductions, threading it through isopiperitenol, pulegone, and menthone to arrive at **(−)-menthol**, the single stereoisomer the receptor reads most coldly. Every step is a place where the design can tune the final balance of menthol against its warmer, harsher cousins, and the stakes of that tuning are real and tastable. **Pulegone,** left unreduced at step 4, is off, acrid, and unwelcome; much of the art of a good mint lives in *finishing* the relay cleanly — driving the whole flux down to menthol rather than letting it pool at an intermediate. A great cooling flavor is as much about what the pathway does not leave behind as about what it makes.

Stereochemistry is the hidden lever, and here the mathematics is combinatorial. Menthol carries **three stereocenters**, so there are 2^3 = **eight** possible stereoisomers — four diastereomeric pairs, known

as menthol, neomenthol, isomenthol, and neoisomenthol — and the tongue is an exacting stereochemist that only answers loudly to one of them, (−)-menthol, the (1R,2S,5R) form. The others read as weaker, mustier, or faintly medicinal. The design problem is therefore to drive flux not merely to *a* menthol but to **one of eight**, and the quality metric is enantiomeric excess, %ee = (R − S) ÷ (R + S) × 100 — how cleanly the relay finished at the correct hand.

The dilution math is unforgiving and first-order. Because the off-isomers are far weaker TRPM8 agonists, a batch that finishes at 90% (−)-menthol, with the remaining tenth split among the wrong isomers and residual pulegone, delivers roughly 90% of the theoretical cool — an IFI-C target of 2,000 slumps to about 1,800 on stereochemical scatter alone, *before* you account for the acrid pulegone, which does not simply disappear but reads as an off-note on other axes, a composition penalty layered on top of the loss. The reductions at steps 5 and 6 are where a program either earns a clean, bright, confident cool or leaves a muddied one on the plate. In our family the model that scored the enzyme variants ran on one pair of hands; the arguing over whether a batch had finished clean — whether the cold arrived whole or came in wearing a trace of something medicinal — belonged to all of us, at the table, spoon by spoon.

The cooling family, read as balance

Mint is not the only plant that counterfeits cold, and the Index reads the whole family's differences as different *balances* on the single axis — not a yes-or-no, but a graded score of how cleanly cool a thing is against whatever warmth and sharpness ride alongside it:

ORGANISM	CONSTRUCTS	COOLING PRINCIPLE	REGISTER ON THE AXIS	DIFFICULTY
Peppermint	2,360	(−)-menthol	bright, clean, high-definition edge	Hard
Lavender	160	linalool / camphor notes	cool wrapped in floral softness	Hard
Eucalyptus	80	eucalyptol (1,8-cineole)	softer, diffuse, faintly medicinal	Hard
Spearmint	80	carvone (cool-adjacent)	aromatic, cool-adjacent, not truly cold	Hard
Camphor	40	camphor	cool shadowed by its own warmth	Hard

Eucalyptol — the 1,8-cineole of eucalyptus — opens TRPM8 to a softer, more diffuse, faintly medicinal cool than menthol's bright edge. **Camphor** brings a cool shadowed by its own warmth and sharpness, a near-contradiction the axis is built to score honestly rather than flatten. **Lavender's** cooling note arrives wrapped in the floral softness of **linalool**, cool and perfume braided together. **Spearmint** carries **carvone** rather than menthol and so sits *adjacent* to true cool — aromatic, herbaceous, more the idea of cool than the fact of it. Each lands at a different point of balance between clean TRPM8 drive and the warmth, sharpness, or scent that travels with it, which is exactly why the axis must be a graded number: it measures not *whether* a thing is cool but *how* cool, and against how much else. Even a citrus belongs on the ladder — orange carries a quiet cooling spoke alongside its bitterness (its axes read B/C in the library), the faint bright lift limonene gives a cold glass of juice.



FIG. 12 — Eucalyptus. Its eucalyptol opens the same channel as menthol to a softer, more diffuse cold — the cooling family is a spread of balances, not a single note.

This is where stereochemistry and family breadth meet in one design decision. A mint is graded not merely by *how much* menthol it holds but by how faithfully its relay finished at the correct-handed molecule and did not leave a scatter of the wrong ones — or a residue of pulegone — behind. The cooling principle you reach for sets the *character* of the cold; the cleanliness of the pathway sets whether that character arrives pure or muddied. Loudness is the EC_{50} ratio; purity is the enantiomeric excess; the two together are what a cool-burn score actually certifies.

Cool beyond the tongue

Cool-burn is the axis most useful far from the dinner plate, and that is worth saying plainly because it shows the Index measures a sensation at large and not only a taste. The same thermal illusion that makes a mint read as cool in the mouth reads as cool on the skin, because TRPM8 rides sensory nerves across nearly the whole body, not merely in the taste buds — the bright clean note of a rinse, the lift of a cooling balm, the finish that makes a confection or a cold drink read as *clean*. We note this as a fact of the receptor, not a claim about it: cool is a sensation the body welcomes almost everywhere, which is why the axis matters to oral care and confectionery and beverage in equal measure, and why the small word *fresh* does so much quiet work across all of them. Freshness, it turns out, is a channel, a stereoisomer, and a pathway finished with care.

There is a gentler reason this axis sits close to our family's heart. Cool is one of the few sensations the body receives as pure welcome — nothing to brace against, nothing to recover from, a cold with no cold and no cost. It is the clean breath after a meal, the rinse that resets the mouth, the sensation of having given the body something bright and uncomplicated. We have always cooked and grown toward food you give your body with intention, and a well-made cool is a small, honest version of that intention: a pleasure with nothing acrid left behind, a finish that is literally the absence of residue. The axis measures how much of that welcome a plant can be made to deliver, and how cleanly.

The axis tuned downward

On the Index, cool-burn carries a subtlety worth naming: it is the axis most often tuned *downward*, and Phoenix is the proof. Here is its full eight-axis profile, the one cultivar in the library we engineered to register on every spoke at once:

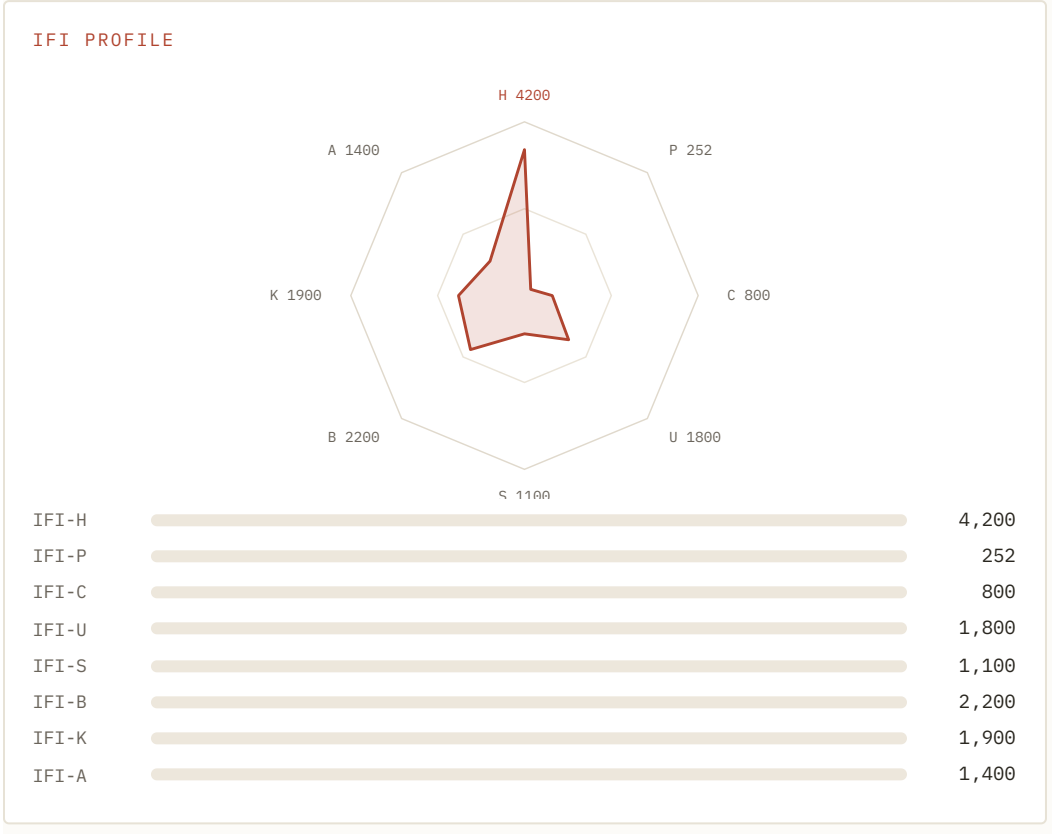

Read the spokes together and the composition announces itself. Heat towers at 4,200 — our model predicts approximately 2.95 million Scoville heat units, a loudness we built gene by gene. Umami, zero-bitter, and kokumi stand in the high hundreds and low thousands, giving the fire a body to sit in. And cool-burn alone drops *below* the reference line, to **800** — the only sub-thousand spoke in the entire profile. That is not a shortfall; it is a decision. Run it back through the formula, EC_{50} = 1.25 × reference, and the number resolves to a deliberate, faint warming lean, because a superhot wants nothing in it that cuts against the heat. A bright TRPM8 cool fighting a 4,200 burn would read as the mentholated muddle the axis is named to warn against. So we pulled one spoke quietly beneath the baseline, on purpose, so the other seven could stand.



FIG. 13 — Phoenix (SS-0096) in the controlled chamber, where every spoke is evaluated against the same reference — including the one we chose to hold below a thousand.

That single choice is the clearest proof in the whole library that the Index is an instrument of *composition* and not merely of measurement. You do not only decide how loud. You decide, axis by axis, which way the flavor leans — and sometimes the most

expressive number on a profile is the one held quietly below a thousand, where the formula can state a softness as precisely as it states a roar.

Honesty requires naming what the number does *not* do, because an axis you trust is an axis whose limits you have mapped. The IFI-C score is a half-maximal potency, an instantaneous reading of how loudly a molecule opens TRPM8 — it is not a verdict on pleasure. A very high cool can still read as harsh, dental, or over-mentholated; icilin's over-bright cold is the cautionary case, loud at the receptor and unwelcome at the table, which is why loudness and hedonics must never be confused on this axis more than any other. Nor does the score carry a clock: TRPM8 can **desensitize** under strong or repeated drive, the cool blunting on the second sip in a way a single midpoint cannot show, so the number describes a first encounter, not a time-course. The absolute EC_{50} is a convention of the shared reference, which means cross-lab comparisons of raw half-doses require that same standard in hand; trust the ratios, not the micromolar. And the cool never travels alone — TRPV1 sits on the same nerve, so a clean receptor reading can still fail at the plate if the warmth it rides with is wrong, the very reason Phoenix holds its cool beneath the line. These are not failures of the Index. They are the edges of an honest instrument, and knowing them is how a number becomes a tool you can design with rather than merely a figure you quote. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. Cool-burn, at its best, is the axis that taught us the Index was never only a ruler. It was always a composer's score — and on this spoke, the quietest note is the most deliberate one of all.

Chapter 11 · Pungency — *the sharp, fast channel.*

If heat is a slow fire and cool-burn its mirror, pungency is the knife — the fast, bright, climbing-the-nose sharpness of fresh wasabi, of mustard that makes the eyes water, of raw garlic and grated horseradish and the first green bite of ginger. It arrives faster than chili heat, peaks higher in the sinuses than on the tongue, and clears sooner, and that whole distinctive signature comes from the receptor it runs through — a channel that does not merely bind its molecules but is *chemically attacked* by them. We grew up at a vegetarian table where this axis did most of the talking: the sting of a mustard seed cracked between the teeth, the clean heat of ginger in a pot of dal, the way a crushed clove could change the mood of a whole dish. For years we could feel that sharpness exactly and say nothing numerical about it. The reason it resisted measurement turns out to be the most beautiful thing about it: pungency is not a molecule settling into a pocket and settling out again. It is a small, deliberate chemical reaction, run to completion on the surface of a living nerve. To put a number on pungency, we first had to learn how to score a reaction.

The receptor that gets attacked

Pungency is read by **TRPA1,** the "wasabi receptor" — a non-selective cation channel riding the same trigeminal and sensory nerve endings as the heat channel TRPV1 and the cold channel TRPM8, and built, like them, as a homotetramer: four identical subunits arranged around a central pore, each contributing six membrane-spanning helices with the ion gate formed between the fifth and sixth. What sets TRPA1 apart before anything culinary is its enormous intracellular skirt of **ankyrin repeats** — fourteen of them stacked into a coiled spring of protein that reaches down into the cytoplasm — a structure found on no other taste channel at this scale, and one that helps set the receptor's exquisite, trigger-hair sensitivity.

But the defining feature of TRPA1 is not its shape. It is its chemistry of activation, which is unlike any other axis in this book. Its agonists are **electrophiles** — reactive molecules that hunt for electron-rich sites and bond to them covalently. On TRPA1's cytoplasmic side, along the linker that leads into the first membrane helix, sits a small cluster of exposed **cysteine residues** — in the human channel, the three best-studied are C621, C641, and C665, with a nearby reactive lysine — and the sulfur atom of each cysteine, in its deprotonated thiolate form, is precisely the kind of electron-rich site an electrophile seeks. The pungent molecule reaches in and forms an *actual covalent bond* to one of them. The channel is not occupied; it is modified.

The reaction takes a different form for each family of pungent principle, and naming them is worth the trouble, because the form of the bond sets the character of the burn:

- **Isothiocyanates** — allyl isothiocyanate from mustard and wasabi, phenethyl isothiocyanate from watercress — carry a central carbon strung between nitrogen and sulfur (R–N=C=S) that is starved for electrons. The cysteine thiolate attacks that carbon and forms a **dithiocarbamate** adduct. This is the sharpest, most durable bond of the three.
- **Michael acceptors** — cinnamaldehyde from cinnamon bark, the gingerols of ginger — present an α,β-unsaturated carbonyl, a double bond conjugated to a carbonyl that pulls electron density away from the far carbon. The thiolate adds across it (a **1,4-Michael addition**) to form a thioether. This bond can run backward, so these pungencies are a touch more fleeting.
- **Thiosulfinates** — allicin from crushed garlic (R–S(=O)–S–R) — react by transferring an allyl-sulfur group onto the cysteine, an **S-thioallylation** that leaves a mixed disulfide behind.

Three different keys; one lock; and every key, having turned, stays bonded rather than drifting out. The covalent event lowers the energy of the channel's open state, the gate swings wide, and calcium and sodium pour through, depolarizing the nerve. The inrushing calcium then binds a site on that cytoplasmic skirt and *potentiates the channel further* — a fast, self-reinforcing kick that makes pungency's onset feel less like a dimmer coming up and more like a switch being thrown. This is why pungency feels like a *reaction* rather than a flavor. It is one: a small, controlled chemistry happening on the surface of your nerve endings, which is a thing worth a little reverence before you decide to enjoy it on purpose.

The covalent signal, and why it lingers

TRPA1 is, before it is anything culinary, the body's general alarm for chemical insult — the sensor that fires at the acrolein in smoke, the volatile tears of a cut onion, the sting of chlorine, the fumes of a solvent. Eating a pungent plant is eating the plant's version of an irritant on purpose and delighting in it, the same strange bargain we strike with chili heat: a protective alarm turned into a pleasure by being dosed small and meant kindly.

Its covalent mechanism gives pungency a distinctive shape in time. Where capsaicin and menthol dock reversibly in a pocket and drift back out as their concentration falls, an electrophile that has formed a real bond to one of TRPA1's reactive cysteines stays put until the cell can repair or replace the channel. *The bond is the signal,* and unmaking it takes the cell time — which is part of why a fierce pungency can leave the nose ringing well after the molecule itself is gone.

The other half of pungency's signature is volatility. These electrophiles are small and airborne, so they do not wait for the tongue; they rise into the nasal passages, where TRPA1-bearing trigeminal nerves meet them directly. That is the mustard "hit" that bypasses the mouth entirely — the retronasal slap of wasabi that lands between the eyes. For the designer, this splits a single axis into two delivery routes at once: the molecule you put on the tongue and the molecule that climbs into the sinus are the same molecule reaching two populations of the same receptor by two paths, on two slightly different clocks. A pungency that scores well in a mouthful

but has already evaporated from the bowl is not well designed. So pungency is as much about controlling the *release* and the *residence* of the electrophile as its amount: a thing stored stable, detonated on cue, and timed as carefully as it is dosed.

Two compartments, one detonation

The most elegant thing about pungency is that the plants store it as a *weapon with the pin pulled only on injury* — a two-compartment system that is harmless until the tissue is damaged.

Mustard and wasabi hold a stable, tasteless glucosinolate — **sinigrin** — in one cellular compartment, and the enzyme **myrosinase** sealed away in another. The two never meet in the intact plant. Bite, crush, or grate the tissue and the compartments rupture together: myrosinase, a thioglucosidase, cleaves the sugar off sinigrin to leave an unstable aglycone, which immediately undergoes a **Lossen rearrangement** — the sidechain migrating from carbon to nitrogen as sulfate leaves — and the sharp **allyl isothiocyanate** is born in the instant of damage, sulfate and glucose falling away as the only byproducts. The pungency you taste never existed in the standing plant. It is manufactured on contact.

Garlic runs the identical trick with different parts. Odorless **alliin** (S-allyl-L-cysteine sulfoxide) sits apart from the enzyme **alliinase**, a carbon–sulfur lyase, until a knife or a tooth breaks the clove; then alliinase cleaves alliin to a reactive sulfenic acid, and two of those condense within seconds into **allicin.** This is why a whole clove is nearly mild and a crushed one is fierce, and why the sharpness mellows as allicin decays into the sweeter sulfur compounds of cooked garlic.



FIG. 12 — Black mustard. Its seed stores a glucosinolate and the enzyme to detonate it; injury releases the sharp isothiocyanate that opens TRPA1.

There is a quiet grace in the cofactors these enzymes need, and it is the kind of thing our table has always paid attention to: myrosinase runs fastest in the presence of **ascorbate** — vitamin C — and alliinase cannot work at all without **pyridoxal-5′-phosphate,** the active form of vitamin B6. The plant builds its detonator around the same vitamins the body wants from the plant. To design this axis is to design that whole two-part system — the stored precursor, the

waiting enzyme, the cofactor that primes it — and to decide how sharply, how fast, and how much of the active molecule is released when the tissue breaks.

PLANT SYSTEM	STORED PRECURSOR	SEQUESTERED ENZYME	COFACTOR	PRODUCT ON INJURY
Mustard · wasabi · horseradish	sinigrin (glucosinolate)	myrosinase (thioglucosidase)	ascorbate (vit. C)	allyl isothiocyanate
White mustard	sinalbin (glucosinolate)	myrosinase	ascorbate (vit. C)	*p*-hydroxybenzyl isothiocyanate
Garlic	alliin (cysteine sulfoxide)	alliinase (C–S lyase)	pyridoxal-5′-P (vit. B6)	allicin (thiosulfinate)

There is a branch point in the mustard reaction that matters enormously to a designer and ruins a careless one: the unstable aglycone does not *have* to become an isothiocyanate. At low pH, or in the presence of a plant protein called the **epithiospecifier protein (ESP),** the same intermediate collapses instead into a **nitrile** — flat, barely pungent, almost mute on TRPA1. A mustard engineered for sinigrin titer but grown in tissue that favors the nitrile branch will test rich in precursor and taste of almost nothing. Designing pungency therefore means designing the *outcome* of the hydrolysis, not merely its fuel: suppress the nitrile branch, hold the pH, and steer every cleaved glucosinolate toward the sharp isothiocyanate we actually want to eat.

The enzyme's own arithmetic

Because pungency is made by an enzyme in real time, its intensity is not fixed by how much precursor the plant holds but by how fast the enzyme can convert it in the seconds after a crush. That rate obeys the oldest equation in enzymology, Michaelis–Menten:

```
v = V_max · [S] / (K_m + [S])
```

where **[S]** is the local precursor concentration, **V_max** the enzyme's top speed when fully loaded, and **K_m** the precursor concentration at which it runs at half that speed. The number that makes pungency *detonate* rather than seep is the relationship between the stored concentration and K_m. Mustard seed holds sinigrin packed to extraordinary local concentrations — tens of millimolar in the ruptured cell — while myrosinase's K_m for sinigrin sits near a few tenths of a millimolar. With [S] running a hundredfold above K_m, the denominator is dominated by [S], the ratio collapses toward one, and the enzyme runs flat out at V_max from the first instant: a near-zero-order burst, the whole payload converted almost as fast as the compartments mix. That is the chemistry behind the word "sharp" — not a slow build but a front-loaded release that is mostly over before the tongue has finished reporting it.

This same kinetics draws the axis's hardest line: **the detonator is fragile.** Myrosinase and alliinase are proteins, and heat denatures them. Toast a mustard seed, roast a clove, or simmer either, and you destroy the enzyme before it can fire — the precursor survives, but there is nothing left to convert it, and the pungency never appears. This is why a cooked mustard is sweet and a raw one ferocious from

the same seed, and it sets a firm rule for anyone designing a pungent cultivar for a cooked dish: either protect the enzyme through the heat step, or move the detonation earlier, before the pan. You cannot score what the stove has already disarmed.

Scoring a reaction

Every axis in the Index answers one question: not how much of a compound a food holds, but how *loudly* that compound speaks to its receptor. The loudness is captured by EC_{50} — the concentration at which a stimulus drives its receptor to half of full response — and the Index inverts it against a fixed reference and scales by a thousand:

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

A smaller EC_{50} means a more potent agonist, so putting it in the denominator makes the score rise with potency, as intuition demands. For pungency we anchor the axis to **allyl isothiocyanate** — the loudest, purest electrophile in the plant kingdom — and set its convention constant to **2.1 µM,** which by definition sets IFI-P = 1,000 for the reference. This convention constant is a stated definition — a fixed anchor, like Scoville's original dilution-to-threshold definition — not a measured literature value. At allyl isothiocyanate's molar mass of 99.15 g/mol, that anchor is worth checking in kitchen units: 2.1 µM is 0.21 mg/L, about **two-tenths of a part per million** in water. The reference pungency is a fraction of a part of the electrophile in a million parts of liquid — a useful reminder of how little of a covalent agonist it takes to be heard.

From the anchor the whole axis follows by algebra. Rearranging the formula, the EC_{50} a candidate must reach to earn a given score is simply $EC_{50}(c) = 1000 \cdot EC_{50}(ref) \,/\, IFI\text{-}P = 2{,}100 \,/\, IFI\text{-}P$ micromolar. That single line turns every target score into a bench specification:

TARGET IFI-P	REQUIRED EC_{50} (MM)	≈ MG/L (PPM)	POTENCY VS. REFERENCE
100	21.0	2.08	0.1× (one-tenth)
252	8.33	0.83	0.25× (Phoenix)
500	4.2	0.42	0.5×
1,000	2.1	0.21	1× (reference)
2,000	1.05	0.10	2×
5,000	0.42	0.042	5× (scorecard ceiling)

Read the Phoenix row as a worked example. Our model predicts our flagship full-spectrum cultivar's pungency at **IFI-P 252.** Running the formula, $1000 \cdot 2.1\ \mu M \,/\, EC_{50} = 252$ gives $EC_{50} = 8.33\ \mu M$ — its designed electrophile reaches computed half-activation of TRPA1 at eight and a third micromolar, about one-quarter the potency of the AITC convention reference, and about 0.83 mg/L in solution. The scorecard ceiling of 5,000 would demand a 0.42 µM EC_{50}, a fivefold gain over the AITC reference compound; nothing we grow approaches it, which is exactly why the ceiling sits there — a headroom, not a target.

The *shape* of the dose–response matters as much as its midpoint, and it is why pungency feels sudden. Receptor activation follows a Hill function, response = $[A]^n / (EC_{50}^n + [A]^n)$, where the exponent **n** — the Hill coefficient — measures cooperativity. Because TRPA1 carries several reactive cysteines and gates through a calcium-driven positive feedback, its effective n is steep, near 2. The practical consequence is geometric: the concentration range that carries the channel from ten percent open to ninety percent open spans a factor of 81^(1/n). At n = 1 that is an 81-fold range, a lazy ramp; at n = 2 it collapses to **ninefold** — a narrow window across which pungency goes from barely there to fully, eye-wateringly present. The steepness of the curve *is* the subjective suddenness. The knife is sharp because the math is.

One honest caveat belongs here, because it governs everything downstream. EC_{50} is an equilibrium idea, and a covalent agonist never reaches equilibrium — the bond accumulates with exposure, so the measured midpoint depends on how long the receptor was exposed. Our pungency EC_{50} is therefore an **operational** value: a potency read at a fixed exposure time, reproducible only because the time is held fixed. Change the exposure window and the number moves. We report the window with the number, always, and we never let a pungency score travel without it.

The organism ladder

With the reference fixed, the garden sorts itself into a ladder — by how many constructs we have built on each organism, and by which electrophile each one speaks:

ORGANISM	CONSTRUCTS	PUNGENT PRINCIPLE	TRPA1 ADDUCT	DIFFICULTY
Black Mustard	2,790	allyl isothiocyanate	covalent · dithiocarbamate	Hard
Garlic	560	allicin	covalent · S-thioallyl	Moderate
White Mustard	200	*p*-hydroxybenzyl isothiocyanate	covalent · dithiocarbamate	Hard
Ginger	120	gingerol / shogaol	covalent · Michael	Moderate
Thyme	120	thymol	reversible · non-covalent	Hard
Watercress	120	phenethyl isothiocyanate	covalent · dithiocarbamate	Hard
Cinnamon	40	cinnamaldehyde	covalent · Michael (reversible)	Hard
Oregano	40	carvacrol	reversible · non-covalent	Hard

The adduct column tells a truth the construct counts hide: the axis is *mostly* covalent but not purely so. The isothiocyanates and allicin bond hard and durably — which is why mustard, horseradish, wasabi, watercress, and raw garlic read loud and linger. The Michael acceptors, cinnamaldehyde and the gingerols, bond reversibly, so their pungency is bright but more willing to let go. And two members — thymol in thyme and its isomer carvacrol in oregano, both phenolic monoterpenoids of molar mass 150.2 g/mol — barely bond at all; they modulate TRPA1 by sitting in it reversibly rather than reacting with it. That makes them the gentlest and most fleeting pungencies in the garden, and also the hardest to pin to a clean,

reproducible EC_{50}, because a reversible modulator's apparent potency drifts with assay conditions in ways a covalent agonist's does not.

The isothiocyanate branch alone spreads remarkably wide across a single receptor. The fierce **allyl** form (99.15 g/mol) dominates black mustard, horseradish, and wasabi; the gentler ***p*-hydroxybenzyl** form carries white mustard; the **phenethyl** form gives watercress its green, climbing bite — each a slightly different key machined to the same TRPA1 lock, each with its own ratio of oral burn to the nasal, tear-drawing hit that is pungency's signature. Black mustard leads the whole axis with **2,790 constructs** — among the most worked organisms in the entire library — and it earns the place, because the sinigrin–myrosinase system is the purest, loudest expression of TRPA1 pungency in any plant, and a precise lever on the sharpness that defines condiments across half the world's cuisines.

Garlic, the double agent

Garlic is the axis's great double agent, and it reveals how several of the eight axes can live inside one plant at once. The same clove that detonates allicin onto TRPA1 when it is crushed raw is also a quiet reservoir of **γ-glutamyl peptides** — small modified amino acids that register on kokumi, the axis of richness and mouthfeel read not by a taste pore but by the calcium-sensing receptor, CaSR. So garlic reads loud on pungency and deep on kokumi both, and which voice dominates is decided entirely by the hand that prepares it.



FIG. 13 — Garlic, cross-sectioned. One clove holds two axes: the alliin that detonates into pungency on a crush, and the γ-glutamyl peptides that speak the slow richness of kokumi.

Raw and crushed, alliinase fires and pungency leads. Slow-cooked, two things happen at once: the alliinase is denatured so no new allicin forms, and the allicin already made decays down a well-mapped chain — into diallyl disulfide, diallyl trisulfide, ajoene, and the vinyldithiins — none of them sharp, all of them rounder and sweeter, while the heat-stable kokumi peptides survive and step forward. That whole chemistry is the difference between a raw-garlic sauce that cuts and a head of garlic roasted soft enough to spread. One plant, two axes, set by the cook — a reminder that a flavor is never only what a plant contains but what a kitchen does to it, and that the Index, in the end, scores the dish and not merely the ingredient.

Ginger, and the instability of sharpness

Pungency shifts under heat and time in a way no other axis quite matches, because its molecules are reactive and therefore unstable: a design that holds still on the vine may not hold still in the pan. Ginger is the clearest teacher. Fresh, its sharpness is **gingerol**, bright and almost lemony. Dry it and gingerol loses a molecule of water, converting to **shogaol** — markedly hotter and more penetrating, which is why ground dried ginger bites harder than fresh. Cook it and both relax toward **zingerone**, rounder and frankly sweet. A dried or cooked ginger is not a weaker fresh ginger; it is a *different pungency altogether,* carried by a different molecule hitting the same channel with a different force.



FIG. 14 — Ginger. One rhizome, three pungencies: gingerol fresh, shogaol dried, zingerone cooked — the same TRPA1 agonist family rewritten by water and heat.

For the Index this is not a complication to apologize for but a rule to obey: we score the *prepared form a recipe will actually use,* never the raw material in the abstract. A ginger engineered to peak as gingerol and then eaten dried would be scored on a molecule it no longer holds. So a pungent cultivar carries, in its record, not one IFI-P but the value for each handling it is meant for — fresh, dried, cooked — because the number is only true for the molecule actually present at the moment of eating.

Where the design breaks

An honest axis names its own failure modes, and pungency has more than most, precisely because it is a reaction rather than a shape.

The number is time-stamped. Because the agonist bonds covalently, its EC_{50} integrates exposure rather than tracking instantaneous concentration. The operational value we report is reproducible only at a fixed exposure window; quote it without the window and it is nearly meaningless. We treat the exposure time as part of the measurement, not a footnote to it.

The nose dose is not the tongue dose. These electrophiles are volatile, so a measurement made in solution does not capture what evaporates into the headspace and climbs into the sinus — the very route that gives pungency its character. A full pungency score has to carry both an oral and a retronasal reading, because the molecule

reaches two populations of TRPA1 on two clocks, and a design tuned only for the tongue will read flat in the nose where wasabi is supposed to live.

The branch can go mute. In the mustard family, the nitrile branch — favored by low pH or epithiospecifier protein — can quietly siphon a well-stocked precursor pool into a nearly tasteless product. Rich in sinigrin, poor in bite. Yield of isothiocyanate, not titer of glucosinolate, is the quantity that must be engineered and verified.

The detonator is heat-labile. Myrosinase and alliinase die in the pan, so a pungency built for a cooked dish can vanish entirely at the stove. The fuel survives; the match does not.

Variance compounds. The platform spec is a batch-to-batch coefficient of variation of **≤7%,** and a single-compound axis like heat can hit it by controlling one pathway. Pungency is a product of three noisy terms — precursor titer × enzyme titer × conversion yield — and variance multiplies, so the axis is intrinsically harder to hold to ≤7% than a one-molecule axis. We control three distributions, not one, which is why so many of our pungency constructs carry the "Hard" difficulty mark in the ladder above.

And there is a ceiling of kindness. Pungency that lingers stops being pungency and curdles into irritation — the covalent bond that would not release, the alarm that will not reset. The design is bounded not only by how loud it can be made but by how fast it must clear to still read as pleasure rather than injury. We tune for the clean departure as deliberately as for the arrival.

The top note

Of the three trigeminal axes, pungency is the fastest, and its speed gives it a particular job in a composition: it is the top note, the sensation that arrives first and clears first, opening a flavor before the slow axes take over. This is why a dab of wasabi or a bright mustard reads as a *lift* rather than a weight, and why pungency, for all its sharpness, rarely tires the palate the way sustained heat can — it is gone before it can wear out its welcome. We treat it as punctuation rather than paragraph: a sharp, clean opening measured not only for its height but for how fast it rises and falls.

Pungency composes most beautifully with heat — the knife over the fire, the way a good hot mustard or a chili-garlic crisp lands twice, once fast and sharp and once slow and warm. Our full-spectrum flagship, **Phoenix,** is designed to carry both at once: heat at the top of the Index and pungency as a distinct bright note, a covalent edge drawn across a deep slow burn so the flavor has both a front and a body.

IFI PROFILE

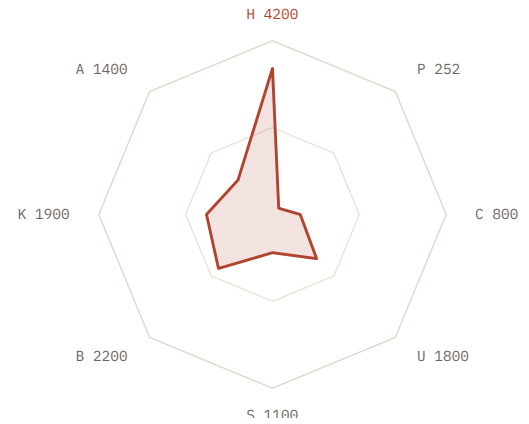

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400



FIG. 15 — Phoenix on the microscope stage. Our model predicts pungency at IFI-P 252 — a bright covalent edge laid over a 4,200 heat, so the flavor keeps moving instead of sitting still.

Read the profile as a composition, not a list. The **4,200** on heat is a slow fire that fills the whole mouth and stays; the **252** on pungency is the knife that opens the bite a half-second ahead of it and is already clearing as the heat arrives. The two together — fast pungency and slow heat — are the whole architecture of why a great spicy food keeps *moving* on the palate instead of sitting still: one sensation departing as the other takes hold. Fast on, fast off, bright and then gone, the top note that makes room for everything under it. One channel attacked and one channel fooled; a knife and a fire, scored side by side at last — and both, in the end, grown to be eaten with intention, and with pleasure, at a table like ours.

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

Chapter 12 · Astringency — *the taste that is a texture.*

We come to the strangest axis, and we kept it on the Index for the most stubborn of reasons: the tongue insists on it. Astringency is the pucker and draw of strong black tea, of young red wine and unripe persimmon and pomegranate pith — the mouth going suddenly dry and tight, drawn inward, stripped of its slipperiness. It belongs to none of the receptor families that carry the other seven axes. It is not, strictly, a taste at all. And it is, unmistakably, part of flavor — so the Index measures it, honestly, as the mechanical event it really is.

That honesty is the whole reason we gave it a number rather than a note. The easy thing would have been to treat astringency as a kind of bitterness, since the same dark polyphenols so often carry both, and to let a receptor score stand in for it. We would not do it. A thing you feel as *friction* should be measured as friction, and the quiet discipline of our table has always been to refuse the convenient fiction — to measure what is actually there, in the units it actually lives in, and only then to build the beautiful account on top. Astringency is the axis where that discipline shows most plainly, because here the Index has to reach past the receptors entirely and anchor itself to a physical collapse.

The film, and its collapse

Begin with the mouth at rest. Your saliva is slick because it is rich in lubricating proteins — above all the **proline-rich proteins, the PRPs**, which come in acidic, basic, and glycosylated families, joined by statherin, the histatins, and the big gel-forming mucins. Together they lay down a thin wet film across every surface of the mouth, a monolayer only nanometers thick that lets the tongue, the teeth, and the palate slide against one another with almost no drag. You never notice it. You notice only its absence.

Into that film come the **tannins** — the large polyphenols of tea, grape, and pomegranate, molecules built to be sticky. A tannin is studded with phenolic hydroxyl groups and flat aromatic faces, and the proline residues of the salivary proteins are shaped, almost perfectly, to receive them: proline's rigid ring stacks against the tannin's aromatic rings while its backbone offers hydrogen-bond partners to the hydroxyls. One tannin, carrying many such faces, binds many points on many protein chains at once. It does not merely stick — it **cross-links**, stitching separate protein molecules into a growing mesh that passes its solubility limit and falls out of solution. The lubricating film, cross-linked and precipitated, collapses in an instant.

What you perceive as astringency is that collapse, read off by the mechanoreceptors and free trigeminal nerve endings of the mouth as a sudden rise in friction — surfaces that should glide now dragging, the tongue reporting roughness where a moment ago

there was none. It is a taste that is a texture: flavor reporting on its own physics, sensed by the very same machinery that feels the body of a sauce or the grit of an underripe pear.

There is a detail here we find genuinely moving, and it is also useful. The proline-rich proteins appear to be a defense the body evolved *precisely for this.* Proline-rich sequences bind tannins so greedily, and so uselessly for any other purpose, that the most plausible reading of their whole existence is as sacrificial scavengers — proteins the body spends freely, throwing them to the polyphenols so the rest of the mouth and gut are spared. The sensation we have spent a chapter learning to design is, at bottom, the feeling of the body deploying a shield. When we tune a tannin, we are tuning how hard it pulls against that shield, and the Index is simply the honest record of the pull.

Anchoring a texture to the same ruler

Every other axis on the Index is anchored to a receptor's **EC_{50}** — the concentration at which a stimulus drives its channel or its G-protein receptor to half of full response — and scored by the master relation

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

so that a louder compound (a smaller EC_{50}) earns a larger number, and a fixed reference sits at exactly 1,000. Astringency has no channel to drive. So we kept the *logic* and changed the *event.* For IFI-A we anchor not to half-activation of a receptor but to **half-maximal precipitation of a standard salivary-protein reference** — the concentration at which a tannin throws down half of all the reference protein it can ever throw down. Call that concentration the **PC_{50}**, the precipitation analogue of an EC_{50}, and the formula carries across untouched:

$$\text{IFI-A}(c) = 1000 \cdot PC_{50}(\text{ref}) \,/\, PC_{50}(c)$$

A tannin that precipitates saliva at a *lower* concentration is a louder tannin and scores higher, exactly as a lower EC_{50} does for heat or sweetness. The same half-maximal reasoning, applied to a binding-and-precipitation event instead of a channel opening, lets astringency take its place on the same ruler as the rest without once pretending to be the same kind of thing.

The measurement behind PC_{50} is a dose–response curve, and it has the familiar sigmoid shape. Titrate a tannin into a fixed pool of reference salivary protein and watch the fraction θ that precipitates:

$$\theta([T]) = [T]^n \,/\, \left(PC_{50}^{\,n} + [T]^n \right)$$

At $[T] = PC_{50}$, the two terms in the denominator are equal and $\theta = 0.5$ — half the protein is down. The Hill exponent *n* measures how cooperative the collapse is, and for astringency it is almost always greater than one: the first tannins to bind nucleate aggregates that make the next binding easier, so large, many-faced tannins show a steep *n* near 3 — a threshold you cross rather than a slope you climb — while small, two-ring tannins show a gentle *n* near 1.2, a draw that comes on softly. Two tannins can share a PC_{50} and still feel different

in the mouth, one arriving as a wall and the other as a mist, and the exponent is where that difference is written. In the assay the collapse is read as light: the reference panel is held at a fixed concentration, the tannin is titrated in, and a nephelometer tracks the haze thrown by the growing mesh — the concentration at half-maximal turbidity is the PC_{50} we report, averaged over replicate titrations so the number carries its own error bar.

Now anchor the reference and do the arithmetic, because the numbers are a form of respect for the reader. We define the IFI-A reference as a standardized grape-seed procyanidin of average chain length four, and we fix its half-maximal precipitation against the standard salivary panel at **PC_{50}(ref) = 100 µM** — the concentration that, by definition, reads 1,000. Then any candidate's score follows directly:

> $IFI\text{-}A(c) = 1000 \cdot (100\ \mu M) / PC_{50}(c) = 100{,}000 / PC_{50}(c)$ [PC_{50} in µM]

Our flagship **Phoenix** carries a predicted astringency of **IFI-A 1,400.** Invert the formula to recover the physics behind the number:

> $PC_{50}(Phoenix) = 100{,}000 / 1{,}400 = 71.4\ \mu M$

So Phoenix's structural tannin precipitates half the reference saliva at 71.4 µM, a little under three-quarters of the reference concentration — it is 1.4× as loud, 1.4× as drawing, as the standard. The short table that follows is simply that formula walked across the range a designer actually works in; the reference row and the Phoenix row are fixed points, and the rest is the curve between and beyond them.

PC_{50} (MM, HALF-MAX PRECIPITATION)	IFI-A = 100,000 / PC_{50}	READS IN THE MOUTH AS
250	400	faint, barely drawing
200	500	soft, round
100 *(reference)*	1,000	the standard draw
71 *(Phoenix, designed)*	1,400	clean, structured frame
50	2,000	firm, briskly puckering
40	2,500	hard, enamel-tightening

One more step turns the molar score into something a formulator can weigh out, and it is worth doing once in full. Phoenix's tannin runs at an average degree of polymerization near six. A flavan-3-ol monomer — catechin or epicatechin — weighs 290.3 g/mol, and each additional unit joined by an interflavan bond adds 288.3 g/mol after the loss of two hydrogens, so a six-unit chain weighs 290.3 + 5 × 288.3 ≈ **1,732 g/mol.** Its half-max in mass terms is then

> $71.4\ \mu M \times 1{,}732\ g/mol = 71.4 \times 10^{-6}\ mol/L \times 1{,}732\ g/mol \approx 0.124\ g/L \approx$ **124 mg/L**

— about a tenth of a gram in a liter, squarely inside the range real brewed and pressed plants deliver. The four-unit reference, at 1,155 g/mol, lands at 100 µM ≈ 116 mg/L. We report the Index in micromolar rather than the field's usual gallic-acid equivalents on purpose: an equivalence number counts phenol groups and cannot tell a brisk

small tannin from a harsh large one, while a molar PC_{50} holds the identity of the molecule doing the drawing. Astringency is a problem of *which* molecule, not merely *how much* phenol, and the unit has to remember that.



Tea leaf at harvest — the catechins and their polymers that draw the mouth tight. Astringency is flavor reporting on its own friction.

Two classes of draw

The tannins come in two great families, and they draw differently, so the axis has two characters living inside it.

Condensed tannins — the proanthocyanidins, polymers built from the flavan-3-ols catechin and epicatechin — give the clean, structured, drying astringency of tea and of grape seed and skin. The plant assembles them along the flavonoid branch that descends from the phenylpropanoid spine, the universal pathway PAL → C4H → 4CL → HCT → COMT that runs in every vascular plant; two terminal enzymes, **leucoanthocyanidin reductase (LAR)** and **anthocyanidin reductase (ANR)**, cast the flavan-3-ol building blocks, which are then polymerized to a controllable length. Their *size* sets their grip: short oligomers read as brisk and lively, long polymers as flatly, harshly drawing.

Hydrolyzable tannins — the gallotannins and ellagitannins, such as the punicalagins that make pomegranate so mouth-puckering — give a sharper, more immediate pull that arrives and releases faster, because their ester-linked galloyl and hexahydroxydiphenoyl groups present their binding faces all at once rather than building draw across a long backbone. Where a condensed tannin grips and holds, a hydrolyzable one strikes and lets go.

That pathway — and especially the degree of polymerization — is exactly where a design reaches in to set the amount of tannin, the length of its chains, and so the precise quality of the draw, brisk or heavy, brief or lingering. The plants we have standing today carry the two classes between them, and their construct counts record where the work has gone:

ORGANISM	CONSTRUCTS	TANNIN CLASS	SIGNATURE POLYPHENOL	DIFFICULTY
Grape	1,610	condensed (proanthocyanidins)	procyanidins · LAR / ANR	Hard
Tea	280	condensed	catechins, EGCG	Hard
Blueberry	120	condensed + anthocyanin	procyanidins + anthocyanins	Moderate
Pomegranate	120	hydrolyzable	punicalagins	Hard
Sorghum	120	condensed	proanthocyanidins	Frontier
Raspberry	40	hydrolyzable	ellagitannins	Moderate

Grape leads with 1,610 constructs, and no wonder: astringency is the backbone of wine, the structural tannin that lets a red grip and age and resolve, the quality a winemaker spends a whole career managing. Tea follows at 280, where the catechins — above all epigallocatechin gallate, EGCG, the most drawing of them — give green tea its brisk clean edge and black tea, with its larger oxidized theaflavins and thearubigins, its deeper pull. Pomegranate stands a little apart as our hydrolyzable specialist, its punicalagins among the largest single astringent molecules in the kitchen.



Pomegranate — the hydrolyzable tannins of the pith; punicalagins that arrive as a sharp, immediate pull and release fast.

The peak in the middle

Astringency is not the single sensation it first seems, and the finer our instruments grow the more it resolves into a small family of textures. Tasters reliably separate a fine, silky drawing from a coarse, grippy roughness from a hard, enamel-stripping pucker — and these track, with real fidelity, the *size* of the tannin doing the drawing. That makes chain length the master dial of the axis, and it has a shape worth knowing, because it is not monotonic.

Two things happen as a proanthocyanidin grows longer, and they pull against each other. Each added flavan-3-ol unit is another binding face, so the avidity for salivary protein climbs — multivalent binding stacks free energy roughly linearly with length, and the PC_{50}

falls, meaning the tannin draws harder. But each added unit also makes the molecule less soluble: more hydrophobic surface, more self-association, until past a critical length the tannin precipitates *on its own*, out of solution before it ever meets the saliva, and the fraction actually available to draw the mouth collapses. Net astringency is the product of these two — avidity rising, available fraction falling — and a product of a rising and a falling term has a peak.

AVERAGE DP	BINDING AVIDITY	SOLUBLE, AVAILABLE FRACTION	NET DRAW (PEAK = 1.0)	CHARACTER
2 (dimer)	low	≈ 1.0	0.35	thin, fleeting
4	moderate	≈ 0.95	0.75	clean, lively
6	high	≈ 0.85	1.00	structured grip
8	very high	≈ 0.60	0.90	firm, drying
12	very high	≈ 0.25	0.45	coarse, hollowing
16+	maximal	≈ 0.05	0.10	precipitates; little draw

The peak sits near a degree of polymerization of six to eight, and this is the quiet reason Phoenix's astringency is a *clean* frame rather than a harsh one. We did not reach for the largest tannin we could build; we set its average chain length right at the peak, near six — which is exactly the DP whose molecular weight, 1,732 g/mol, we used to turn its PC_{50} into milligrams a moment ago. At that length the draw is strongest *and* fully soluble, structured without ever turning to grit. Push past the peak and you buy nothing: the molecule only gets harder to keep in solution, the astringency actually *falling* at the heavy end while the mouth-feel turns coarse and then hollow. The Index scores the net draw, but the design aims at the peak, because the *quality* of an astringency is set by the length of the chains, and the length is a thing the polymerizing enzymes can be tuned to hold.

The numbers make the lesson concrete. A brisk green-tea catechin fraction that precipitates at PC_{50} 125 µM scores IFI-A = 100,000 / 125 = 800 — lively, and clearly drawing. Oxidize it toward the heavy thearubigin aggregates of a long-steeped black tea and intuition expects a louder draw, but much of that added mass has climbed past the solubility peak; the *effective* PC_{50} rises to roughly 170 µM and the score *falls* to 100,000 / 170 ≈ 590. Bigger is not louder once you are over the hill — and the Index records the drop plainly, exactly where a chemist counting phenol groups would have predicted a climb.



Grape in cross-section — seed and skin, where the proanthocyanidins concentrate; the structural tannin that lets a wine grip, age, and resolve.

What the mouth loses — build-up and aging

Two dynamics give astringency its peculiar behavior over time, and both follow from its being a *depletion* rather than a signal. A sweet compound you taste as strongly on the tenth bite as the first; the receptor resets. Astringency does not, because every draw carries off a portion of the very film it draws against.

The first dynamic is that it **builds.** Each sip binds and precipitates a share of the mouth's lubricating protein, and the glands take real time to replace it — salivary PRP replenishment runs on a time constant of minutes, not seconds. Model it plainly. Let each sip remove a fraction *f* of the film available at that moment, and let the deficit recover between sips by the fraction $r = 1 - e^{(-\Delta t/\tau)}$, where Δt is the spacing and τ the replenishment time. The available film settles toward a steady state

```
A* = r / [ f + r(1 − f) ]
```

For a tannic wine sipped every $\Delta t = 30$ s, with $f = 0.2$ and $\tau = 240$ s, the recovery per interval is $r = 1 - e^{(-30/240)} = 0.118$, and

```
A* = 0.118 / [ 0.2 + 0.118 × 0.8 ] ≈ 0.40
```

The film stabilizes near **40 % of full** — the mouth meeting each later sip already well more than half stripped. Since the perceived draw scales inversely with the film still present, the draw by the fifth or sixth pour runs roughly 1 / 0.40 ≈ **2.5× the first.** Astringency is the one axis that accumulates against you as you drink, and the arithmetic says plainly why a long flight of young reds leaves the mouth sandpapered while a single glass does not.

The second dynamic is that it **softens with age**, and this is the whole romance of the cellar read as chemistry. In a young red the tannins are reactive and sit near the peak, DP four to six, grippy and hard. Across years in the bottle they slowly polymerize — with one another and with the wine's own pigments — and the mean chain length drifts upward, DP five climbing toward eleven or twelve. But eleven is *past the peak.* The soluble, available fraction falls from roughly 0.85 to roughly 0.25, so the net draw more than halves even as the total tannin barely changes — the excess is literally falling out

of the wine as the sediment you decant away. The hard young draw rounds into a supple one not because the tannin leaves but because it grows too large to pull. A winemaker manages that softening across a decade; our design manages it at the source, setting the tannin's length and reactivity in the fruit so the frame is right from the first pour rather than the fifth year. Astringency, uniquely among the eight, is a flavor you feel the mouth *lose* — and the design is the quiet art of losing exactly the right amount.

Texture against taste

Astringency and bitterness are easy to confuse and important to separate, because the same dark polyphenols so often do both at once. A single tannin can draw the mouth tight *and* trip a bitter TAS2R receptor on its way past, so a harsh green tea or a raw young wine arrives as one rough impression that is really two axes braided together. The untrained palate feels a single unpleasantness; the Index refuses to.

We score them on different events on purpose. Astringency lives on IFI-A, read as the physical draw against salivary proteins — the PC_{50}, the collapse of the film. Bitterness lives on IFI-B, read as receptor activation — a true EC_{50} at a true channel. Because the two numbers come from two independent measurements, a design can move one without moving the other, and that separation is one of the quiet powers of an eight-axis index. We can hold a tea *structurally brisk* — its astringency kept near the peak for that clean clearing draw — while easing its harsh bitter edge, tuning the galloylation and the receptor-tripping shape of the same polyphenol pool so the draw survives and the bite does not. We can frame a wine with firm tannin and spare it the green, stalky bitterness that so often rides alongside. On no two axes does the untangling matter more, because nowhere else do two genuinely different sensations hide so completely inside one word.

It is worth saying where this reaches its honest limit. The two axes are measured apart, but in the living mouth they interact — a cleared, drawn palate reads a bitter note differently than a slick one does — and the Index reports the two draws, not every cross-term between them. We would rather hand you two exact numbers and the truth that they converse than one blurred number pretending the conversation away.

Where the number slips

No axis is only its triumphs, and astringency has failure modes a designer has to respect. The first is that the collapse can happen in the wrong place. A tannin set too long, or a product left too cold, precipitates in the bottle or the cup before it is ever drunk — the chill haze of an over-tannic tea, the sediment in a too-young wine — spending its draw on the glass rather than the mouth, so the computed PC_{50} and the swallowed experience part ways. The second is matrix masking: fat coats the salivary film and shields it, sugar and high viscosity blunt the drag, warmth loosens the whole system, so a tannin that reads 1,400 in water can draw far more gently in a rich, sweet, warm dish. The draw is real; the matrix negotiates how much of it reaches you.

The deepest limit is the mouth itself. We anchor IFI-A to a *standard* salivary panel, but no real saliva is standard — PRP levels vary with genetics, with diet, with the hour of the day, and a person rich in proline-rich protein meets a tannin with a larger shield and feels less draw than a person poor in it. The single number is therefore an honest measure against a conventional mouth, not a promise about every mouth. We would rather name that variation plainly than launder it into a false universality: the Index measures the tannin exactly, and tells you, at the end, where the tongue keeps the last word.

The frame, not the picture

It is fair to ask why a texture earns a place on an index of *taste*, and the answer is that astringency does what the chemical tastes cannot: it **structures** the others, in time and in relief. A rich food with no astringency to cut it turns cloying and flat; a sweet one turns sickly; the clean draw of a tannin scrapes the palate and resets it, so the next bite lands fresh and the delicate aromatics of a tea or a wine stand out sharp against cleared ground. Astringency is the frame rather than the picture, the rest between the notes — and a flavor without it, however loud on the other seven axes, reads as somehow unshaped, a sound with no silence around it.

That is why astringency is among the last axes to come online as we build a flavor outward. Across our flagship cultivars the earliest express two or three axes and carry no real draw at all; the full frame arrives only at the top tier, the full-spectrum design, where a flavor has grown dense enough to *need* framing. **Phoenix** — a fruit our model predicts at approximately 2.95 million Scoville, with eight axes at once — carries its astringency at a predicted 1,400, and in its complete profile you can read the frame sitting beneath everything else:

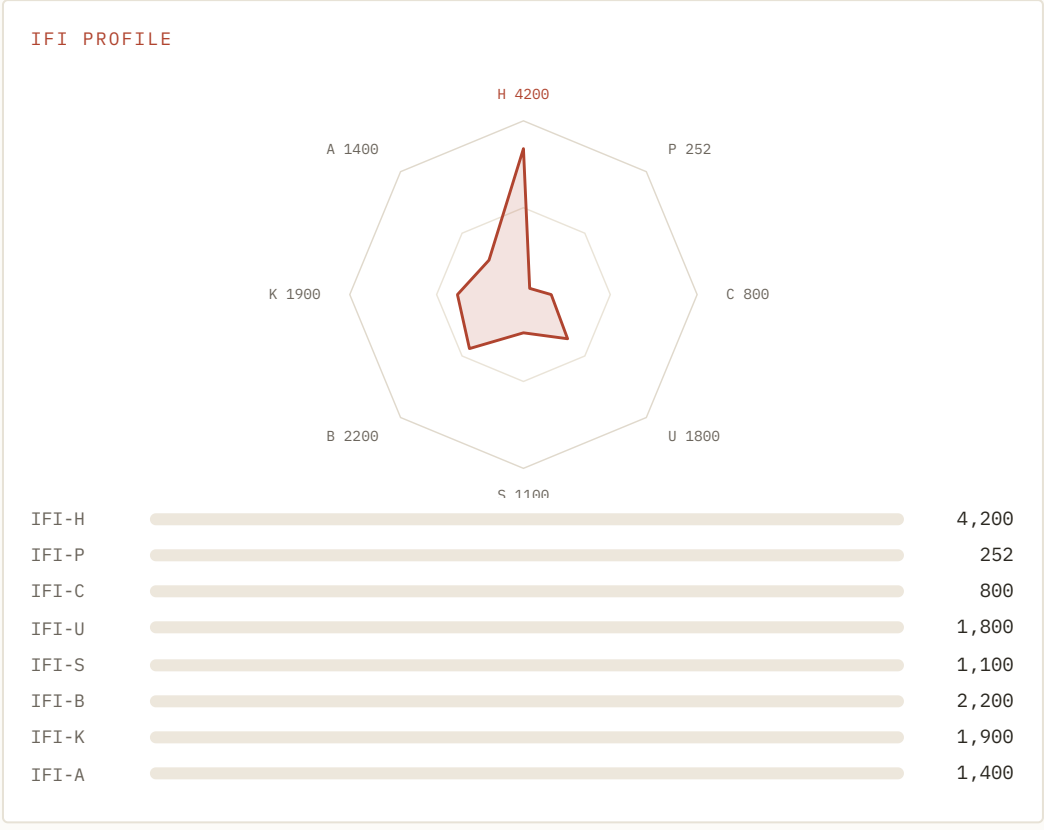

That 1,400 is not loud against the 4,200 of its heat, and it is not meant to be. It is the structural tannin laid under an enormous density — the tightening draw that gives even a computationally predicted 2.95-million-Scoville fruit a clean, resolving finish rather than a heavy, flabby one, the frame that keeps the whole from

reading as cloying. We set it, deliberately, right at the degree-of-polymerization peak, so the draw is structured and soluble and clean: present in every bite, strident in none.

The precision there is the point, and it is a design specification. Phoenix's astringency is designed to hold a batch-to-batch coefficient of variation at or under **7 %** — a one-sigma spread of about ± 98 points on that 1,400, so our design target places better than nine fruits in ten inside roughly 1,200 to 1,600. A wild tannin source runs nearer **42 %** CV: the same nominal frame would scatter from faint to savage bottle to bottle, one pour brisk and the next stripping the mouth bare. And because the Index scores a ratio against a tightly held reference, the measurement adds almost nothing of its own — propagate the error through $\mathrm{IFI\text{-}A} = 1000 \cdot \mathrm{PC}_{50}(\mathrm{ref})/\mathrm{PC}_{50}(\mathrm{c})$ and the relative uncertainties add in quadrature, $\sqrt{(0.02^2 + 0.07^2)} \approx 7.3$ %, so the number you read is essentially the batch CV and not the ruler's noise. A frame is only a frame if it holds the same shape every time, and holding that shape is the whole quiet promise of the work.

So we measure the draw the mouth cannot name. Astringency rarely leads a flavor; it frames one — the pucker that cuts a rich dish and clears the way for the next bite, the clean draw that makes a tea's aromatics legible by sweeping the ground around them, the structural tannin a plant builds for its own defense and we have learned to read, and to tune, and to set exactly where a flavor needs it. It is the one axis we feel as a loss, scored by the same half-maximal logic as all the rest, carried on the same ruler, and anchored — honestly, as it must be — to the physics of a film that falls.

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



PART III OF V

Composition

Full-spectrum design · the combinatorics of flavor · how axes interfere

Chapter 13 · Full-Spectrum — *all eight at once.*

Eight axes measured one at a time is an analysis. Eight axes sounded together, on purpose, in a single fruit, is a composition — and it is the thing the whole Index was built to make possible. This chapter is about the regime where no axis dominates the conversation, where the profile fills out toward a broad, balanced octagon, and where the engineering is not any one number but the standing of all eight in relation to each other. We call it **Full-Spectrum, IFI-F,** and it is the furthest edge of what designed flavor can presently reach.

Full-Spectrum is not a ninth axis and not an average of the other eight; it is a *verdict.* A flavor earns the IFI-F designation when it registers meaningfully on all eight pathways at once — when heat, pungency, cool-burn, umami, sweetness, zero-bitter, kokumi, and astringency are each present above the noise and composed into a coherent whole rather than fighting one another into mud. That is far harder than it sounds, and the difficulty is the whole story of this chapter: axes interfere, pathways draw on the same precursors, and a plant pushed loud on one often goes quiet on another whether you wanted it to or not. To hold all eight in balance is to have solved, at once, every trade-off the platform knows — which is why the designation is rare, and why it means something when it appears.



Phoenix at dawn under glass — the first cultivar built to register on all eight axes, and the proof that a whole flavor can be specified and then grown.

There is exactly one organism in the library that has reached the full eight in a single plant, and it is the cultivar that recurs through this book because it is the proof of the entire thesis: **Phoenix, SS-0096,** a *C. chinense* design our model predicts at approximately 2.95 million SHU, whose heat is, somehow, the least interesting thing about it. Read its profile once more — not as a scorecard, but as a composition:

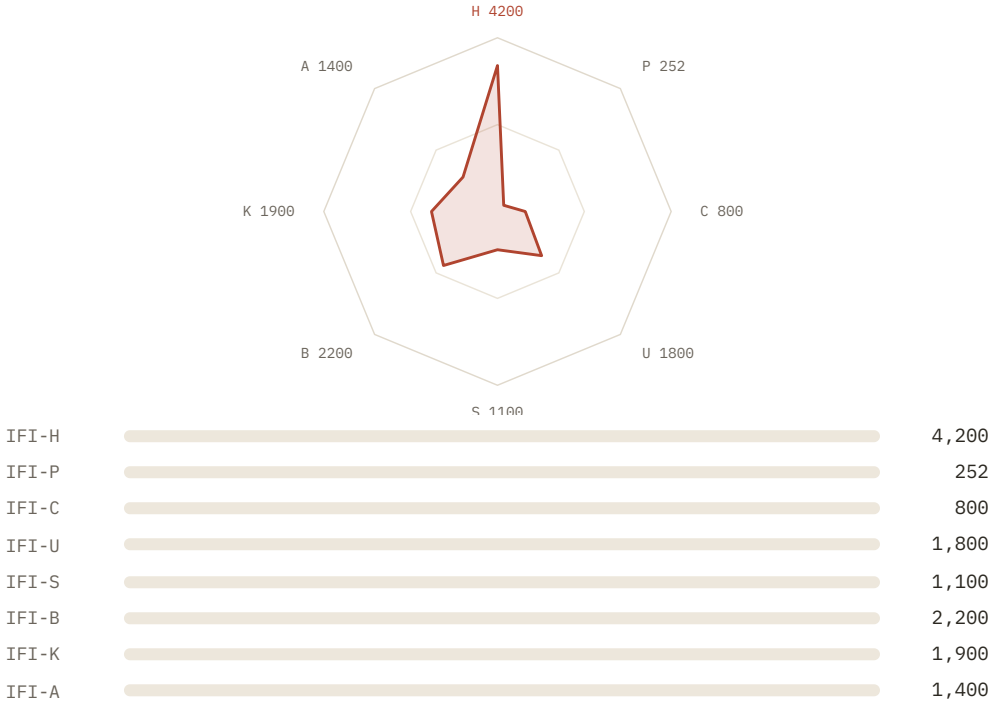

Every number there is a decision, and the decisions hang together. But before we read them as taste, we should be honest about what the Index is actually counting, because the whole of Full-Spectrum rests on it.

What a single number on an axis really measures

Each axis score is anchored to a receptor, not to a chemistry. The quantity underneath it is EC_{50} — the concentration at which a stimulus drives its receptor to half of full response — and the Index is simply that half-activation, inverted against a fixed reference and scaled by a thousand:

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

The inversion is the point. A *lower* EC_{50} means the receptor reaches half-response at a *lower* concentration — the compound speaks more loudly — and inverting turns that quiet physical number into a score that climbs with loudness. At $EC_{50}(c) = EC_{50}(ref)$ the formula returns exactly 1,000, which is why 1,000 is the native-reference anchor on every axis. Double the loudness — halve the EC_{50} — and the score doubles to 2,000. The relationship is clean because it is a simple reciprocal, and it lets us run it backwards: any score on the Index *is* a statement about a concentration.

Each reference constant is a **stated convention** — defined in the same spirit as Scoville's dilution-to-threshold: not a value drawn from literature, but a fixed anchor point that makes the scale meaningful. Setting IFI = 1,000 at a chosen reference EC_{50} gives every axis its own calibrated origin. For the pungency axis (TRPA1), the convention constant is **2.1 µM**, so the formula reads: IFI-P = 1000 · 2.1 µM / $EC_{50}(c)$. IFI-H carries a different convention constant, as does each other axis. Because each axis is anchored to its own defined reference, scores across axes cannot be directly compared as if they shared a single common scale — each is a loudness relative to its own convention, not to a universal baseline.

Take Phoenix's heat at 4,200. Rearranging the formula,

$$EC_{50}(c) = 1000 \cdot EC_{50}(ref) / IFI(c) = EC_{50}(ref) \times (1000 / 4200) = 0.238 \cdot EC_{50}(ref)$$

so the capsaicin in Phoenix half-activates TRPV1 at **23.8 % of the reference concentration** — it is 4.2× louder than native. Capsaicin's half-activation of TRPV1 sits, as a matter of ordinary pharmacology, in the hundreds of nanomolar; a 0.238× shift moves it down into the low hundreds. The Index never asks us to carry those nanomolar in our heads. It asks only for the ratio, and the ratio is exact.

This is also why the Index measures loudness rather than quantity. A dose–response curve has two independent facts in it: how *far left* it sits (its EC_{50}, its potency) and how *tall* it rises (its maximum). Written as a Hill function,

$$R([L]) = R_{max} \cdot [L]^h / (EC_{50}^h + [L]^h)$$

the half-maximal point falls at [L] = EC_{50} regardless of R_max or of the steepness h. By anchoring to that point, the Index reads the curve's horizontal position — the thing a tongue experiences as *loudness* — and stays blind to how much compound a food happens to hold. A trace of a very loud molecule and a flood of a quiet one can land on the same axis score. That independence is what makes eight axes composable at all: each is a position on its own curve, and positions can be set one at a time.

It is worth separating this from Scoville heat directly, because the two are often confused. A Scoville value is a dilution-to-threshold: our model predicts approximately 2.95 million SHU for Phoenix, meaning its heat is designed to remain detectable after roughly 2.95 million-fold dilution — a measure dominated by *total pungent load.* IFI-H is a per-molecule loudness. A food can move one without the other: dilute Phoenix's flesh and its SHU falls while the capsaicin's EC_{50}, and so its IFI-H, does not change at all. The Index tracks the molecule's voice; Scoville tracks how much of it is in the room.

Reading Phoenix as a composition

Now the profile reads cleanly. Run each axis back through the same inversion — EC_{50} ratio = 1000 ÷ IFI — and the design intent becomes visible as a column of concentrations, one per receptor.

AXIS	RECEPTOR	IFI	EC_{50} ÷ EC_{50}(REF)	WHAT IT DOES ON THE TONGUE
H Heat	TRPV1	4,200	0.238	the fire, leading and loudest — 4.2× native
B Zero-Bitter	TAS2R38	2,200	0.455	presses the pepper's alkaloid harshness down hard
K Kokumi	CaSR	1,900	0.526	continuity and body; the heat rides on depth
U Umami	mGluR1 / TAS1R1·R3	1,800	0.556	savory ground so the burn isn't empty
A Astringency	salivary PRPs	1,400	0.714	a clean, drying frame on the finish
S Sweetness	TAS1R2 / TAS1R3	1,100	0.909	a thread of kindness, just above reference
C Cool-Burn	TRPM8 / TRPV1	800	1.250	the one axis pulled *below* reference — TRPV1's opponent, held quiet
P Pungency	TRPA1	252	3.97	a mild pungent frame, deliberately restrained to frame the heat rather than contest it

Six axes sit at a ratio below 1.0 — louder than their convention reference, lifted in careful proportion. Two do not. Cool-burn carries a ratio of 1.250, meaning its half-activation requires 25 % more concentration than the TRPM8 reference — it is deliberately held quiet. Pungency sits lower still, at a ratio of 3.97, requiring nearly four times the AITC convention concentration to half-activate TRPA1 — also a deliberate choice. The physiological reason is the same in both cases. TRPM8, the cold-and-menthol channel, is a functional opponent of TRPV1; a loud cooling signal fights the fire, and the two blurred together read not as complexity but as confusion. Pungency, arriving as a sharp front-note at high intensity, would compete with the heat's slow bloom for attention and muddy the leading arc of the experience. So we pulled both below their reference lines, and let the other six rise around a heat that nothing fights against. Eight axes; six raised, two pressed below the line; a complete flavor, specified in advance and grown to the specification. That octagon is the most ambitious object our family has ever designed, and the surprise in it is never the heat. The surprise is the composition built underneath the heat.

Why the composite is a verdict and not an average

Here is the temptation, and the trap: given eight numbers, average them. Phoenix's eight scores sum to 13,652, so the arithmetic mean is 1,707. Report 1,707 as the "full-spectrum score" and move on.

We do not, because the arithmetic mean is blind to the one thing the designation is *about.* An average lets a loud axis buy back a silent one. A profile that screams on three pathways and whispers on five can post the same mean as a profile that speaks evenly on all eight — and those are not the same flavor. One is balanced; one is a spike wearing a balanced number.

The honest operator is the **geometric mean,** the eighth root of the product of the eight scores:

```
IFI-F = ( H · P · C · U · S · B · K · A )^(1/8)
```

Its virtue is exactly its sensitivity to the quiet axis. Because every factor multiplies, a single small number drags the whole product down, and if any axis falls to zero the geometric mean collapses to zero — mathematically encoding the rule that *full-spectrum requires all eight present.* An average can hide a missing axis; a product cannot. Watch it separate two profiles built to have the identical arithmetic mean of 1,707:

PROFILE	H	P	C	U	S	B	K	A	Σ	ARITH. MEAN	GEOM. MEAN	COHERENCE Φ = GM/AM
Phoenix SS-0096	4200	252	800	1800	1100	2200	1900	1400	13,652	1,707	**1,330**	0.779
Illustrative three-spike	4700	80	80	80	80	4600	4400	80	14,100	1,763	**369**	0.209

Both profiles post a moderate arithmetic mean. The geometric mean puts 1,330 next to 369 — a 3.6× gap — because the second profile pays for three axes at or near the ceiling by letting five fall to the noise floor, and the product punishes every one of those floors. The ratio of the two means, **Φ = GM / AM,** is itself a clean coherence index: it runs from 0 (one spike, the rest silent) toward 1 (a perfectly regular octagon). Phoenix sits at 0.779 — a strong, well-distributed eight — while the spike sits below 0.21. IFI-F is awarded through two gates, not one: a *presence* gate, every axis above a floor set where a reading is cleanly resolved from assay noise, and a *coherence* gate, Φ high enough that the eight hang together. The three-spike profile clears neither. It is, by construction, exactly the failure the composite exists to catch.

The precision of the whole is better than the precision of the parts

There is a reward buried in that product, and it is the kind of result that makes the mathematics feel like a gift. Our per-axis design specification carries a batch-to-batch coefficient of variation target of **≤7 %** — tight, the fruit of controlled-environment growth and the same reproducibility work that pulled wild-type's ragged 42 % CV down to single digits. The question is what happens to that error when eight axes are folded into one composite.

Take logs, because the geometric mean is an arithmetic mean in log-space: $\ln(\text{IFI-F}) = (1/8) \cdot \Sigma \ln(\text{IFI}_i)$. For independent axes the variance of a sum of scaled terms is the scaled sum of variances, and for small relative error the variance of a log is the square of the CV, so

> $\text{Var}(\ln \text{IFI-F}) \approx (1/8^2) \cdot \Sigma\, (0.07)^2 = (0.07)^2 / 8 \rightarrow \text{CV(IFI-F)} \approx 0.07 / \sqrt{8} \approx \mathbf{0.025}$

The composite is known to about **2.5 %** — roughly three times tighter than any single axis that went into it. This is the familiar $\sqrt{n}$ shrink of averaging, and it means the full-spectrum verdict is not merely a summary of eight noisy numbers but a *steadier* number than any of them. We will be honest about the assumption: it holds cleanly only if the eight axes vary independently, and they do not entirely, because they compete for the same metabolism and their receptors cross-talk. Real correlations inflate the figure, so 2.5 % is

a floor, not a promise. But the direction is real, and it is why a composed flavor can be held to a tighter standard than its own ingredients.

Eight receptors, eight syntheses, one finite budget

To build all eight loud in one plant is to run eight expensive syntheses at once in one metabolism — and to make them share. The receptors are genuinely different machines, and so are the molecules that speak to them.

Heat is **capsaicin** on **TRPV1,** a non-selective cation channel on the trigeminal endings of the tongue; the molecule does not warm the tissue, it lowers the channel's thermal threshold so the body reports a fire that is not there. Capsaicin is assembled when vanillylamine condenses with an 8-methyl branched-chain fatty acid at the Pun1/AT3 acyltransferase. Pungency is a faster, sharper channel entirely — **TRPA1,** gated by the allyl isothiocyanates of mustard and wasabi, by allicin, by cinnamaldehyde — which is why a pungent compound, when present and loud, reads as a bright edge arriving *ahead* of H's slow bloom. Cool-burn is **TRPM8,** the menthol-and-cold channel, TRPV1's opponent. Umami is glutamate read by **mGluR1** and the TAS1R1·TAS1R3 pair — and it carries its own piece of hidden arithmetic, because 5′-ribonucleotides bind an adjacent site and multiply the glutamate signal rather than add to it; a little inosinate can swing umami loudness several-fold with no more glutamate present at all, an allosteric potentiation the Index reads as a genuine leftward shift in EC_{50}. Sweetness is the **TAS1R2·TAS1R3** heterodimer answering to sugars and to sweet proteins like brazzein (UniProt P56552) and thaumatin. Zero-bitter works the other way — not a flavor added but a flavor removed, a 20–65 % attenuation of **TAS2R38** signaling that lifts the alkaloid harshness off the palate. Kokumi is the quiet marvel: γ-glutamyl peptides on the **calcium-sensing receptor, CaSR,** contributing no taste of their own but lending continuity, mouthfeel, the sense of a flavor that lasts. And astringency is barely a taste at all — tannins and polyphenols precipitating the **salivary proline-rich proteins,** stripping the tongue of lubrication so the percept is mechanical, a drying, a frame.

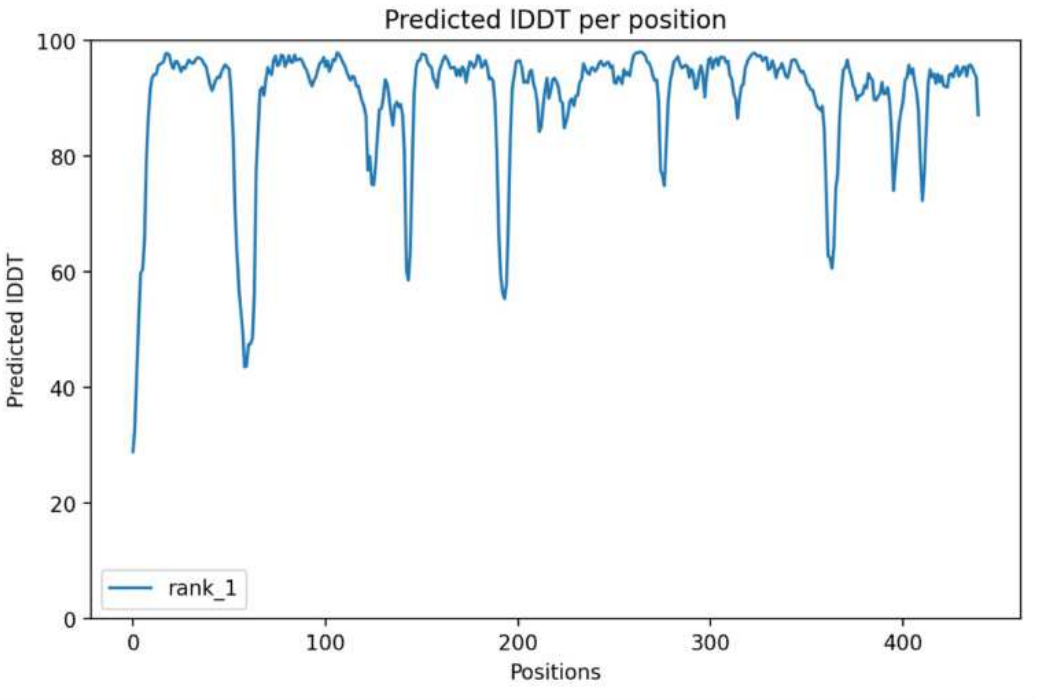

Pun1/AT3 coloured by model confidence — the acyltransferase at the end of the heat pathway, and the branch point where carbon bound for capsaicin is carbon not bound for tannin.

The budget problem is that these syntheses do not start from separate pantries. The phenylpropanoid branch — **PAL → C4H → 4CL → HCT → COMT** — runs upstream of both the vanillylamine that becomes capsaicin (heat) and the flavonoid route that becomes condensed tannin (astringency). Push H and A loud at the same time and they divide one pool of p-coumaroyl carbon between them. Kokumi's γ-glutamyl peptides draw on the same glutamate pool that umami spends; sweet proteins draw on the amino-acid and translation budget; cooling terpenes and pungent isothiocyanates each demand their own carbon, sulfur, and reducing power. Eight loud axes is eight claims on one finite supply of carbon, nitrogen, sulfur, and ATP — and if the plant honors all of them past what it can afford, it does not make a fuller flavor, it makes a sick plant with a ragged CV. Phoenix is not merely a plant that does eight things; it is a plant in which eight competing demands were balanced against one metabolic budget and made to coexist without any of them collapsing.

The climb to eight

That balancing is why the full eight sat at the *top* of the design ladder and nowhere lower. The flagship line is a record of axes accumulating as the heat climbed — each cultivar a step, each step adding pathways the plant could afford to carry.

CULTIVAR	SKU	PREDICTED SHU	IFI-H	EC_{50} ÷ REF	× LOUDER	AXES
Ember	SS-0087	866K	1,570	0.637	1.57×	2 / 8
Blaze	SS-0091	1.72M	2,630	0.380	2.63×	4 / 8
Wildfire	SS-0092	1.96M	2,930	0.341	2.93×	5 / 8
Firestorm	SS-0094	2.47M	3,490	0.287	3.49×	6 / 8
Solara	SS-0095	2.72M	3,980	0.251	3.98×	6 / 8
Phoenix	**SS-0096**	**~2.95M**	**4,200**	**0.238**	**4.20×**	**8 / 8**

The progression from Ember to Phoenix is not a certified ascending measurement; it is an increase in computational design scope and axis activation — each step adding pathways our model predicts the plant can sustain without its CV degrading past the ≤7 % specification. The heat was designed, not found. An ESM2 650-million-parameter protein language model scored every single-residue substitution in Pun1 by predicted effect on catalytic efficiency × fold stability; a genome-scale flux-balance model then mapped where metabolic control actually sits — and our flux analysis found it upstream, at **4CL and HCT/CCoAOMT** in the phenylpropanoid pathway, not at the terminal Pun1/AT3 step. Pun1 is the final capsaicin synthase; it labels the product, it does not govern the rate. Our Metabolic Control Analysis placed Pun1's flux control coefficient at effectively zero across 5,000 Monte Carlo draws; the upstream phenylpropanoid steps carry 93–99 % of pathway control. This is the same flux accounting that governs whether a second or third axis can be afforded at all. Two substitutions emerged from the protein language model scoring — **S39L** (ΔΔG −2.1, binding affinity up at the N-terminus) and **L345G** (ΔΔG −1.8, steric relief in the

active-site cleft) — forming the **double variant S39L+L345G** that is the designed Pun1 construct. Our model predicts this design achieves approximately 2.95 million SHU at a CV of ≤7 % against wild-type's 42 %. Reaching IFI-H 4,200 was the price of entry; the real work was that, at the summit, there was finally enough metabolic headroom and enough control over the upstream pathway to raise seven more axes around the fire without any of them guttering. The summit is proof that the mountain can be designed. It is not how most goods get carried.

Co-delivery: how full-spectrum becomes practical

Because the single-plant full eight is a landmark and not a logistics plan, the library's broadest work travels a second road: **co-delivery.** Rather than burden one organism with every axis, we compose a flavor by combining designed inputs — each organism contributing the axes it speaks most cleanly — and, at the frontier, by **nanoparticle delivery:** packaging the responsible molecules into tailored carriers so they arrive together, in the right proportion, at the right place, releasing in the sequence the composition calls for. This is the engine behind the roughly **300 full-integration products** in the catalog, each assembled to hit a target profile across all eight axes without demanding that any single plant carry the whole load — drawn from a library of approximately **225 distinct construct designs** deployed across 42 host organisms, 1,000 products live and purchasable today.



Golden-hour under glass — co-delivery is what turns one summit cultivar into a working range.

The allocation is the quiet half of the craft, and it can be written out like a budget. Suppose the target is a full-eight profile and no single plant is asked to carry it. We assign each axis to the organism that speaks it most cleanly: heat to **Pepper,** which speaks all eight but is loudest on H; pungency and kokumi to **Garlic,** whose native register is exactly K and P; cool-burn to **Peppermint;** sweetness and umami to **Tomato,** restored along the S and U it was bred away from; zero-bitter and astringency to **Grape,** whose home axes are B and A. Five designed organisms, each contributing only what it does best, together cover all eight — H · P · C · U · S · B · K · A — without any one of them straining its metabolism past a clean CV. The burden

that broke across one plant is spread across five, and the quiet axes stay quiet because no organism is being pushed to speak an axis it was never built for.

The nanoparticle side deserves more than a phrase, because it is how a composed flavor actually reaches a tongue — and it is also what re-times an allocation like the one above, since five organisms harvested separately do not arrive on the palate in any particular order until the carriers give them one. The responsible molecules differ wildly in nature — capsaicin is oil-loving and stubborn, a sweet protein is water-soluble and fragile, a γ-glutamyl peptide is small and quick to wash away — and simply blending them yields a flavor that falls apart in handling and arrives in the wrong order. So we encapsulate each molecule, or each small set, in a carrier sized in the tens to hundreds of nanometers and engineered for *when* and *where* it should release.

Here the mathematics turns from receptors to kinetics. Model a carrier's release as first-order — the simplest honest model, a constant fractional loss per unit time — and the fraction released by time *t* is

```
f(t) = 1 − e^(−t/τ), with half-release at t½ = τ · ln 2 ≈ 0.693 τ
```

where τ is the carrier's time constant. Sequencing a flavor's arc is then just staggering τ across carriers. Give pungency a fast carrier, τ ≈ 0.7 s, and its bright edge half-arrives in about 0.49 s — the front note. Give sweetness τ ≈ 2 s (t½ ≈ 1.4 s) so a thread of kindness lands just behind it. Hold umami and kokumi on τ ≈ 5 s (t½ ≈ 3.5 s) for the body. Hold the heat longest, τ ≈ 10 s (t½ ≈ 6.9 s), so the fire blooms only *after* the sweetness has registered — reproducing by design the rise and fall a great dish has by nature. The time constant is in turn a matter of geometry: for a surface-controlled carrier the release rate scales with surface-to-volume ratio, which for a sphere is 3/r, so a 25-nm-radius carrier sheds its cargo roughly twice as fast as a 50-nm one. Tune the radius, tune the timing. Controlled release is the difference between a chord struck flat all at once and a flavor that unfolds — and unfolding is much of what makes a composed flavor feel alive rather than assembled.

Failure modes, edge cases, and honest limits

The composer's freedom is also the composer's burden, and it is worth naming exactly how full-spectrum goes wrong.

The greedy octagon. The first failure is wanting everything loud. Nothing in the method stops us from pushing every axis toward the per-axis ceiling of 5,000, and the result is not richness but mud — eight signals at full cry, none of them legible, a plant bankrupted paying for all of them. The 5,000 ceiling is not a target; it is a reminder that the tongue saturates, that past a point more concentration buys no more perceived loudness, and that a profile which maxes all eight is noise wearing an octagon. A full-spectrum flavor is not one that maxes every axis; it is one in which all eight are

right, which almost always means several deliberately restrained — Phoenix's cool-burn and pungency held below their reference lines are the model, not the exception.

Axes that fight. Some pairs are physiologically opposed, and loudness on both cancels rather than composes. Heat and cool-burn are the clearest case — TRPV1 and TRPM8 pulling against each other — which is why C is the axis Phoenix pulls quiet. A composite can post eight respectable numbers and still fail if those numbers quarrel: a heat that cancels a cool that muddies an umami. This is precisely why the coherence gate, Φ, is not optional. Eight present axes are necessary; eight *agreeing* axes are what IFI-F certifies.

The limit the profile cannot see. The honest edge of the whole system is this: the eight numbers are a static snapshot, and taste is not static. The Index scores the chord; it does not score the arc. Two products with identical IFI-F profiles — the same eight numbers, the same geometric mean — can taste different if their molecules release in a different order, because the profile encodes *how loud* each axis is and says nothing about *when* it speaks. That is not a flaw to apologize for; it is the exact seam where the nanoparticle kinetics above do their work. The profile specifies the destination; the carriers specify the journey. A complete account of a composed flavor is the eight-axis target *and* its release schedule, and only the two together determine what a person tastes.

Why there is still only one. All of which is why, for all the method's reach, exactly one organism has been designed to reach the full eight in a single plant. The metabolic budget is real, the receptor cross-talk is real, and co-delivery is simply the wiser road for almost every product we make. The summit cultivar proves the thing is possible in one organism; co-delivery makes full-spectrum flavor practical across the whole range of food. That is the difference between a landmark and an industry, and we are content to have designed both.

What Full-Spectrum finally demonstrates is that flavor is composable — that the eight axes are genuine independent handles, that a profile can be drawn on paper as a target shape and then realized, and that "a food that tastes exactly like this" is now a sentence with an engineering answer. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. It is a different relationship to food than cooking has ever had. The cook adjusts what already exists; the composer specifies what will. There is a humility that has to ride inside that ambition, and it is the oldest thing our family believes about a table: to specify a flavor is to take responsibility for every axis of it, including the ones you choose to keep quiet, and the best compositions, like the best cooking, are known as much by what they leave soft as by what they make loud. A whole flavor, grown with intention, to nourish the body that receives it — that was always the point. The octagon is only the proof that it can be written down.

Chapter 14 · The Combinatorics of Flavor — *axes interfere.*

The clean premise of the eight axes is that they are independent — separate receptors, separate references, separate curves, each one read on its own ruler. At the point of measurement that premise is exactly true, and it is the whole reason the Index works. But it stops being true the instant you try to *build* two or three axes loud in one fruit, because a plant is not eight instruments standing in a rack. It is one metabolism — one pool of precursors, one finite budget of carbon and light-made energy, one living body that has to go on growing while it flavors — and the axes compete for every bit of it. This chapter is about that competition: the real, countable reason that most cultivars speak two or three axes and not eight, and the arithmetic we use at our table to decide which combinations are worth the fight.

The distinction is the whole chapter, so let us set it down precisely. The interference is not in the *reading.* At the receptor there is no cross-talk worth the name: capsaicin arriving at TRPV1 does not know, and could not care, what menthol is doing at TRPM8 a few microns away; a dose–response curve measured for one axis is unmoved by the presence of the others. That is why the Index can isolate each axis, score it, and trust the number even when the thing on the tongue is a dense chord. The interference lives one step upstream, in the *making* — and it lives in three particular places: the shared pathways that feed more than one axis, the finite metabolic budget every axis is drawn from, and, now and then, the palate's own perceptual cross-talk. A responsible design has to model all three before it commits a living organism to a multi-axis target.

Where independence is real — and where it isn't

Begin with where independence genuinely holds, because the Index leans its whole weight on it. Each axis is anchored to a different receptor — heat to TRPV1, sweetness to TAS1R2·TAS1R3, umami to mGluR1, zero-bitter to TAS2R38, kokumi to the calcium-sensing receptor CaSR, cool-burn to TRPM8, pungency to TRPA1, astringency to the salivary proline-rich proteins. These are separate proteins reading separate ligands, and their activation curves are, to the precision we can measure, orthogonal. Score a fruit's heat and you have learned nothing and spoiled nothing about its sweetness. This is the quiet miracle that lets one ruler span a shishito and a superhot, a cup of tea and a square of chocolate: the reading of any axis is clean because the receptor doing the reading is deaf to the rest.



FIG. 15 — A multi-axis run laid out as a specimen tray. Two axes loud in one fruit is a flux problem long before it is a flavor.

The trouble is that the fruit was never grown one axis at a time. The three sites of interference are easy to name and harder to tame. The first is **shared precursors:** several axes draw their molecules from the same upstream metabolism, so loudness in one can quietly starve another of the carbon it needed. The second is the **budget:** every flavor molecule is carbon and energy the plant did not spend on growth, and there is a ceiling on how much of its total metabolism a healthy plant will divert before it suffers. The third, smaller and occasional, is **perceptual cross-talk** at the palate — a loud bitterness masking a sweetness the receptors both reported honestly, or kokumi lending its gain to umami so the broth reads louder than its own curve predicts. The first two are the ones a design must solve on paper, before an edit is written, and they are where the real mathematics lives.

The Index as arithmetic — and what "combined" means

Everything in this chapter sits on one formula, so hold it where you can see it. An axis score is an inverse of a concentration:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the half-activation point. The reference for each axis is a fixed native standard defined by convention — each reference constant establishes IFI = 1,000 on its axis, in the same spirit as the dilution threshold that defines a Scoville unit — whose EC_{50} sets the scale. Divide the reference EC_{50} by the candidate's, scale by a thousand, and loudness becomes a pure number. The arithmetic has two consequences worth saying out loud. First, the reference scores itself exactly 1,000 — set candidate equal to reference and the ratio is one. Second, because EC_{50} sits in the denominator, loudness and concentration run opposite ways: a candidate that half-activates its receptor at *lower* concentration is the *louder* one, and halving an EC_{50} doubles the score.

Work it forward on the axis everyone knows. Our full-spectrum flagship **Phoenix (SS-0096)** carries IFI-H of 4,200. Read the formula backward and that score *is* a statement about concentration: 4,200 = 1,000 × EC_{50}(ref) ÷ EC_{50}(Phoenix), so EC_{50}(Phoenix) = EC_{50}(ref) ÷

4.2 = 0.238 × EC_{50}(ref). Our model predicts Phoenix's capsaicinoid design reaches half-activation of TRPV1 at a little under a quarter of the reference concentration — it is, in our computation, 4.2 times louder than the native standard. The score is not a rating anyone assigned; it is a precisely computed ratio of half-activation concentrations — for our designed cultivars, the candidate EC_{50} is a computational prediction — and it inverts cleanly in either direction.

Run the inversion down the whole superhot ladder and a second story appears beside the first. We express everything in *reference units* — EC_{50} as a multiple of that axis's reference, so EC_{50}/E_ref = 1,000 ÷ IFI — which keeps the arithmetic honest and free of any borrowed physical constant. (For a felt sense of scale: capsaicin's half-maximal concentration at TRPV1 sits in the low hundreds of nanomolar, the kind of potency that makes a receptor a reasonable thing to build a ruler on.)

CULTIVAR	SKU	PREDICTED SHU	IFI-H	EC_{50} ÷ E_REF	AXES
Ember	SS-0087	866K	1,570	0.637	2/8
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

Read the last two columns together and the chapter is half-written already. As heat climbs from 1,570 to 4,200, the EC_{50} falls from 0.64 to 0.24 of reference — the design grows more than twice as loud at the receptor. But the predicted SHU, the raw capsaicinoid content, climbs faster than the loudness does: from Ember to Phoenix the predicted Scoville figure triples (866K to 2.95M, 3.41×) while IFI-H only 2.68×. The receptor saturates; the last capsaicinoid molecules are heard less loudly than the first. That gap between *how much a fruit holds* and *how loudly it speaks* is exactly the gap the Index was built to measure — and it is the first hint that loudness is not a thing you simply add.

Which brings us to the word that makes this a chapter on *combinatorics.* When a design engages more than one axis, we speak of **combined loudness** — the plain sum of the engaged axis scores, L = Σ IFI_i. In the reading, this sum is a design budget metric: each axis was scored at its own deaf receptor against its own defined reference convention, so the values add cleanly as a combined budget. Each axis carries its own scale; a score of 1,000 on the heat axis is not the same magnitude as a score of 1,000 on the pungency axis, because each reference is separately defined by convention. In the *making,* the sum is the bill, and the bill is where the metabolism pushes back.

The shared spine — competition at the precursor

The deepest interference is molecular, and it is worth drawing step by step, because the whole claim of feasibility rests on it. Heat and astringency are not built from scratch in separate corners of the cell;

they branch off the same upstream highway — the phenylpropanoid spine that every vascular plant runs. It opens with **PAL** (phenylalanine ammonia-lyase) stripping phenylalanine to cinnamic acid, then **C4H** hydroxylating it to *p*-coumaric acid, then **4CL** activating that to **p-coumaroyl-CoA.** That CoA-thioester is the fork in the road, and both of our axes want it.

Down one branch, toward heat, p-coumaroyl-CoA is carried through **HCT** and **COMT** and on to vanillin, which **pAMT** transaminates to vanillylamine. Vanillylamine then meets an 8-methyl branched-chain acyl-CoA — built separately from branched-chain amino acids by **BCAT** and extended by the fatty-acid synthase **KAS** using **malonyl-CoA** — and the two halves are condensed by **Pun1/CS**, capsaicin synthase, into capsaicin itself. Down the other branch, toward astringency, that same p-coumaroyl-CoA is handed to **chalcone synthase,** which extends it with *three* molecules of malonyl-CoA into naringenin chalcone, the gateway to the flavonoids; downstream, **ANR** (anthocyanidin reductase) and its neighbors deliver the flavan-3-ols — catechin and epicatechin — that polymerize into the condensed tannins the salivary proteins read as astringency.

Now see the competition plainly. The two branches share *two* currencies, not one. They both draw from the **p-coumaroyl-CoA** pool at the fork, and they both spend **malonyl-CoA** — heat for the acyl chain of capsaicin, astringency for the three-fold extension of every flavonoid. Push heat loud and you pull p-coumaroyl-CoA and malonyl-CoA toward capsaicin; the tannin branch, bidding for the same two pools, comes up short, and astringency falls whether or not you intended it. The reverse holds just as firmly. This is not a metaphor about scarcity. It is stoichiometry: two sinks on a shared source, and a source that can only be divided, not multiplied by wishing.

Feasibility before edits — the flux-balance program

Because the competition is stoichiometric, we treat it as stoichiometry — literally, as a system of equations solved before a single edit is designed. We build a genome-scale metabolic model of the host: a great matrix **S** with one row per metabolite and one column per reaction, each entry recording how much of that metabolite the reaction makes or consumes. Let **v** be the vector of reaction rates. At metabolic steady state nothing accumulates inside the cell, so production must equal consumption for every internal metabolite at once, which is to say

$$\mathbf{S} \cdot \mathbf{v} = 0$$

— a wall of simultaneous equations that any valid metabolism has to satisfy. Each reaction also has capacity bounds, $v_{lower} \le v \le v_{upper}$, set by enzyme levels and thermodynamics. Within that feasible space we pose an objective — maximize the flux toward a target flavor while holding the plant's growth rate above a floor so the organism stays vigorous — and solve the resulting **linear program** for the flux through every pathway together. The answer is not a guess; it is the simultaneous, mass-balanced best the metabolism can do.



FIG. 16 — The bench where feasibility is decided. We find the flavors a plant *can* make, on paper, before we reach for CRISPR.

Two readings of that program do the real work. The **shadow prices** — the dual variables attached to each metabolite's balance equation — report the marginal value of relaxing each constraint by one unit, which is a precise way of naming the true *rate-limiting step:* the one reaction whose capacity actually caps an axis. Our flux model predicts 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. Edit at the upstream control points and the number moves; edit only at the terminal synthase and the metabolism simply reroutes around you. Much of why our designs land is that the model first tells us where *not* to bother. The second reading is **flux-variability analysis:** axis by axis, we maximize and then minimize each output flux subject to S·v = 0, the capacity bounds, and the growth floor, and the pair of answers draws a *feasible envelope* — the range an axis could occupy before some shared constraint bites. A target profile that falls outside that envelope is rejected on paper. We find the achievable flavors first, and only then grow one.

A toy version shows the exchange rate falling straight out of the algebra. Let a single precursor **P** be produced at rate F and consumed only by the heat branch (flux v_H) and the astringency branch (flux v_A), one unit of P each. Steady state on P is a single equation: F − v_H − v_A = 0, so **v_H + v_A = F.** The feasible frontier is a straight line. If the shared flux supports, say, F = 100 units, then every unit of P routed to heat is a unit denied to astringency: dv_A/dv_H = −1, an exact one-for-one trade in *flux* space. The real models carry thousands of reactions rather than two, but the shape is the same — a shared source divided between competing sinks, and a frontier the design must respect rather than wish away.

Pricing the trade — the heat–astringency exchange rate

A one-for-one trade in flux is not a one-for-one trade in *loudness,* and the difference is where the craft lives. The receptors answer their ligands along saturating curves — sigmoids, not straight lines — so a unit of flux buys a large jump in IFI low on the curve and a small one near the top. The loudness exchange rate is therefore the flux exchange rate (−1, from the shared pool) scaled by the ratio of the two axes' *marginal receptor gains,* dIFI/dv. Written out,

dIFI_A/dIFI_H = −(dIFI_A/dv_A) ÷ (dIFI_H/dv_H): the price of a point of heat, measured in astringency, is set by how steep each axis's receptor curve is at the moment you make the trade.

Put illustrative numbers on it with a symmetric receptor map and the behavior is immediate. When both branches sit mid-curve, driven to equal moderate loudness, their gains match and the trade is gentle — **one point of astringency for one point of heat,** a fair and quotable exchange. Drive heat up toward saturation, though, and its curve flattens: now each additional point of IFI-H demands a large slug of flux, and that flux is torn from an astringency branch still on the steep part of *its* curve, where a little flux was worth a lot of loudness. The exchange rate steepens past two- and three-to-one — so the last few hundred points of IFI-H cost a visible slice of IFI-A, while the first few hundred cost almost nothing. This is why the flux model is consulted *before* a plant is committed and not after: the top of every axis is the expensive stretch, and only the model, reading its shadow prices, can tell you the exchange rate before the ripened fruit tells you for free. Multi-axis design is a problem of **voicing, not stacking.** You are not adding independent loudnesses; you are spending one budget across competing demands, and the whole art is in knowing the exchange rates before you spend.

The budget, and the modes it buys

Behind every exchange rate stands a ceiling. Every molecule a plant makes for flavor is carbon and energy it did not spend on growth, and a healthy plant will divert only so much of its total metabolism before it begins to suffer — and a suffering plant makes a poor, inconsistent flavor, which is no flavor a formulator can build on. Load eight axes loud and you are asking one fruit to carry eight expensive syntheses at once, past the point the body can bear. This is why the deliberate multi-axis designs — the **IFI-M** family — are scored on their *relationship,* not on any single number, and why the budget shows its hand so clearly across the three modes.

MODE	AXES ENGAGED	COMBINED LOUDNESS	AVG PER AXIS	CHARACTER
M2 — dual-axis	2	up to 8,400	4,200	two loud axes, sharply composed
M3 — triple-axis	3	≈ 2,800 balanced	≈ 933	three axes held in proportion
IFI-F — full-spectrum	8	13,652 (Phoenix)	1,707	the complete octagon

Read the average column and the budget speaks. A dual-axis **M2** design reaches a combined loudness near **8,400** across its two engaged axes — which, split evenly, is 4,200 apiece. Hold that against Phoenix: an M2 can run *two* axes each as loud as Phoenix's single loudest axis, precisely because it is paying for only two and not eight. A triple-axis **M3** sits far lower, near **2,800** in balanced combination — an average of roughly 933 per axis, which is *below* the reference loudness of 1,000. That looks like a paradox until you see what M3 is buying. It is not buying amplitude; it is buying *proportion.* Holding three axes in a deliberate ratio means spending budget on regulation and balance — fighting the metabolism's

natural tendency to let one branch run away — so less of the budget survives as raw loudness. You buy breadth and poise with loudness, and the Index shows you the exchange to the point.

A swicy design makes the two-axis case concrete — heat and sweetness composed into a single sharp chord, the register of a great chili-chocolate or a true hot honey. The target handed to such a design is not "maximize heat" but "set heat and sweetness equal at the M2 ceiling." Below, we render that target; the six unengaged axes are held at a nominal baseline, well under the reference of 1,000 and the per-axis scorecard ceiling of 5,000, to keep the composition sharp. (This is a design target, a brief, not a measured fruit.)

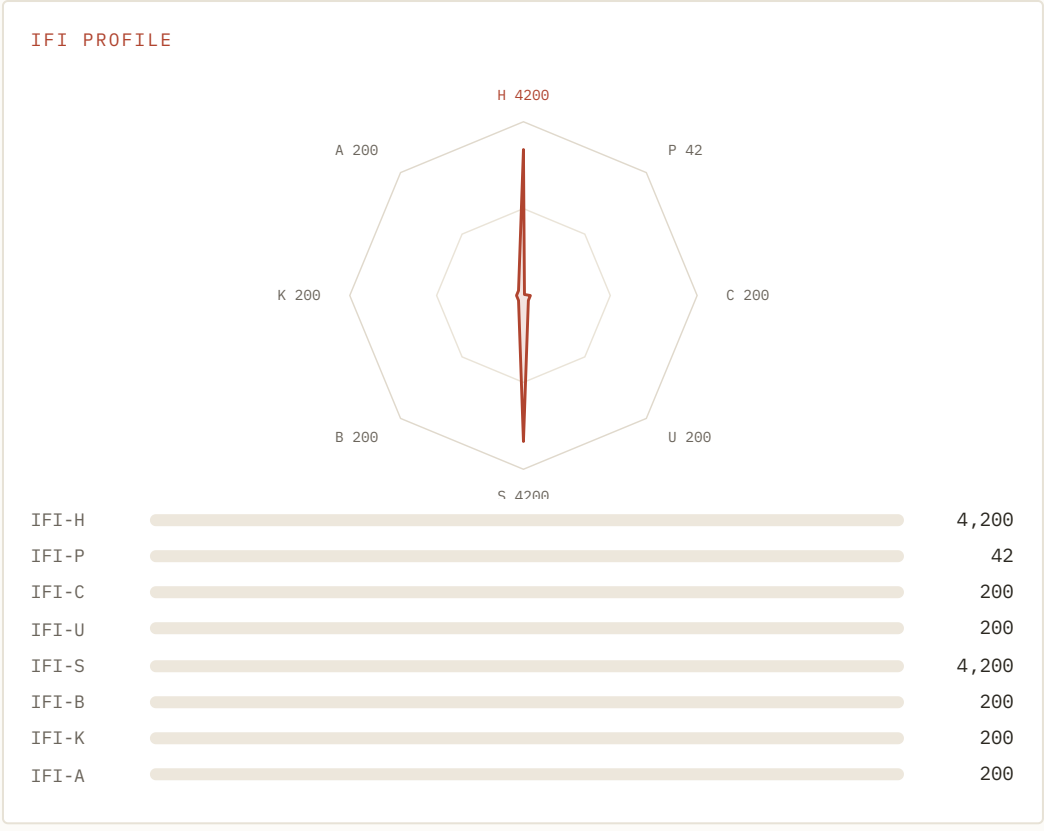

The two engaged axes sum to exactly the 8,400 ceiling, and each lands where Phoenix puts its single loudest note. That is the whole economics of M2 in one picture: skip six axes and you can afford to shout the other two.

The full octagon — paying the whole bill

To reach all eight at once is to pay the bill nature always refused, and it has been done exactly once. Phoenix — SS-0096, our model predicts approximately 2.95 million SHU, every axis above baseline

— is the one fruit that carries the complete octagon, and its profile is the clearest single object in the whole Index.

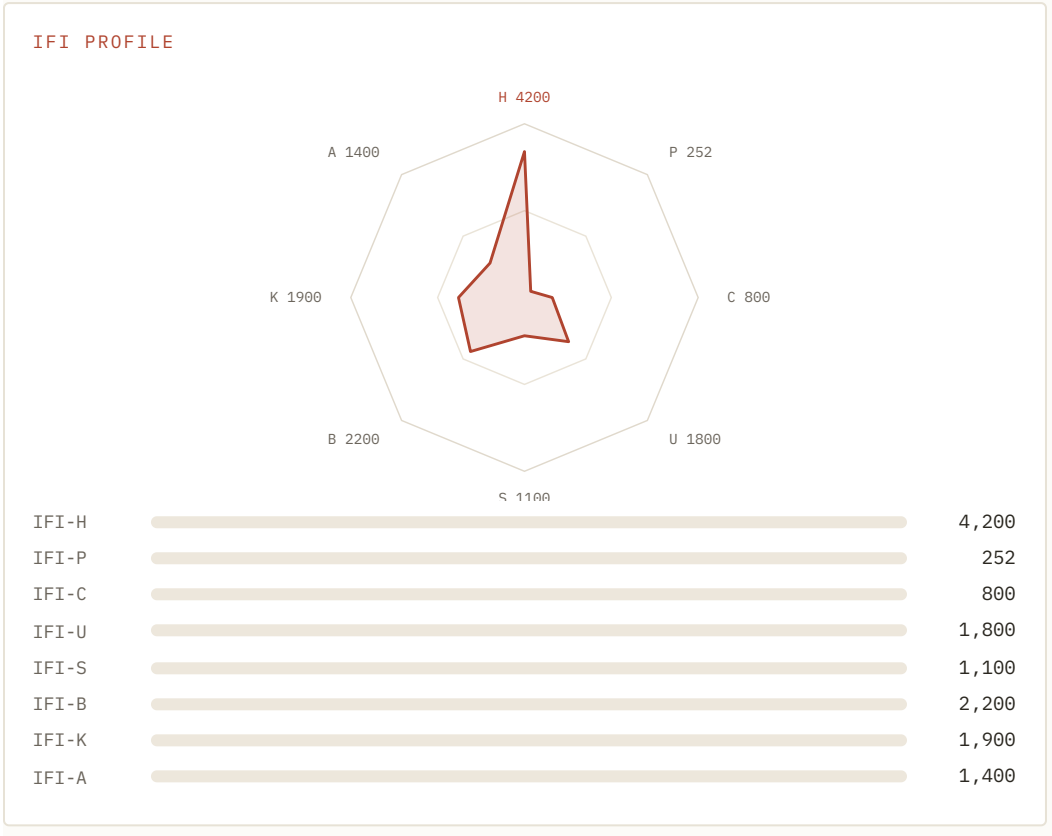

Sum those eight and the combined loudness is **13,652** — far above the M2 ceiling, and the reason it took years rather than months. Expressed as half-activation concentrations, the octagon spreads across a wide band of the receptors it plays:

AXIS	RECEPTOR	IFI	EC_{50} ÷ E_REF
H · Heat	TRPV1	4,200	0.238
B · Zero-Bitter	TAS2R38	2,200	0.455
K · Kokumi	CaSR	1,900	0.526
U · Umami	mGluR1	1,800	0.556
A · Astringency	salivary PRPs	1,400	0.714
S · Sweetness	TAS1R2·R3	1,100	0.909
C · Cool-Burn	TRPM8	800	1.250
P · Pungency	TRPA1	252	3.968
Total		**13,652**	

Two readings reward the eye. Heat, at 4,200, our model predicts half-activation of TRPV1 at 0.24 of reference — four axes' worth of loudness poured into one, sitting at 84% of the per-axis scorecard ceiling. Cool-burn and pungency are the two axes running below the reference line: cool-burn, at 800, is designed to need 1.25× the reference concentration to half-activate TRPM8, and pungency, at 252, sits at nearly four times the reference concentration — both quiet, both remarkable for a pepper to register on those axes at all. The octagon is not eight axes shouted equally; it is eight axes held, each at the loudness the shared budget could spare it once the others were paid for.



FIG. 17 — Phoenix in the controlled chamber: the one fruit engineered to register on all eight axes at once, and to pay for all eight at once.

It is worth saying that nature faces this same budget and mostly, wisely, declines to pay it. A wild plant invests in the one or two flavor-and-defense axes that serve its own survival — heat to turn away a browsing mammal, bitterness to warn of a toxin, a burst of sweetness to recruit a disperser — and spends almost nothing on the rest, which is why so few fruits anywhere are loud on more than a couple of axes. The near-even octagon is simply not a shape the wild has any reason to grow. It is a human intention, a flavor composed for the eater's pleasure rather than the plant's defense, and when we build a full-spectrum fruit we are not improving on nature's design so much as pursuing a goal nature never held — and paying, in metabolic budget, the bill nature always refused.

Error, reproducibility, and the honest limits

A design tool is only as trustworthy as its error bars, so trace them through the formula. IFI = 1,000 × E_ref ÷ E_cand is a ratio, and for a ratio the *relative* variances add: $CV(IFI)^2 = CV(E_{ref})^2 + CV(E_{cand})^2$. The reference is measured once, deeply, and held fixed, so its contribution is small; the reproducibility of a score is set by the construct's own batch-to-batch variation. Our designs hold that at **≤7%**, against **42%** in wild-type — and that single number is the difference between a flavor you can formulate against and one you cannot. A score of 4,200 carries a batch band of about ±294 at 7%; the same axis in a wild plant would swing ±1,764 at 42%, a quarter of the way to meaningless. Consistency is not a cosmetic virtue here. It is what makes a number worth printing.

The combinatorics change the error in two instructive ways. When you *sum* loudness — combined loudness across engaged axes — independent batch errors partly cancel: for the even M2 pairing of 4,200 + 4,200, each ±294, the sum's error is $\sqrt{(294^2 + 294^2)} \approx 416$, or about ±5% on 8,400 — *tighter,* in relative terms, than either axis alone. But when you hold a *ratio,* which is the real target of every M2 and M3 design, the errors compound instead: CV of a ratio is $\sqrt{(7^2 + 7^2)} \approx 10\%$. A chord specified by proportion is intrinsically noisier than a single note specified by loudness, which is the quantitative reason multi-axis design is the harder craft. The budget makes the chord expensive to build; the error propagation makes it expensive to *hold.*

And the honest limits, named plainly, because a reader deserves them. The flux model assumes metabolic steady state and known capacities; it prices competition well and regulation imperfectly, so a target near the feasible envelope's edge is a prediction, not a promise, and we treat it as one. An over-budget design announces itself as a sickly plant and a wandering CV long before the tasting, which is mercy, not failure — the model's rejection and the plant's distress are the same signal arriving twice. The one interference the receptor-level score genuinely cannot see is **perceptual cross-talk:** a loud bitterness can mask a sweetness that TAS1R2·R3 reported with perfect honesty, and kokumi can lend gain to umami so the broth eats louder than its curve predicts. The Index keeps the axes clean by scoring each at its own deaf receptor — but the eater tastes the sum, not the spreadsheet, and the designer has to carry both pictures at once: the eight honest numbers, and the single chord they will become on the tongue.

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

That doubled vision is finally what makes the Index a design tool and not merely a measuring stick. It tells you not only how loud each axis *is,* but how loud it is *allowed* to be, given everything else you have asked one living plant to do at once. To compose a flavor is to spend a finite budget wisely across eight competing demands — and to do it with enough care that the fruit which carries the chord is a healthy, nourishing thing, grown to hold the most it can and give the most it can. That is the arithmetic, and it is also, for our family, the point.



PART IV OF V

From Number to Plant

The design loop · the phenylpropanoid spine · forty-two hosts · the roadmap of sixty-four

Chapter 15 · From Score to Sequence — *designing the molecule to hit the number.*

We begin, always, with a number. Say it is heat held at the furthest honest edge — an IFI-H of 4,200, a little under three million Scoville, and not heat standing alone but heat set inside a whole flavor, the other seven axes sounding beneath it like strings under a melody. The Index gives us that number. The receptor biology tells us which molecule must carry it and at which channel it must be heard. A genome-scale model of the plant's metabolism tells us the pathway that builds the molecule and where that pathway chokes. What remains is the hardest and most beautiful step of all the work: turning the number into a *sequence* — editing a living plant's own genome so that the plant, left to grow in soil and light, becomes on its own the flavor we specified. This is the design loop. It is the place where one pair of hands at a keyboard, held inside the longer patient work of our family, did in a season of computation what the breeders of ten thousand years could not.

Every breeding program before ours worked from the outside in. You chose a phenotype — the hottest fruit on the plant — saved its seed, and *hoped the genotype followed,* crossing and backcrossing across years, reading a plant's visible traits and guessing at the genes beneath them. It worked, slowly, and it left the heat of a pepper hostage to chance, wildly variable from fruit to fruit, because the genetics underneath were never actually controlled, only leaned toward. We inverted the arrow. We work from the inside out — from the gene to the molecule to the number — and we never guess at the genotype, because we are the ones who write it. The whole of this chapter is that inversion, slowed down and shown in full: how a score on a page becomes an edited base pair in a chromosome, and why every turn of the method is an act of restraint as much as ambition.

Reading the number backwards into a molecule

The Index is built on one equation, and the design loop runs that equation in reverse. A flavor's score is the half-maximal activation of its receptor, inverted against a fixed reference and scaled by a thousand:

$$\mathrm{IFI} = 1000 \times \mathrm{EC}_{50}(\text{reference}) \div \mathrm{EC}_{50}(\text{candidate})$$

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the honest measure of *loudness,* because it asks not how much of a compound a food holds but how little of it the tongue needs to hear it. A smaller EC_{50} means a louder flavor and a higher score. For heat the anchor is set so that IFI-H of 1,000 is the

half-activation of the native wild-type reference; every number above it is a flavor that speaks more loudly than the plant we started from.

To design *to* a number, we solve the equation for the one unknown that matters — the EC_{50} the finished fruit must present:

$$EC_{50}(\text{target}) = 1000 \times EC_{50}(\text{reference}) \div IFI(\text{target})$$

Put the target in. For IFI-H 4,200, EC_{50}(target) = 1000 ÷ 4,200 × EC_{50}(reference) = 0.238 × EC_{50}(reference). The designed fruit's extract must half-activate TRPV1 at a shade under a quarter of the reference concentration. Said the other way, the ratio EC_{50}(reference) ÷ EC_{50}(target) is exactly 4.2 — the fruit must be 4.2 times louder than the plant we began with. The brief "IFI-H 4,200," which arrives as a mark on a scorecard, is in truth a demand for a precise molecular potency, and because capsaicin is capsaicin — one fixed molecule at one fixed channel — a 4.2-fold rise in loudness is a 4.2-fold rise in how much capsaicin the fruit actually presents to the receptor. The number has become a flux target. That is the first translation, and it is pure arithmetic.

It helps to see what 4.2-fold means on the shape of a real dose-response, because receptors do not answer in straight lines. TRPV1 opens cooperatively: two capsaicin molecules push on the channel together, so we model its activation with a Hill coefficient near n = 2, and the fraction of channels open follows

$$\theta(c) = c^n \div (EC_{50}{}^n + c^n)$$

At the one concentration where the wild-type reference sits exactly at its half-max — θ = 0.5 by definition — what is the designed fruit doing? Its EC_{50} is 4.2 times smaller, so with n = 2:

$$\theta = 1 \div (1 \div 4.2^2 + 1) = 1 \div (1 \div 17.64 + 1) = 1 \div 1.057 = 0.946$$

At the very dose that barely half-opens the reference channel, the designed fruit has already driven TRPV1 to 94.6 percent of saturation. That single figure is the whole ambition of the chapter made visible: a modest-sounding 4.2× in EC_{50} is, on the steep middle of a cooperative curve, the difference between a warmth and a roar. And to bring the designed extract back down to a polite half-max, you would have to dilute it 4.2 times further than you dilute the reference — the loudness is real, calculable, and would be reversible on the bench.

The loop, in four turns

Nothing in what follows was invented for peppers. The loop has four turns, and each is an ordinary instrument of modern computational biology pointed, for once, at flavor instead of at disease. Model the enzyme. Score every edit it could take. Confirm the bottleneck. Write the winner in. Those four verbs are the whole method, and their quiet universality is the reason a single chili could become a platform rather than a product.

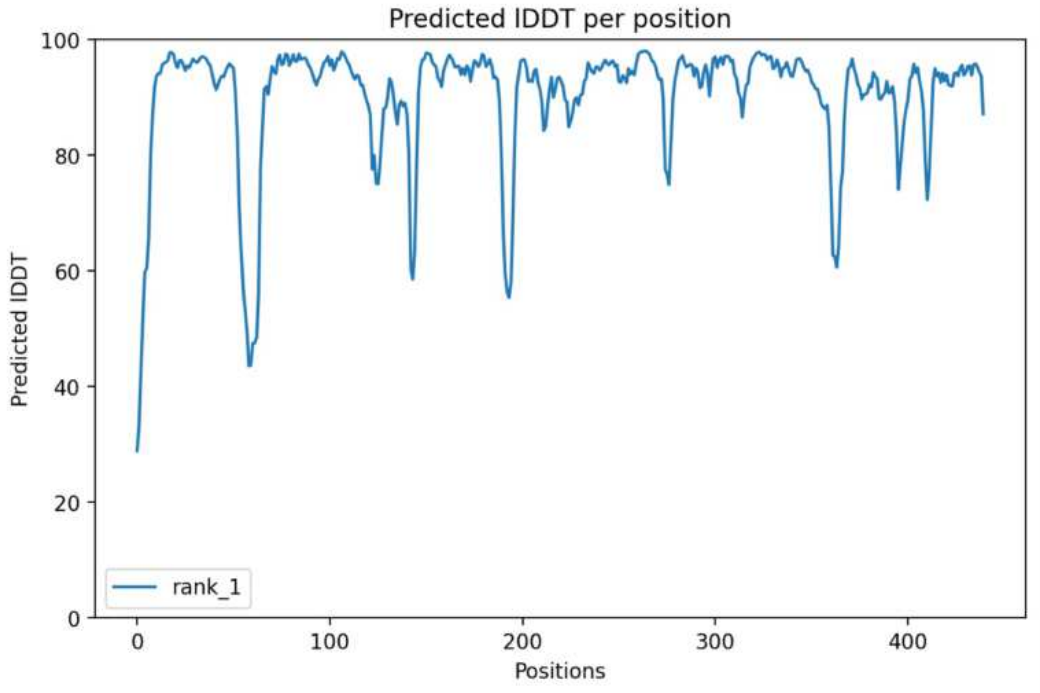

FIG. 16 — The capsaicin-synthase (Pun1/AT3) model, coloured by pLDDT confidence. The binding domain and active-site cleft resolve high-confidence; the design reads them directly.

For heat, the enzyme at the end of the line is capsaicin synthase, the Pun1/AT3 acyltransferase that condenses vanillylamine with an eight-methyl branched-chain acyl group to make capsaicin. The length and saturation of that acyl chain set the balance of capsaicin against dihydrocapsaicin, and with it the very *shape* of the burn — its rise, its peak, its slow honest fade. If we can read that enzyme's structure, score its every possible mutation, understand how it sits within the flux hierarchy our platform model reveals, and rewrite the letters that govern its product character and reliability, then we can set the molecule, and through the molecule the EC_{50}, and through the EC_{50} the number. Each turn earns the next; none can be skipped without lying to ourselves.

Turn one · modelling the enzyme, and learning where it is safe to cut

We predict the enzyme's three-dimensional fold and — this is the part that matters more than the fold itself — we read the *confidence* of that prediction, residue by residue, so that we act only where the model is sure. The first confidence map is the **pLDDT** score, a per-residue estimate from 0 to 100 of how much to trust each position. In Pun1, reassuringly, the two regions we need are the two that resolve at high confidence: the substrate-binding domain and the active-site cleft, where the chemistry actually happens. The floppy connecting loops the model cannot pin down come back low, and we leave them exactly alone. A prediction we do not trust is not a foundation we will build on — the same clean-living instinct that keeps us from putting into food what we cannot vouch for keeps us from cutting a gene where the model only shrugs.

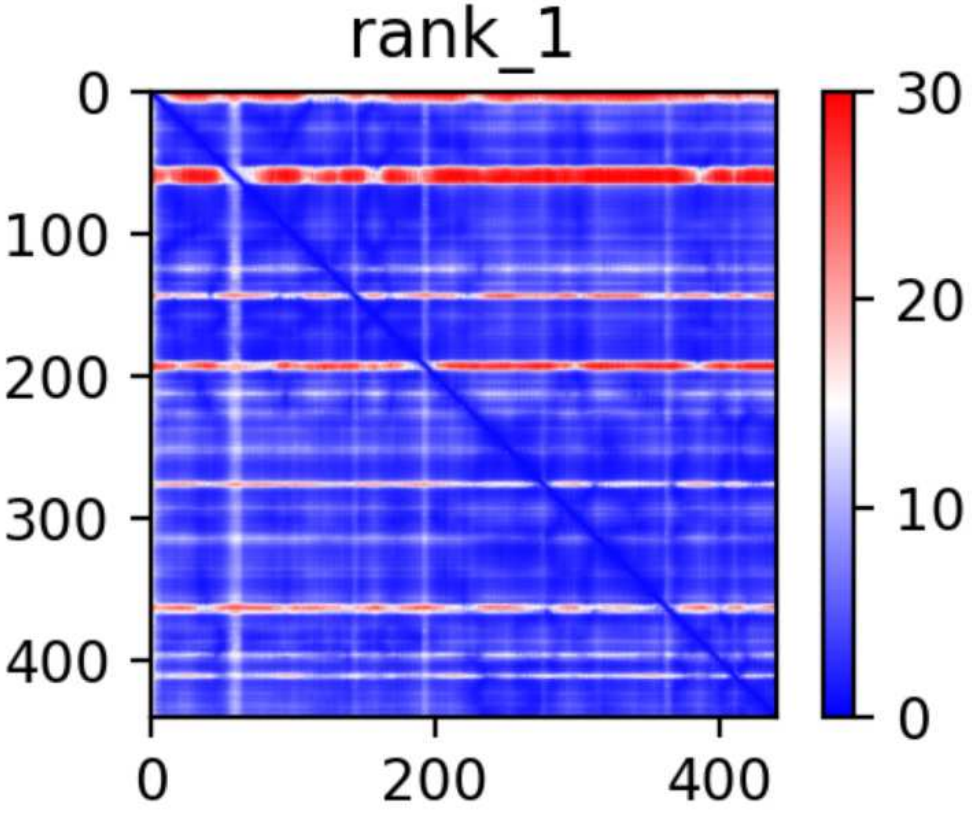

FIG. 17 — The predicted aligned error map for the same model: a residue-by-residue picture of which parts of the fold are pinned relative to each other. Dark blocks move as rigid units; the pale off-diagonal flags the loops that float.

The second map is the **predicted aligned error**, and it answers a different question — not "is this residue placed correctly?" but "is this residue placed correctly *relative to that one?*" Read as a heatmap, it shows which stretches of the enzyme move together as rigid bodies and which flex on their own. This matters enormously, because an edit meant to open the catalytic pocket has to be aimed at a residue that genuinely controls the pocket's geometry, not at one that merely sits near it on the page and drifts independently of it in space. The aligned-error map is how we tell the load-bearing wall from the curtain. Only once both maps agree — a residue that is confidently placed *and* rigidly coupled to the machinery we mean to change — do we let a position onto the shortlist. Everything downstream is as trustworthy as this first, most conservative act of reading.

Turn two · scoring every edit the enzyme could take

Now the search. We take a protein language model — **ESM2, in its 650-million-parameter form** — and ask it to evaluate every single-residue substitution the enzyme could possibly take. The idea sounds like alchemy until you see what the model is. A protein language model is trained on the whole observed universe of protein sequences, hundreds of millions of them, with residues masked out and the model asked to guess what belongs; in learning to fill those blanks it absorbs the deep statistics of what proteins actually are — which residues co-vary down the generations, which positions are sacred and which are free, what the grammar of a working fold sounds like. Ask such a model to read a proposed substitution and it returns, in effect, how *surprising* the change is. Formally it returns a log-likelihood ratio, the log-probability it assigns the mutant residue at that position minus the log-probability of the wild-type residue:

```
LLR = log P(mutant | context) − log P(wild-type | context)
```

A positive LLR means the model finds the change fluent, the kind of substitution real evolution has tolerated a thousand times in proteins like this one; a negative LLR means it reads as a stumble. That surprise predicts function remarkably well with no experiment and no label at all — the field calls it zero-shot variant-effect prediction — and the 650M model gives us the score for every possible single substitution in Pun1 in effectively one sweep, a complete fitness landscape read straight off the model's own fluency.

The scale of that sweep is the reason we need the model in the first place. An acyltransferase like Pun1 runs to a few hundred residues — call it L ≈ 430, the size of a typical BAHD-family enzyme. The number of single-letter substitutions is 19 × L, on the order of 8,000 — already more than any lab could make and assay by hand. But the two shipped edits were never found one at a time; they are a *combination,* and the combinatorics of combinations are merciless:

- single mutants: 19 × 430 ≈ **8,000**
- double mutants: C(430, 2) × 19^2 ≈ 92,000 × 361 ≈ **33 million**

The winning pair lives somewhere in a space of roughly thirty-three million doubles, no two of which any breeder had ever combined. You do not find a point in a thirty-three-million-member space by growing plants and tasting fruit. You find it by scoring the landscape in silico, ranking, and testing only the handful the model stakes its fluency on.

And we rank on two things, never one. The language model's fluency tells us whether a mutation makes the chemistry *better;* a structure-based calculation of folding free energy tells us whether the protein still *folds.* We multiply them:

```
score(variant) = catalytic efficiency × fold stability
```

The product, not the sum, is deliberate, and it is the quiet ethic of the whole method. A sum would let a spectacular catalytic gain paper over a folding catastrophe; a product refuses. Picture two candidates in illustrative relative units. One is fast but fragile — catalytic efficiency 3.0, fold factor 0.2 — and scores 3.0 × 0.2 = 0.6. The other is balanced — efficiency 1.8, fold 1.1 — and scores 1.8 × 1.1 = 1.98. The balanced edit wins by more than threefold despite being the slower chemist, because a faster enzyme that sometimes unfolds buys loudness at the cost of consistency, and consistency is the entire point of a standard. We would rather have a pepper that is reliably hot than one that is occasionally spectacular.

The folding half of the score deserves its own arithmetic, because it is where the "small" edit earns its keep. The structure calculation returns ΔΔG, the predicted change in folding free energy in kilocalories per mole, where negative means *more* stable. Folding is an equilibrium between the correctly folded chain and the unfolded or misfolded one, and a change in free energy shifts that equilibrium by a factor you can read straight off the Boltzmann relation. At room

temperature, 298 K, the thermal energy RT is 0.592 kcal/mol, so a stabilization of ΔΔG shifts the folded-to-unfolded ratio by exp(−ΔΔG ÷ RT). Work it for each of the edits considered:

EDIT	SITE & EFFECT	ΔΔG (KCAL/MOL)	FOLD-EQUILIBRIUM SHIFT
S39L	N-terminal; binding affinity ↑	−2.1	exp(2.1 ÷ 0.592) ≈ **35×**
L345G	active-site cleft; steric relief	−1.8	exp(1.8 ÷ 0.592) ≈ **21×**
C175S	disulfide-misfold prevention (structural annotation; not in shipped construct)	−0.9	exp(0.9 ÷ 0.592) ≈ **4.6×**
Shipped double variant	S39L+L345G combined	**−3.9**	exp(3.9 ÷ 0.592) ≈ **700×**

Taken together — and to first order, for mutations this far apart in space, the free energies simply add — the two shipped edits tip the folding equilibrium roughly 700-fold toward the correctly folded protein. Look what that does to misfolding, which scales as the inverse of the equilibrium constant: were the reference enzyme to misfold in one chain in twenty, a 700× shift would drop that to about one chain in fourteen thousand. The protein that now almost never folds wrong is the reason the fruit is now almost always the same heat. That is a thread we will pull all the way through.

Turn three · confirming the bottleneck

Before we touch a single base pair, a genome-scale flux-balance model of the whole plant has to answer one unglamorous question: is Pun1 actually the step that limits output, or is the true ceiling somewhere upstream? Flux-balance analysis treats the plant's metabolism as a vast system of mass-balance constraints — every metabolite produced must be consumed — and solves for the distribution of flux that the network can actually sustain. It is Liebig's barrel in algebra: a barrel holds water only to its shortest stave, and a faster capsaicin synthase raises yield only if capsaicin synthase is the shortest stave. The model reports a shadow price for each reaction — how much the whole system's output would change if that one step were loosened — and the answer came back upstream.

Control coefficients, which by the Metabolic Control Analysis summation theorem must add to exactly 1.000, placed Pun1's share at effectively zero across 5,000 Monte Carlo draws. Our model predicts control sits in the shared phenylpropanoid branches at 4CL and HCT/CCoAOMT — steps that feed not just capsaicinoids but the whole of plant secondary metabolism. Pun1 is the final capsaicin synthase; it labels the product, it does not govern the rate. A faster Pun1 that arrives at an upstream-constrained junction simply waits.

What the Pun1 edits do buy, decisively, is *product character* — the balance of capsaicin against dihydrocapsaicin that determines the burn's shape, rise, and fade — and *reliability*, the fold stability that holds that balance constant from fruit to fruit and batch to batch. Total capsaicinoid ceiling is a question for the upstream pathway;

what Pun1 edits govern is the character and consistency of what arrives at the receptor. The design logic shifts with this finding, but it does not collapse: we edit only where the analysis confirms the edit will be felt, and the Pun1 edits earn their keep not through raw flux but through product fidelity. This is the turn that keeps the method honest. It is the difference between designing a plant and merely hoping at one.

Turn four · writing it in

Only now, with the target residues chosen under two confidence maps, scored on two objectives, and cleared by the flux model, do we write them into the genome — and the writing is as restrained as the reading was. A point substitution is installed by CRISPR, increasingly by the gentler base and prime editing that change a single genetic letter without cutting both strands of the DNA at all. We are not inserting foreign machinery or bolting a borrowed gene onto the plant; we are correcting two letters of the plant's own text, the way you would fix two typos in a manuscript rather than stapling in someone else's pages. Where the edit allows, we choose a codon the plant already uses heavily, so the ribosome reads the new instruction as fluently as the old. What grows from that seed carries no scaffold, no marker, no passenger DNA — only its own genome, two characters truer to the flavor we specified.



FIG. 18 — The sequencer bench where every edited line is read back, base by base, and checked against the design before a single seed goes to soil.

Then we read it back. Every edited line is sequenced and checked against the design, letter for letter, before it is ever planted, because a method whose whole claim is precision has no right to assume its own hand was steady. The language model proposes; the structure disposes; the flux model permits; the sequencer verifies; and the plant, in the end, decides.

The shipped edits, and why we score two things and never one

Out of all the substitutions the gene could have taken, two emerged as the shipped double variant that no breeder had ever found and no pair had ever been combined.

EDIT	EFFECT	ΔΔG (KCAL/MOL)
S39L	binding affinity ↑ (N-terminal)	−2.1
L345G	steric relief, active-site cleft	−1.8
C175S	disulfide-misfold prevention (structural annotation; not in shipped construct)	−0.9

Read them one at a time. **S39L** replaces a serine with a leucine near the N-terminus, tightening how the enzyme grips its substrate and, as the free-energy column shows, stabilizing the fold by 2.1 kcal/mol — the single largest structural gain of the two. **L345G** swaps a bulky leucine for the smallest residue there is, glycine, relieving a steric clash that had been crowding the active-site cleft; the pocket opens, substrate turns over faster, and the fold gains a further 1.8. **C175S** — a potential change from cysteine to serine to remove a thiol prone to forming the wrong disulfide bond and dragging the whole chain into a misfold — appeared as a structural annotation from our modelling (ΔΔG −0.9). It was held out of the final shipped design; it lives in our modelling library, clearly labelled as a structural annotation and not a synthesis target. The table retains it to show the full scope of the structural analysis.

The two shipped edits together are the purest argument for why we multiply two objectives instead of chasing one. A method that optimized catalytic efficiency alone would hunt only for faster turnover; a method that optimized fold stability alone would hunt only for buried hydrophobics. Neither discipline, pursued alone, finds S39L+L345G — a pair that scores well on *both.* Flavor built as a *standard* is not about peak loudness; it is about saying the same thing every time, and the thing that lets a protein say the same thing every time is folding the same way every time. Score only speed and you design a pepper that is sometimes magnificent. Score speed and stability together, as their product, and you design one that is magnificent on schedule. Our model predicts approximately 2.95 million SHU for Phoenix (SS-0096) — but the deeper achievement is not in that predicted peak at all. It is in the fold stability that holds it, and it does not show up on any heat scale.

The same hot, every time

Here is what the loop actually buys, and it is invisible to anyone looking only at how hot the pepper gets. A wild-type pepper's heat wanders from fruit to fruit with a coefficient of variation — the standard deviation divided by the mean — of about **42 percent.** The designed cultivar holds to **7 percent or better** as a design specification. That is not a small tightening; it is a collapse. Variance is the square of the standard deviation, so the ratio of variances is $(0.42 \div 0.07)^2 = 6^2 = \mathbf{36}$**.** The designed pepper's heat is thirty-six times less *scattered* than the plant we began with, its standard deviation six times smaller. To put scale on it in predicted Scoville units, taking our model's Phoenix prediction of approximately 2.95 million SHU:

	IF IT CARRIED WILD-TYPE SCATTER	PHOENIX AS DESIGNED (SS-0096)
Coefficient of variation	42%	≤ 7%
Standard deviation σ	1,233,938 SHU	205,656 SHU
95% fruit-to-fruit band	≈ 0.52M – 5.36M SHU	≈ 2.53M – 3.34M SHU
Fruit needed to fix the batch mean to ±1% (95%)	≈ 6,800	≈ 190

Sit with the middle row. A pepper with a predicted mean near three million Scoville but a 42 percent coefficient of variation would deliver, across its fruit, a ninety-five percent spread running from barely half a million Scoville to well over five million — a chili that is a mild surprise on Tuesday and a medical event on Thursday, the same plant, the same vine. At 7 percent the same mean yields a band from roughly 2.5 to 3.3 million: still formidable, but *knowable,* a number a formulator can build a recipe on and a cook can trust. The bottom row is the same truth read as labor. To pin a batch's true mean heat to within one percent with ninety-five percent confidence, you sample until 1.96 × CV ÷ √n falls under 0.01, which takes about 6,800 fruit at wild-type scatter and about 190 at the designed cultivar's — the same thirty-six-fold ratio, now measured in the hours of a quality bench. Reproducibility is not a luxury laid on top of the flavor. It *is* the flavor, in the only sense that lets a taste become a standard.

IFI PROFILE

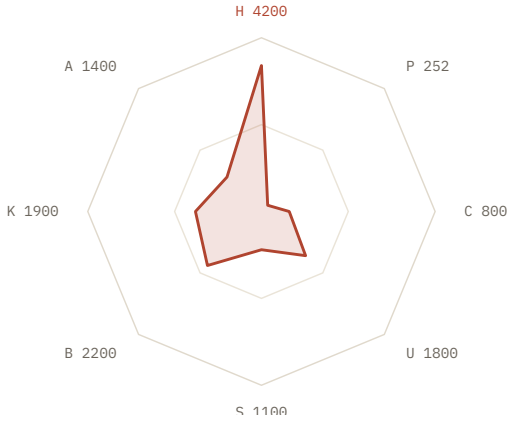

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

That profile is Phoenix, SS-0096 — the first cultivar we designed to register on all eight axes at once, heat at 4,200 with pungency, cooling, umami, sweetness, zero-bitter, kokumi and astringency all sounding in designed proportion beneath it. It did not arrive at the top of a staircase by luck; it arrived by the loop, run again and again up a progression of flagships, each one reaching further in computational design scope and axis activation than the last. The table below gives our model's predicted figures for each cultivar in the series:

CULTIVAR	SKU	PREDICTED SHU	IFI-H	SHU PER IFI PT	AXES
Ember	SS-0087	866K	1,570	552	2/8
Cinder	SS-0088	1.05M	2,290	459	3/8
Kindle	SS-0089	1.25M	2,000	625	3/8
Flare	SS-0090	1.48M	—	—	4/8
Blaze	SS-0091	1.72M	2,630	654	4/8
Wildfire	SS-0092	1.96M	2,930	669	5/8
Scorch	SS-0093	2.21M	3,220	686	5/8
Firestorm	SS-0094	2.47M	3,490	708	6/8
Solara	SS-0095	2.72M	3,980	683	6/8
Phoenix	SS-0096	2.95M	4,200	700	8/8

The last column is quietly instructive, and it is the kind of honesty the Index was built to make possible. If Scoville heat units and IFI-H measured the same thing, the ratio between them would be a constant; it is not. It climbs from the mid-500s at the foot of the progression to a plateau near 700 at the top, with real scatter along the way — Cinder low at 459, Kindle high at 625. That drift is not sloppiness; it is physics. Scoville is a perceptual scale, read by human panels diluting an extract until the burn disappears, and perception compresses — the tongue, like the ear, answers to ratios more than to sums. IFI-H is a receptor scale, read as an EC_{50} at the channel, and it runs closer to linear in the molecule. The two are kin, and they rise together, but they are not the same instrument, and the honest column that shows them diverging is worth more than a tidy one that hid it.

The loop is the product

Nothing in those four turns is specific to heat, or to peppers, and that is the whole of why a chili became a company. Model the enzyme, score every edit, confirm the bottleneck, write the winner: change the target enzyme and the objective, and the method holds unchanged. Point it at the glutamate machinery of a tomato and it designs umami; at the glycosyltransferases of a stevia and it designs sweetness without sugar; at the tannin-polymerizing enzymes of a grape and it designs the astringent grip of a great wine. The upstream phenylpropanoid branch — PAL to C4H to 4CL to HCT to COMT — runs in every vascular plant on earth; only a few terminal enzymes are pepper's own. Which is why a technique first applied to a single acyltransferase in a single chili now runs across forty-two living host organisms and eight sensory axes, with sixty-four more concept organisms mapped on the roadmap ahead, without being reinvented for each. The loop does not care which flavor it is chasing. It asks only that the flavor reduce, in the end, to an enzyme whose sequence can be read, scored, and rewritten — and nearly every flavor does.

And the loop closes on itself, which is the part that compounds. Every edited plant we grow and observe feeds its true, in-planta result back into the next round of scoring — the model learns where its fluency was right and where the living plant corrects it — so the method does not merely produce cultivars, it gets *better* at producing them with every cultivar it produces. Heat was never the destination. Heat was the proof. The loop was the product.

Where the loop strains — the honest limits

A method is only as trustworthy as its account of its own failures, so here are ours, plainly. The folding arithmetic assumed the two edits' free energies simply add; real mutations are not perfectly additive, and when two edits sit close enough to feel each other the assumption of independence breaks down into epistasis — the combined effect is not the sum, and occasionally it is the opposite of what either predicted alone. We treat additivity as a first estimate to be confirmed at the bench, never as a law. The language model's fluency, likewise, is a prior and not a promise: a mutation it finds plausible can still fail in the one plant we care about, because the model has read the universe of proteins but not the particular cellular world this enzyme folds into. That is the whole reason the plant gets the last vote — the loop ends at a living fruit and a sequencer, not at a confident number on a screen. Where the confidence maps come back low we do not edit at all, which means some genuinely useful edits in the floppy loops stay out of reach, a cost of our own caution we accept with open eyes. The flux model can be fooled when a cell quietly reroutes metabolism around an edit we thought decisive, relieving the bottleneck we aimed at and opening another. And the gap between the receptor scale and the perceptual one — the drifting ratio of Scoville to IFI-H — is a standing reminder that we design against the receptor but people eat with the whole apparatus of perception, so the number we write and the pleasure it becomes are near relations, not identical twins. None of these is a reason the method fails. Each is a reason we run all four turns and verify every one, and a reason the plant, not the model, is allowed the final word.

Number in, sequence out, and a living thing grown to the specification — grown, in the end, not only to taste exactly as we meant but to be food you can give the body with intention, the same every honest time. That is the whole method. Heat was only the first place we proved it. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 16 · The Phenylpropanoid Spine — *one pathway under almost everything.*

The capsaicin work was never really about peppers. We said so from the beginning, and this chapter is the biochemical reason we could say it with a straight face: beneath the eight axes and the forty-two organisms runs a single shared pathway, present in every vascular plant on Earth, and the pepper-specific machinery is only a short cap bolted onto its end. Learn the spine, and you have learned the part of the method that transfers — which is most of it, and which is why a program that began with a chili can credibly promise a volcanic tomato and a sweet-heat strawberry.

We are a vegetarian family by life and by conviction, and there is a particular, quiet joy in the fact at the center of this chapter: that the fire we design and the stiffening that holds a flower upright are two readings of one ancient sentence of metabolism. The plant wrote that sentence long before we arrived to read it. Our work was to read it carefully enough to add a word at the end.

The spine is the **phenylpropanoid pathway**, one of the oldest and most central routes in all of plant metabolism. It is the source of the lignin that lets a stem stand, of the flavonoids that color a petal and screen it from ultraviolet, of the polyphenols behind astringency, of the vanillin in vanilla and the eugenol in clove — and, at its far end in peppers, of the vanillylamine half of capsaicin. Every plant in the library already runs it, right now, because every plant needs its products to live. That is the whole leverage of the method, and the rest of this chapter is the mechanism, the mathematics, and the honest limits of it.

The spine, named — the universal five-enzyme relay

The pathway begins with the amino acid phenylalanine and proceeds through a conserved relay. We draw it as five steps, naming the enzyme, the cofactor it cannot work without, and where in the cell it happens — because the cofactors and the locations are what make the relay universal rather than merely common.

STEP	ENZYME	CONVERTS →	COFACTOR / CHEMISTRY	WHERE
1	PAL — phenylalanine ammonia-lyase	phenylalanine → *trans*-cinnamic acid	MIO electrophile (self-made from an Ala-Ser-Gly loop); non-oxidative deamination	cytosol
2	C4H — cinnamate-4-hydroxylase	cinnamic acid → *p*-coumaric acid	cytochrome P450 (CYP73A); O_2 + NADPH	ER membrane
3	4CL — 4-coumarate:CoA ligase	*p*-coumaric acid → *p*-coumaroyl-CoA	ATP → adenylate, then CoA thioester	cytosol
4	HCT — hydroxycinnamoyl transferase	*p*-coumaroyl-CoA → caffeoyl esters	BAHD acyl transfer onto shikimate/quinate	cytosol
5	COMT — caffeic acid *O*-methyltransferase	caffeic acid → ferulic acid → vanillin	*S*-adenosyl-methionine (SAM) methyl donor	cytosol

PAL opens the pathway, and it opens it with a flourish worth pausing on: the enzyme manufactures its own cofactor. Three of its own residues — alanine, serine, glycine in sequence — cyclize autocatalytically into an electrophilic ring called MIO, 4-methylideneimidazole-5-one, and that homemade electrophile is what lets PAL strip the amino group from phenylalanine without oxygen or metal, by *anti*-elimination, yielding *trans*-cinnamic acid and free ammonia. This is the committed first step, the gate through which carbon enters the entire phenylpropanoid world, and the plant guards it heavily: PAL is feedback-inhibited by its own downstream cinnamate and transcriptionally governed by the MYB family of regulators that decide, in effect, how much of the plant's carbon budget is spent on this chemistry at all.

C4H is a cytochrome-P450 monooxygenase anchored in the endoplasmic-reticulum membrane. Using molecular oxygen and electrons delivered from NADPH through cytochrome-P450 reductase, it installs a single hydroxyl at the *para* position to make *p*-coumaric acid. Because it sits in the membrane while PAL floats in the cytosol, the two often work as a loosely tethered pair — a metabolon — handing intermediate between them so the reactive cinnamate never wanders. **4CL** then activates *p*-coumarate as its coenzyme-A thioester in two motions: it spends one ATP to form a *p*-coumaroyl-adenylate, then swaps the adenylate for coenzyme A. That thioester, *p*-coumaroyl-CoA, is the single most important molecule in the chapter — the activated currency from which lignin, flavonoids, and the pepper's heat all diverge. **HCT**, a BAHD-family acyltransferase, shuttles the activated acid onto shikimate so a second hydroxyl can be added, then hands back caffeoyl species; **COMT** methylates with S-adenosyl-methionine to walk the intermediates onward through ferulic acid toward vanillin.

Five enzymes, four cofactor systems, one universal result. It is already running in every tomato and grape and coffee plant and blade of sorghum we have ever sequenced, because each of them needs lignin to stand and flavonoids to survive the sun. Nothing in this relay is a pepper's private property. That is not a convenience of our method; it is the premise of it.

Where the trunk forks — one precursor, three destinies

The spine is really a trunk with great boughs, and seeing where they branch is seeing why the pathway governs so much of what a plant tastes like. From *p*-coumaroyl-CoA, the activated product of the third step, the metabolism forks three ways.

One bough runs to the monolignols, polymerized into **lignin** — the stiffening that let plants leave the water and stand upright. A second runs through chalcone synthase into the **flavonoid** branch: the pigments, the ultraviolet screens, and, further down, the flavan-3-ols and condensed tannins that register on our astringency axis and much of our bitter axis. A third, in peppers alone, carries on toward vanillin and the aromatic half of **capsaicin**. One trunk, three destinies — the structure of a stem, the color of a petal, the fire of a fruit — all drawn from the same few activated intermediates.

This fork is not a diagram; it is a budget, and budgets are arithmetic. The carbon that enters at PAL is finite, and whatever flows down one bough cannot flow down another. Write the constraint plainly: if a unit of *p*-coumaroyl-CoA partitions into lignin, flavonoid, and (in peppers) capsaicinoid fractions, then

$$f(\text{lignin}) + f(\text{flavonoid}) + f(\text{capsaicinoid}) + f(\text{other}) = 1.$$

This is why heat and astringency are not independent dials but cousins drawing on one shared pool — the competition between axes is written into the biochemistry, not imposed on it from outside. Here is an illustrative partition (arbitrary conserved units, 100 µmol of committed phenylalanine per gram dry fruit) contrasting a heatless wild bell pepper, whose capsaicinoid branch is switched off, with a full-spectrum design that pulls carbon *into* the heat bough without starving the others to zero.

BRANCH FROM *P*-COUMAROYL-CoA	WILD BELL (PUN1 OFF)	FULL-SPECTRUM DESIGN	WHAT IT BECOMES
Lignin (monolignols)	55	48	structure — the plant stands
Flavonoid → tannin	40	20	pigment, UV screen, astringency / bitter
Capsaicinoid (vanillyl branch)	0	28	heat
Other soluble esters	5	4	aroma, defense
Total	**100**	**100**	conserved

Read the two columns and the whole thesis of competing axes becomes concrete. To light the heat bough from zero to 28, the design draws the flavonoid bough down from 40 to 20 — it does not abolish it, which is exactly why a full-spectrum fruit still registers astringency and bitterness, only quieter than a tannin-first grape would. The constraint is conservation; the art is choosing the partition. A number raised on one axis is quietly borrowed from a neighbor, and the ledger always balances.

The vanillin crossroads

It is worth stopping at vanillin, because it sits at a crossroads that holds the entire argument in one molecule.



FIG. 16·1 — Vanilla. One enzyme past the shared spine's end lies vanillin; one amination further lies the vanillylamine that becomes capsaicin's aromatic half.

Vanillin is the aroma of vanilla — a flavor in its own right, one of the library's own organisms, the backbone of custard and the most faked taste on Earth. It is also exactly one amination short of vanillylamine, the aromatic half of capsaicin. The same compound, a single enzyme apart, is the scent of a dessert and the backbone of a fire, and both are reached down the same spine every plant already runs. When we say the method generalizes, this adjacency is precisely what we mean: flavors that seem worlds apart on the tongue are often neighbors in the metabolism, separated by an enzyme or two at the end of a shared road. Vanilla registers on our kokumi axis — the rounding, mouth-filling continuity a CaSR agonist gives — and the pepper registers on heat; yet the two share a parent molecule. To learn the road is to be able to travel to any of its destinations, and the destinations do not resemble one another at all.

The four-enzyme cap — what makes a pepper a pepper

Onto that universal spine, peppers bolt just **four terminal enzymes** to finish capsaicin. These four are the only genuinely pepper-specific parts of the whole heat machine, and capsaicin's two halves are built in two different kitchens that meet only at the final step.

ENZYME	ROLE IN THE HEAT BRANCH	COFACTOR / CHEMISTRY	HONEST EDGE CASE
pAMT	aminates vanillin → vanillylamine (aromatic half)	PLP-dependent transamination	lose it, and the fruit makes non-pungent **capsinoids** instead
BCAT	transaminates valine / leucine → branched keto acids	PLP-dependent	feeds the acyl half; shared with primary metabolism
KAS	elongates the branched acyl chain	iterative Claisen with malonyl-ACP	chain length sets capsaicin : dihydrocapsaicin
Pun1 / CS	condenses the two halves → capsaicin	BAHD acyl transfer; amide-bond formation	natural loss-of-function → sweet bell pepper

The **aromatic half** descends the phenylpropanoid spine: vanillin, then **pAMT** — a pyridoxal-5′-phosphate (PLP) transaminase — turning it into vanillylamine. The **fatty-acyl half** comes from somewhere else entirely. **BCAT**, the branched-chain aminotransferase, takes the carbon skeletons of valine and leucine and hands them off as branched keto acids; the fatty-acid

machinery, **KAS** among its enzymes, elongates these by repeated condensation with malonyl-ACP into an 8-methyl-branched acyl chain. Finally **Pun1/CS** — the capsaicin synthase that gives the heat locus its name — forms the amide bond that joins amine to acyl and makes capsaicin. A pepper, in other words, assembles its fire from two metabolisms at once, an aromatic one and a lipid one, welded only at the last joint. Pun1/CS is the seam where the two halves finally join — the place that gives our designs their signature. Our flux model, however, places metabolic control further upstream: at 4CL and HCT, where the spine decides how much carbon enters the heat branch at all. Pun1 runs last in the queue; it labels the product and does not set the rate.

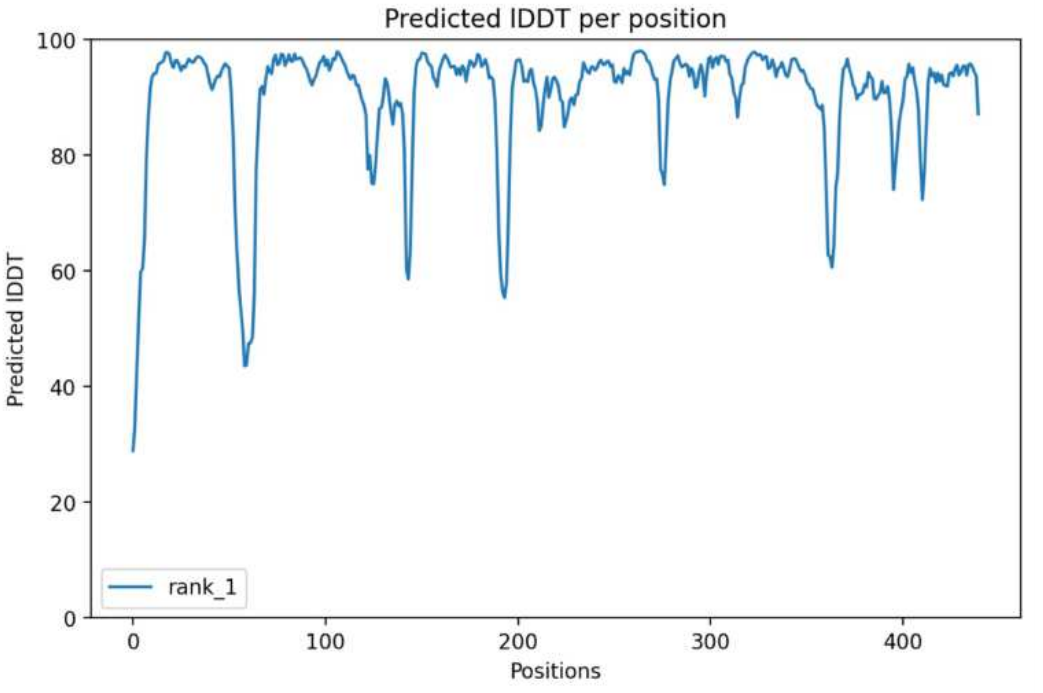

FIG. 16·2 — Capsaicin synthase (Pun1/AT3), colored by model confidence (pLDDT). The binding domain and the active-site cleft resolve at high confidence; the surface loops do not.

The chain chemistry is not a detail — it shapes the sensation. Capsaicin carries one double bond in its acyl tail (from 8-methyl-6-nonenoic acid); dihydrocapsaicin is fully saturated. The two reach TRPV1 at slightly different rates, so the *ratio* KAS sets between them governs the *shape* of the burn: a front-loaded flash that clears, or a slower warmth that lingers and climbs. Two fruits at the same Scoville number can burn like different events because their acyl chemistry differs.

And the enzymes fail in instructive ways. Break **pAMT**, and the plant cannot make vanillylamine; it esterifies vanillyl alcohol to the acyl chain instead and produces **capsinoids** — capsiate and its relatives, nearly heatless, the quiet chemistry behind the naturally sweet CH-19 pepper. Break **Pun1** itself — a natural recessive deletion at that locus — and you get the sweet bell pepper on every grocery shelf: the full spine runs, the fruit swells, and no capsaicin is ever condensed. Nature already knew how to switch heat off at this locus, twice, by two different lesions. Our design simply learned to switch it *up*, and to hold it there batch after batch.

The shipped construct carries two AI-proposed substitutions in Pun1 — never before combined — that lift and stabilize output together. A third substitution, C175S, is retained in our structural model as a disulfide-misfold annotation, but was not included in the final synthesized design:

EDIT	EFFECT	ΔΔG (KCAL/MOL)	IN SHIPPED CONSTRUCT
S39L	binding affinity ↑ at the N-terminus	−2.1	yes
L345G	steric relief in the active-site cleft	−1.8	yes
C175S	prevents a disulfide misfold (structural annotation — not synthesized)	−0.9	no

The two substitutions are additive by design — a tighter grip on substrate and a roomier cleft to turn it over — and together they are the difference between a fire that flickers and one that holds.

Reading the fire as a number — the index on the spine's output

A pathway is only half the instrument. The other half is the ruler we read its output with, and the ruler is the Iyengar Flavor Index. Every axis is a single number anchored to the half-maximal activation of a receptor — not how much of a compound a food holds, but how *loudly* that compound speaks to the nerve. The anchor is an EC_{50}, the concentration of a standardized extract at which the receptor reaches half of its full response, and the score is that EC_{50} inverted against a fixed reference and scaled by a thousand:

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

Each reference EC_{50} is a **stated convention** — fixed by definition on that axis, like Scoville's original dilution-to-threshold definition — not a value drawn from published assay data. The H axis (heat/capsaicin/TRPV1) and P axis (pungency/AITC/TRPA1) each have their own defined reference constant; a score of 1,000 on H and a score of 1,000 on P are both defined reference points, but they anchor to different conventions and are not on a single common scale.

Three things follow from the formula, and they are worth deriving rather than asserting. First, **why invert.** A loud stimulus reaches half-activation at a *low* concentration — a little goes far — so loudness runs opposite to EC_{50}. Dividing by EC_{50}(candidate) turns "needs little" into "scores high," which is the direction a taster's intuition already runs. Second, **the reference scores exactly 1,000, by construction.** Set the candidate equal to the reference and the ratio is one; times a

thousand is a thousand. The zero-point of every axis is therefore the same gesture — the native wild-type reference, half-activating its receptor — and it looks like this:

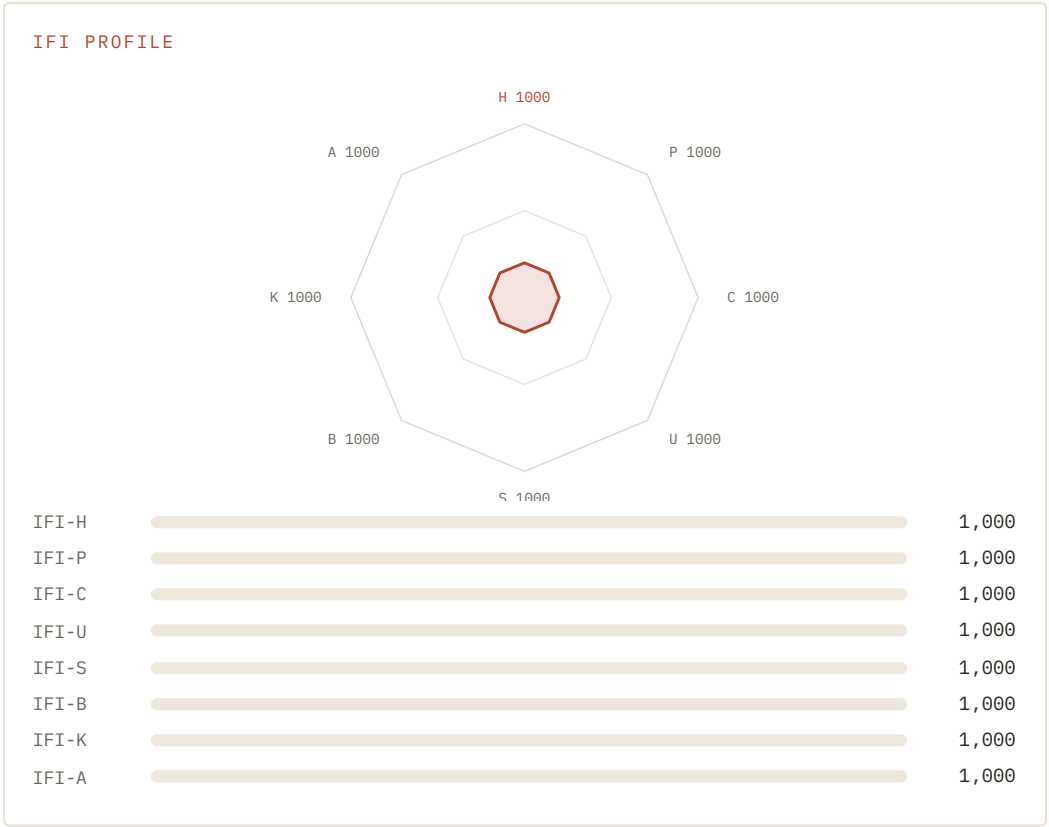

Third, **the score is a pure ratio**, which means the absolute EC_{50} — the plate, the cell line, the day — cancels, provided the reference is run alongside the candidate in the same assay. This is the quiet reason the Index is portable across organisms that share no chemistry: we are never reporting an EC_{50}, only the ratio of two of them measured together.

Work it on real heat. Anchor the reference extract at an EC_{50} of 3.00 µM capsaicin-equivalent (illustrative, stated so the arithmetic is checkable), and the formula runs backward to give every rung of the heat ladder its implied EC_{50} — EC_{50}(candidate) = 1000 × 3.00 ÷ IFI-H. The SHU and IFI-H values below are computational predictions from our flux and binding models; the progression from Ember to Phoenix reflects increasing design scope and axis activation, not a certified ascending measurement:

CULTIVAR · SKU	PREDICTED SHU	IFI-H	IMPLIED EC_{50} (MM)	95% CI AT ≤7% CV
reference anchor	—	1,000	3.00	–
Ember · SS-0087	866K	1,570	1.91	1,355 – 1,785
Blaze · SS-0091	1.72M	2,630	1.14	2,269 – 2,991
Firestorm · SS-0094	2.47M	3,490	0.86	3,011 – 3,969
Solara · SS-0095	2.72M	3,980	0.75	3,434 – 4,526
Phoenix · SS-0096	**2.95M**	**4,200**	**0.71**	**3,624 – 4,776**

There is a small marvel hiding in the last row. At full design, our model predicts Phoenix's standardized extract would half-activate TRPV1 near 0.71 µM — essentially the concentration at which *pure crystalline capsaicin* half-activates the same channel in a cell assay, roughly 0.7 µM. Our design targets a fruit extract that speaks to the nerve almost as loudly as the pure molecule — computed potency converging on that of purified capsaicin itself. And Phoenix carries all eight axes at once (each axis scored against its own stated reference convention; values on different axes are not on a single common scale):

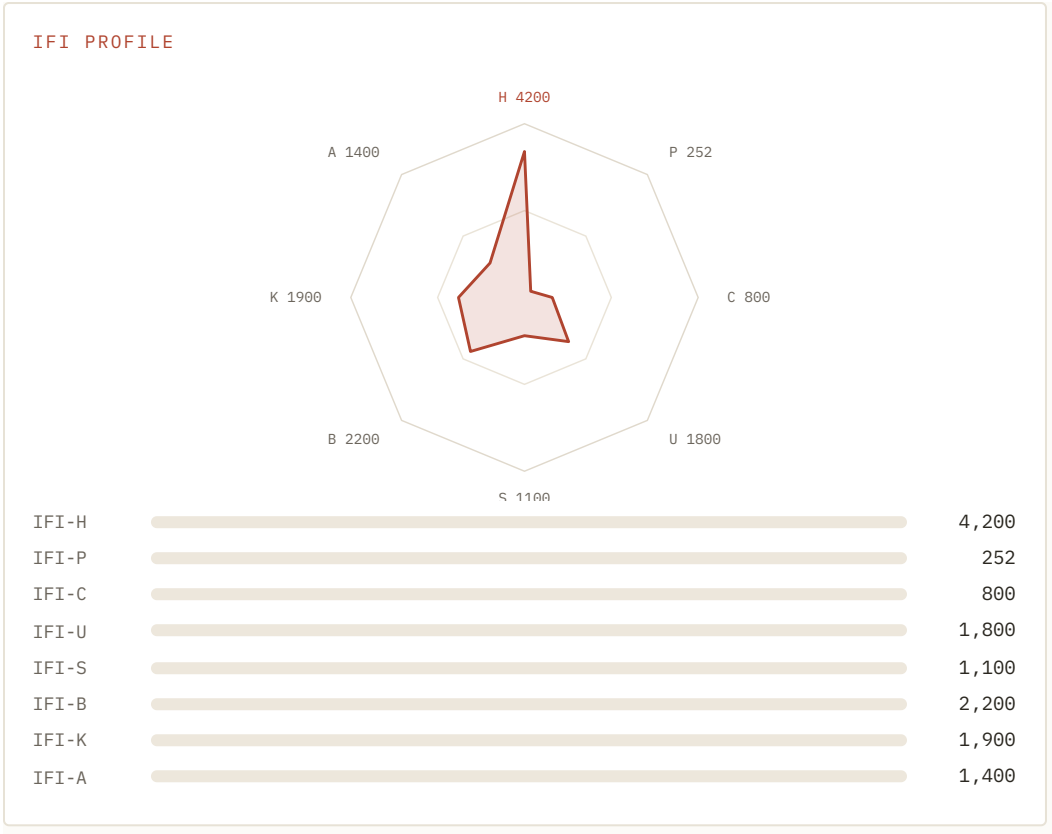

Four of those eight values trace straight back to the spine of this chapter. Heat (4,200) is the vanillyl branch condensed to capsaicin. Astringency (1,400) and much of the bitter score (2,200) are the flavonoid bough — flavan-3-ols and condensed tannins drawn off *p*-coumaroyl-CoA. Even the kokumi reading (1,900) leans on vanillin-adjacent chemistry at the spine's end. One pathway, four of the eight numbers on the card.

A word on units, because the Scoville scale and the micromolar scale have to be made to talk. Pure capsaicin is defined at **16,000,000 SHU**, which fixes the conversion **1 ppm capsaicinoid ≈ 16 SHU**. Capsaicin's molar mass is 305.4 g/mol, so **1 ppm = 1 mg/L ÷ 305.4 g/mol ≈ 3.27 µM**. Those two constants let any predicted Scoville figure be walked into molarity: our model predicts approximately 2.95 million SHU for Phoenix; at that predicted heat level, 2,937,948 ÷ 16,000,000 = 0.184 — a designed extract computed to speak at roughly eighteen percent of crystalline capsaicin's full voice. The organoleptic unit and the receptor unit are the same measurement read in two alphabets.

One more derivation, because steepness matters as much as position. A receptor's dose–response follows a Hill curve, response $R(c) = R(max) \cdot c^n \div (EC_{50}^n + c^n)$, where the exponent n is the cooperativity. For capsaicin at TRPV1, n is near 2. Solve the curve for the concentration that gives a fraction f of full response and you get $c = EC_{50} \cdot (f \div (1–f))^{(1/n)}$. Put f = 0.1 and f = 0.9 and the span from barely-there to nearly-saturated is $c(90)/c(10) = 9^{(1/n)} \div (1/9)^{(1/n)} = 9^{(2/n)}$. With n = 2 that is a nine-fold concentration range; with n = 1 it would be eighty-one-fold. The higher the Hill coefficient, the more switch-like the burn — a threshold crossed rather than a slope climbed. It is why heat arrives the way it does: not a dimmer, but closer to a door.

The arithmetic of error — why ≤7% is the whole game

An index is only as real as its smallest resolvable difference, and that is a question about variance, not about means. Here the design earns its keep. Because the IFI is a ratio, its relative error is the two input errors added in quadrature:

```
CV(IFI) ≈ √[ CV(reference)² + CV(candidate)² ].
```

Our engineered constructs hold a batch-to-batch CV of **≤7%**, against **42%** in wild-type material. Run a candidate at 7% against a tightly held reference standard near 3%, and the index itself carries $\sqrt{(0.07^2 + 0.03^2)} \approx 7.6\%$ — call it eight. Run the *wild* material the old way, at 42%, and the index inherits essentially all of it. Translate that into a score. At 7% CV, Phoenix's 4,200 sits inside a 95% interval of about 4,200 ± 576, roughly 3,620 to 4,780. At 42% CV the same 4,200 spreads to 4,200 ± 3,460 — a band from about 740 to 7,660, which is to say no information at all. Consistency is not a manufacturing nicety here; it is the precondition for the number meaning anything.

Push the point to where it gets honest. Two IFI values are cleanly distinguishable (non-overlapping 95% intervals) only when their difference exceeds about $1.96 \times \sqrt{(SD_1^2 + SD_2^2)}$. At 7% CV, Blaze (2,630) and Firestorm (3,490) differ by 860, comfortably past the ~600 needed — the ladder is readable at its lower rungs. But Solara (3,980) and Phoenix (4,200) differ by only 220, well under the ~790 that would separate them by heat alone. **At the ceiling, heat saturates the resolution of the ruler.** The reason Phoenix is nonetheless unmistakable is not its heat but its breadth: it registers on eight axes where Solara registers on six. When the first number stops discriminating, the other seven begin — which is the quiet structural argument for a full-spectrum index rather than a single hot number. At 42% CV, by contrast, even the 860-point gap between Blaze and Firestorm vanishes into the noise; the entire ladder collapses to a smear. The method did not invent the heat ladder so much as it made the ladder *measurable*, by driving the variance down until the rungs stood apart.

Averaging helps, within limits: the CV of a mean of k replicates falls as $CV/\sqrt{k}$, so three independent runs pull a 7% construct to about 4%. We use it, and we do not oversell it — $\sqrt{k}$ buys precision slowly, and it cannot rescue a process whose batch-to-batch spread is 42% to begin with.

Three honest limits belong on the same page as the math. First, the IFI deliberately folds *potency* and *abundance* into one effective concentration: a food can be loud because its agonist is intrinsically strong or because there is simply a lot of it, and a single EC_{50} cannot tell those apart. We consider this a feature — it matches what the tongue actually experiences — but it is a modeling choice, and we say so. Second, the absolute EC_{50} is assay-dependent; only the *ratio* is portable, so a profile is meaningful only relative to its reference, never as a free-floating number. Third, a receptor in a dish is not a mouth: the Index measures the opening move of flavor, the lock and key, and not the full sentence the brain writes from it. We built the number to design against, not to replace the tasting.

The spine runs everywhere — the arithmetic of a platform

Here is the thesis made quantitative, and the plain reason forty-two organisms were reachable and sixty-four more are on the horizon. Of the seven functional modules in the capsaicin-design platform, **five transfer unchanged** to any organism with a sequenced genome, because five of them act on the universal spine every plant already possesses.

MODULE	WHAT IT DOES	TRANSFERS?
1 · Committed entry	PAL and its MYB regulators	universal
2 · Activation	C4H hydroxylation + 4CL CoA-thioester	universal
3 · Methylation	HCT / COMT toward ferulate and vanillin	universal
4 · Regulatory logic	carbon commitment into the branch	universal
5 · Measure & design	the EC_{50} → IFI loop, species-indifferent	universal
6 · Terminal condensation	Pun1/CS makes the amide	pepper-specific
7 · Fatty-acyl branch	BCAT / KAS build the tail	pepper-specific

Five already in place, two to port: that is the arithmetic of a platform, and the reason the organism count could climb from one to forty-two without the method being rebuilt each time. The universality is not rhetorical — it is visible in the library's own hosts, each one running the spine and expressing a different one of its boughs.



FIG. 16·3 — Coffee. Its signature bitterness is chlorogenic acid — a caffeoylquinate ester formed one transfer off the shared spine, the same caffeoyl chemistry a pepper routes toward vanillin.

ORGANISM	CONSTRUCTS	PRIMARY AXES	SPINE PRODUCT IT DRAWS ON
Pepper	1,060	all	capsaicinoids (full cap)
Grape	1,610	B / A / M	condensed tannins (flavonoid bough)
Tea	280	U / B / A	catechins / flavan-3-ols (flavonoid bough)
Coffee	40	B	chlorogenic acid (caffeoyl ester)
Sorghum	120	A	lignin + 3-deoxyanthocyanidins
Vanilla	40	K	vanillin (spine's end)

Read that table and the claim stops being a slogan. Coffee's bitterness *is* chlorogenic acid, a caffeoylquinate formed one BAHD transfer off the trunk — the very caffeoyl chemistry a pepper routes

the other way toward vanillin. Tea's astringency and grape's tannin are flavan-3-ols and their condensed polymers, straight off the flavonoid bough. Sorghum's astringency is lignin and 3-deoxyanthocyanidins, two more products of the same few intermediates. Vanilla is the spine's own endpoint. Six organisms, six different numbers on six different axes, every one of them a reading of the same pathway this chapter has been tracing. The forty-two share a trunk; only their caps differ.

So the method is not a pepper technology that happens to generalize. It is a *plant* technology that happened to be proven first on peppers. The same measure-and-design loop that built a computationally predicted nearly three-million-Scoville chili can restore umami to a tomato, ease bitterness from a coffee, draw tannin into a grape, or finish menthol cleanly in a mint — different terminal chemistry each time, the same spine and the same five transferable modules underneath. We say the honest part too: the terminal chemistry is genuinely different for each flavor, and the frontier organisms are harder the further they sit from a pepper's own family, which is why the roadmap is tiered by gene count and years rather than promised all at once. But the hard part is bounded and known, and the easy part — most of the pathway — is already alive in the plant before we touch it.

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

The universality of the phenylpropanoid pathway is the quiet structural fact on which the entire library, and the entire roadmap, is built. It is also, if you let it be, a small daily wonder: that the stem holding up a flower and the fire inside a chili are two readings of one ancient sentence of metabolism, that the scent of a custard and the backbone of a burn are a single enzyme apart, and that we have learned — carefully, reverently, and with the arithmetic shown — to read the sentence and add a word at its end.

Chapter 17 · Forty-Two Hosts — *the organism library.*

A method is a promise; an organism is a fact. A design loop that turns a sequence into a flavor, and a biosynthetic spine that carries that flavor from one genome into the next, are elegant arguments — but an argument proves nothing until a living plant, grown from seed in our own light, carries the number we asked it to carry and hands it to the tongue. This chapter is the library as it actually stands: forty-two host species growing, carrying roughly nineteen thousand engineered constructs between them, a thousand of those live and purchasable today. It is the place where the thesis stops being a thesis. Forty-two genomes, each one persuaded — some gently, some over years — to speak a taste it was never born to speak, and to keep speaking it, true to the number, generation after generation.

The atom of the work — *what a construct is.*

The figure nineteen thousand hangs entirely on one word, so we define it before we spend it. A **construct**, in our usage, is one designed genetic change-set engineered into a host: an edit, or a coordinated group of edits, together with the regulatory sequence that governs when and where in the plant it switches on — a single complete unit of designed function. A cultivar may carry one such unit or many. A deeply worked organism like the tomato, at two thousand and forty constructs, stands for two thousand-odd distinct designed variants explored and held in the library, not two thousand products on a shelf. The construct is the smallest complete thing the design loop produces, the atom beneath every cultivar, and counting constructs — rather than products — is the truest measure of how much flavor-engineering a genome has actually absorbed. It reads the thinking done, not the things sold.

Counted that way, the library's center of gravity becomes legible at a glance. Here is the full ledger, organism by organism, with the axes each host was built to speak and the difficulty of making it speak them:

ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Pepper	1,060	all	Native
Tomato	2,040	S · U · F · M	Easy
Habanero	3,190	F · M	Hard
Black Mustard	2,790	P · M	Hard
Peppermint	2,360	C · M	Hard
Grape	1,610	B · A · M	Hard
Stevia	910	S · F · M	Moderate
Garlic	560	K · P	Moderate
Broccoli	380	B · M	Moderate
Soybean	360	U · K	Frontier
Tea	280	U · B · A	Hard
Broccoli-Kale	280	B	Hard
Yeast	240	U · K	Hard
Sugarcane	200	S	Frontier
White Mustard	200	P	Hard
Lavender	160	C	Hard
Blueberry · Ginger · Orange · Pomegranate · Sorghum · Thyme · Watercress	120 each (840)	A / P / B·C / A / A / P / P	Moderate–Frontier
Potato · Hops · Koji · Monk Fruit · Eucalyptus · Shiitake · Spearmint	80 each (560)	B / B / U / S / C / U / C	Easy–Hard
Coffee · Cabbage · Kale · Onion · Raspberry · Camphor · Cinnamon · Mushroom · Oregano · Radish · Rapeseed · Vanilla	40 each (480)	B / B / B / K / A / C / P / U / P / B / B / K	Moderate–Hard
42 organisms	**18,500**	**all 8 axes**	**Native → Frontier**

Add the column honestly and the forty-two named hosts sum to eighteen thousand five hundred constructs. The program reports roughly nineteen thousand as its working total; the few hundred between the two are constructs in active development, not yet posted against any one public organism line — the live front of the work, counted but not yet filed. We print both numbers rather than smoothing them into one, because a library that will seed a shelf of books cannot afford to let its own arithmetic drift, and the surest way to keep a figure honest is to show the sum it comes from.

Reading the library by its numbers.

Eighteen thousand five hundred constructs spread across forty-two genomes average four hundred and forty per organism — but the average is the least truthful number in the whole ledger, because the distribution is steeply skewed. The median host carries one hundred and twenty constructs, less than a third of the mean. A handful of deeply worked genomes drag the average far above the middle of the pack, which is exactly what a program concentrating its effort should look like.

Six organisms tell most of the story. Habanero, black mustard, peppermint, tomato, grape, and the pepper itself together hold thirteen thousand and fifty constructs — **seventy percent of the entire library in fourteen percent of its organisms.** The long tail is

the mirror image: twelve hosts sit at forty constructs apiece, four hundred and eighty in all, a bare two and a half percent of the work spread across more than a quarter of the species. The shape is deliberate. Where an axis could be pushed to its edge, we poured in constructs by the thousand; where a genome needed only a flag planted — proof that the method reaches it at all — a few dozen sufficed.

Sorted by engineering difficulty, the same skew resolves into something sharper still:

TIER	HOSTS	CONSTRUCTS	SHARE OF LIBRARY	MEAN / HOST
Native	1	1,060	5.7%	1,060
Easy	2	2,120	11.5%	1,060
Moderate	10	2,330	12.6%	233
Hard	26	12,310	66.5%	473
Frontier	3	680	3.7%	227
All tiers	**42**	**18,500**	**100%**	**440**

Two-thirds of every construct in the library lives in a genome rated **Hard.** That single row overturns the natural assumption that easy organisms would be the most developed and hard ones left thin — the table says almost the opposite. And the Hard tier is itself two things wearing one label. Four of its twenty-six members — habanero, black mustard, peppermint, grape — carry nine thousand nine hundred and fifty constructs between them, a mean of nearly two thousand five hundred each, and **more than half of the whole library in four organisms.** The other twenty-two Hard hosts average one hundred and seven. The tier is a mountain range with four peaks and a wide foothill, and the peaks are not where the work was easiest. They are where it was most worth doing.

The five tiers — *where a genome fights back.*

The difficulty tiers are worth naming carefully, because they map the real frontier of the work — not how valuable an organism is, but how hard its genome is to make speak a given axis.

Native is the pepper, where the terminal capsaicinoid machinery already lives and the work is tuning rather than building. **Easy** hosts — the tomato, the potato — are close relatives or simple, well-understood metabolisms where the spine is accessible and the edits land cleanly the first time. **Moderate** hosts take more coaxing: a less tractable genome, a pathway that needs a step or two rebuilt. **Hard** hosts demand substantial pathway construction, often several introduced enzymes working in concert. And **Frontier** hosts — soybean, sugarcane, sorghum — are the genomes at the very edge of what the platform can presently reach, where a single axis is a multi-year campaign and every success teaches the next.



Soybean, a frontier genome — even at the hard edge of the library, the same spine and the same loop apply, only slower, and against more resistance.

Difficulty is not a verdict on an organism; it is a measurement of distance, and the distance is counted in steps where a genome fights back. A tractable host accepts an edit, builds it into its genome, switches it on where we ask, carries the flux to the new chemistry without starving itself, and passes the trait to its seed — all on the first honest try. A Frontier host resists at one or more of those stations: it is slow or recalcitrant to transform, or it refuses to regenerate a whole plant from an edited cell, or it silences the new gene a generation after accepting it, or its metabolism simply cannot spare the carbon the new pathway demands. Soybean at three hundred and sixty constructs and sugarcane at two hundred are not thin because the work was neglected; they are thin because each construct there is bought at a far higher price in time and failed attempts. The tier, in the end, is just the program being honest about where in the long pipeline a given plant makes us earn it.

From a sequence to a plant that breeds true.

Getting a design from a sequence into a living, breeding plant is its own discipline, and it is the discipline the timelines are really spent on. The journey has a beginning, a slow middle, and a set of gates that nothing crosses without passing each in turn.

The beginning is delivery. An edit must be carried into plant cells, and classically that courier is *Agrobacterium tumefaciens* — a soil bacterium that has been transferring a loop of its own DNA into plant genomes, for its own ends, since long before we borrowed the trick. We disarm it, load our construct onto the stretch of DNA it naturally transfers, and let it do what it has always done, only now in our service. Increasingly we also edit directly: protoplasts, plant cells stripped of their rigid walls, taken up to the editing machinery and returned, so that the change is written into the genome with no foreign scaffold left behind. Which courier we use is itself part of the difficulty — some genomes welcome *Agrobacterium* and some rebuff it, and a host that must be edited as bare protoplasts is a harder host by that fact alone.

The slow middle is regeneration. A single edited cell is not a plant; it is a promise of one. It must be coaxed through tissue culture — hormone by hormone, week by week, in trays under steady light — to divide into a mass of callus, then to form a shoot, then a root, then

a whole plantlet hardy enough to leave the sterile vessel for soil. This is the station where Frontier genomes most often stall. A pepper or a tomato regenerates almost eagerly; a legume or a grass may refuse to form a shoot at all, so that an entire axis waits years not on the chemistry of the flavor but on the stubbornness of a cell that will not become a plant.



Trays under lights in the growth facility — the slow middle of the pipeline, where an edited cell is coaxed, week by patient week, back into a whole plant.

Then come the gates, and the plant is designed to pass through them one at a time. The edit is sequenced in the genome, exactly where it was designed to land. The designed enzyme is built to express in the tissue we intended. The flavor molecule is designed to reach the concentration the model predicted. The IFI profile is designed to read where the model placed it. And finally the whole plant is bred to grow true and stable, holding its profile across generations rather than drifting back toward its wild self. Only a plant that clears all five is allowed to call itself grown.

Each gate is also a way to fail, and naming the failures is the most useful thing a method can do. A construct can be delivered but never integrate. It can integrate but be silenced a generation later by the plant's own methylation. The enzyme can express yet produce nothing, because the true rate-limiting step lay upstream where we were not looking and the precursor pool runs dry. The molecule can be made yet read wrong on the Index, because the acyl chain came out a carbon too long and the burn arrived in the wrong shape. The trait can be perfect in the first plant and segregate away in the second. Every one of these has happened; every one is why a line's honest timeline is measured to the stably breeding plant and not to the first green flush of expression. The first plant is a demonstration. The stable line is a fact. The years between them are stewardship, and stewardship keeps no schedule.

The index every host is measured on.

Every organism in this library, from the native pepper to the frontier grain, is read on the same ruler, because a library that could not be compared against itself would not be a library at all. The ruler is the Index, and its whole definition is a single ratio:

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

The IFI is a stated convention: each axis defines its own reference constant, and IFI = 1,000 on that axis by definition — the same way the Scoville scale is defined by dilution-to-threshold rather than by a single measured value. Because each axis has its own reference constant, an IFI-H score and an IFI-P score each measure loudness relative to their own axis's defined reference, not against a common absolute scale; the Index makes organisms comparable *within* an axis, not across axes.

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the half-activation point, the moment a compound has opened half the locks it will ever open. For heat, that receptor is TRPV1, the trigeminal cation channel that also answers to temperatures above 43 °C and to acid; capsaicin does not warm the tissue, it lowers the channel's thermal threshold, so the body reports a fire that is not there. A low EC_{50} means a stimulus that speaks loudly, needing only a little of itself to half-open the channel. The Index inverts that concentration, anchors it against a fixed reference constant, and scales by a thousand, so that loudness — not how much of a compound a food holds, but how loudly that compound speaks — becomes a number a pepper, a tomato, and a cup of tea can each be read on within their respective axes.

Two things fall straight out of the formula. First, the anchor: when a candidate's half-activation concentration equals the reference's, the ratio is one and the score is exactly 1,000. A thousand is therefore the loudness of the native reference itself, the zero-point of the scale, the sound an unedited plant makes. Second, you can read any score backward into the physical demand it implies. Rearranged, $EC_{50}(\text{candidate}) = 1{,}000 \times EC_{50}(\text{reference}) \div \text{IFI}$, so the full-spectrum pepper our family designed, Phoenix, with an IFI-H of 4,200, is designed to half-open TRPV1 at 1,000 ÷ 4,200 = 0.238 of the reference concentration — a bit under a quarter. By the Index's own definition it is designed to be 4.2 times louder on the heat axis than the plant it was built from.

Here is the one profile the whole library is measured against — the design target that shows a single plant can be built to register on all eight axes at once:

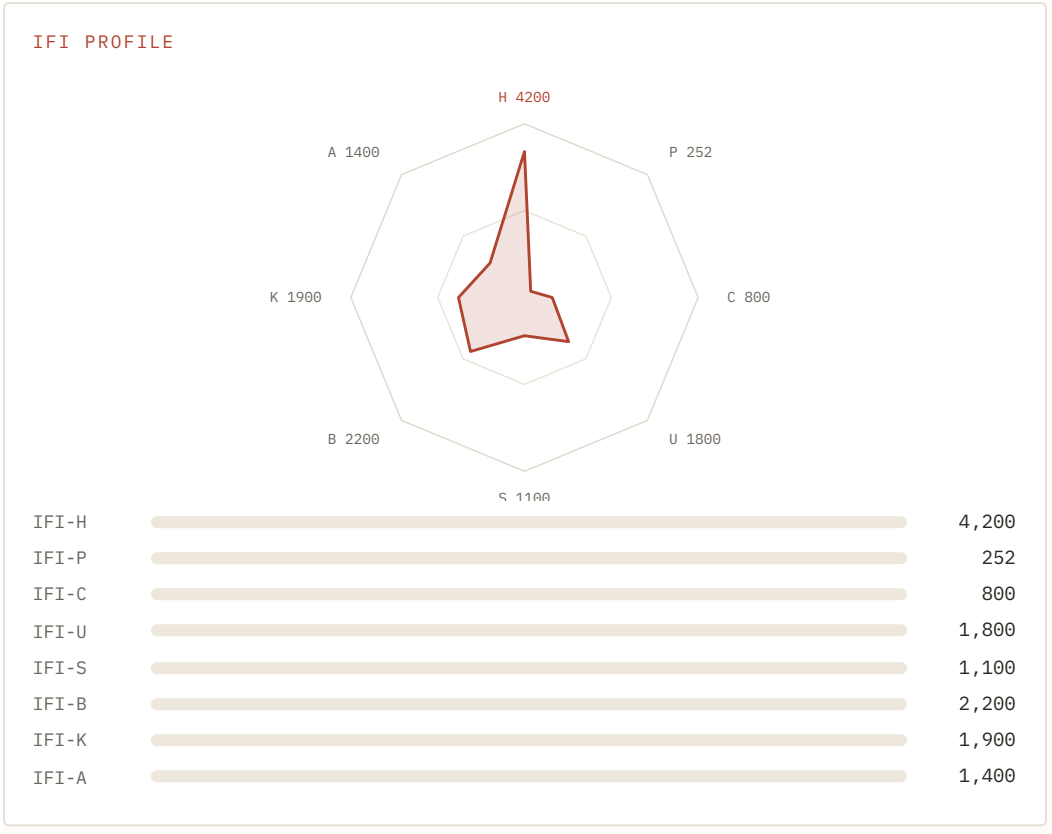

That is Phoenix, SS-0096 — our model predicts approximately 2.95 million SHU — the crowning cultivar of the native pepper line and the first designed to speak on all eight axes together. Its predicted heat sits at the top of the scale; its pungency, cool-burn, umami, sweetness, zero-bitter, kokumi, and astringency fill in beneath, each a designed number, each read on the same ruler as every other host in the ledger.

The loudness scale is readable at all only because the receptor compresses its own input, and that compression is worth seeing as mechanism, not magic. Receptor activation follows a dose–response curve — in the simplest honest form, response $\theta = [A]^n / (EC_{50}^n + [A]^n)$, where [A] is the stimulus concentration and n is the Hill coefficient that sets the curve's steepness. At $[A] = EC_{50}$, θ is exactly one half, which is what makes the half-activation point the natural anchor. The useful consequence is the dynamic range: carrying a receptor from one-tenth to nine-tenths of its full response takes a concentration change of 81^(1⁄n). For a cooperatively gated channel like TRPV1, with a Hill coefficient near 2, that is roughly a ninefold change in concentration spanning almost the whole range of sensation. Enormous swings in how much compound a fruit carries are folded, by the receptor, into a far gentler swing in what the tongue reports — which is precisely why loudness, not content, is the thing worth scoring.

The flagship pepper line makes the compression visible in canonical numbers. Read it as a worked ladder, with two columns computed from the index — the candidate's half-activation concentration as a fraction of the reference's (1,000 ÷ IFI-H), and the fruit's pungency as a fraction of pure capsaicin, which by the Scoville scale's own definition tops out at 16,000,000 units:

CULTIVAR · SKU	SHU (PREDICTED)	IFI-H	EC_{50}(CAND) ÷ EC_{50}(REF)	% OF PURE-CAPSAICIN PUNGENCY
Ember · SS-0087	866,000	1,570	0.637	5.4%
Blaze · SS-0091	1,720,000	2,630	0.380	10.8%
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%

Across that line the half-activation concentration falls from 0.637 of the reference down to 0.238 — each cultivar designed to answer the channel at ever-lower dilution, heard with less of itself. But notice the two spans do not match. From Ember to Phoenix the pungent concentration climbs 3.39-fold, while the loudness the Index records climbs only 2.67-fold. The receptor quietly swallows about a quarter of the concentration gain — the saturating curve flattening near its top, exactly as the Hill form predicts. This is the whole reason the Index is an EC_{50} measure and not a content assay: it reads what the tongue integrates, not what a chemical panel counts, and the two part company precisely where it matters most, at the loud end.

Precision is what makes it a library.

A construct is only an asset if it reproduces. A genetic change that gives 4,200 one harvest and 6,000 the next is not a library entry; it is an anecdote. What turns eighteen thousand five hundred constructs into a *library* — a set of dependable, reusable units a formulator can build on — is that each one is designed to hold its number, batch to batch, at a coefficient of variation of seven percent or better.

The coefficient of variation is just the standard deviation divided by the mean, the spread of a thing expressed as a fraction of its target, and seven percent is a tight band. On a construct aimed at IFI-H 4,200, a seven-percent CV is a one-sigma spread of ±294 points — roughly 3,900 to 4,500, a window narrow enough to design a recipe inside and trust the result. That reproducibility is not a cosmetic virtue; it is the precondition for everything else in this book. You cannot stack what you cannot predict, and you cannot sell what you cannot repeat.

Reproducibility also compounds in a way that rewards the fuller cultivars rather than punishing them, and the arithmetic is worth doing. When you combine n independent, equally variable units — replicate runs of a construct, or independently engineered edits averaged toward one target — the variations partly cancel, and the CV of the combination falls as the original CV divided by the square root of n. Two independent seven-percent units combine to $7\% \div \sqrt{2} \approx 4.9\%$. Four combine to $7\% \div \sqrt{4} = 3.5\%$, half the scatter of a single one. A cultivar built from several constructs is therefore not more fragile for its complexity; it is, on this measure, steadier than its parts. The precision does not merely survive being stacked — it improves under it, which is the quiet mathematical reason an eight-axis plant can be designed to hold its whole profile at once rather than wobbling on every axis it adds.

Ambition, not difficulty, decided where the work went.

Return, with the numbers now in hand, to the thing the ledger first seemed to say and the tier table flatly contradicts. One might assume the easy organisms would be the most developed and the hard ones left thin. Instead, two-thirds of the library sits in Hard genomes, and four of them — all rated Hard — hold more than half of everything. Difficulty did not decide where the work went. **Ambition did.**

We poured constructs into the organisms where an axis could be pushed to its very edge, and accepted the hard genomes as the price of that edge: habanero and the superhot peppers for the full-spectrum heat work, black mustard for pungency, peppermint for cool-burn, grape for astringency. Each of those is a genome that fought us, and each is where a flavor could be taken further than anywhere else, so we paid the price and kept going. The depth of work follows the depth of ambition on an axis, never the size of a market or the ease of a host.



The tomato, deeply worked — two thousand and forty constructs restoring the umami and sweetness that distance and refrigeration take, grown back in through the plant's own routes.

The Easy tier makes the same point from the other side. The tomato is rated Easy, yet it carries two thousand and forty constructs — the deepest body of work in the library outside the superhot peppers — because its sweetness and umami and full-spectrum potential were worth that many tries, and because an accommodating genome let each try land cleanly. Easy did not mean neglected; it meant the ambition met little resistance, so the constructs accrued quickly. And the long tail of forty-construct hosts — coffee, kale, onion, vanilla — is the breadth of the thesis made concrete: each a single axis proven in a new genome, each a small flag planted on the claim that the

method generalizes, each a toehold from which a deeper campaign can later climb. Depth and breadth are both in the ledger, and both are honest; they simply answer different questions.

A moving front — *why the library compounds.*

The library grows the way a craft grows, each hard-won organism paying for the next. A capability earned on a Frontier genome — a way through a wall of regeneration, a trick for stabilizing a reluctant edit, a promoter that will not be silenced — does not stay with the one plant that taught it. It lowers the difficulty of every related genome behind it, so that today's multi-year campaign becomes next season's routine port. The three Frontier hosts carry only six hundred and eighty constructs between them, the thinnest tier in the ledger, and that thinness is not failure; it is the cost of tuition. What those genomes teach is paid forward into the Hard and Moderate tiers, where the lesson, once learned, is cheap to reuse.

This is why the forty-two are not a finished set but a front that keeps moving, and why the roadmap of sixty-four concept organisms still ahead is a design document rather than a fantasy: it is drawn by people who have already watched *impossible* become *hard* and *hard* become *easy*, forty-two times over. Each organism is at once a product and a lesson, and the lessons compound — which is the real reason a program that began with a single pepper can speak seriously of grains and berries and the far synthetic edge.

Honesty requires naming what has not yet yielded. Some Frontier axes remain multi-year campaigns with no shortcut yet found. Some hosts will likely never reach the top of a given axis, because the precursor pool their metabolism can spare is simply too small, and no amount of edited enzyme can pour from an empty cup — a limit we would rather state than engineer around with a claim. And the nineteen thousand constructs are not nineteen thousand products: a thousand are live today, and the rest are design knowledge held in the library, waiting for the cultivar or the season that will call for them. A construct proven and shelved is not a failure; it is capital. But it is not a thing for sale, and we will not let the two numbers blur.

The numbers we hold.

Together the forty-two hosts resolve to the canonical figures the whole series is locked against, so that nothing drifts from book to book or page to page: **roughly 19,000 engineered constructs — 18,500 of them attributed across the forty-two named organisms here, the balance in active development; 1,000 live and purchasable today; 78,800 SKUs across all forty catalog views; 42 host organisms; 8 sensory axes; and a batch coefficient of variation of 7 percent or better.** Hold those figures. They recur throughout this library, and they are deliberately the same everywhere they appear, because a single drifting number read across a shelf of books is the one thing that reads as carelessness.

Underneath the arithmetic runs the conviction that has governed the work from the first pepper: that a plant is something to revere and to grow toward the fullest thing it can carry — the fullest flavor, and the

fullest nourishment woven through it — and that designing the one and designing the other are, in our family's hands, a single act. Forty-two organisms is where that work *is*: growing, counted to the construct, and read on one honest ruler. The sixty-four still unbuilt are where it is going. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 18 · The Sixty-Four — *what could be built.*

Forty-two organisms are grown. This chapter is about the sixty-four that are not.

They are a roadmap, but we would rather call them a design document, because every organism named here is reachable by a method established in roughly nineteen thousand catalogue designs built on some two hundred and twenty-five distinct constructs, in the thousand cultivars available today, in a single pepper our family designed gene by gene until it registered on all eight axes at once. The sixty-four are not wishes. They are specifications waiting for their sequence: each one a host species, a target profile on the eight axes, a gene count, and a span of years, written down the way you write down the name of a child not yet born — carefully, with reverence, because naming a thing precisely is the first honest act of bringing it into being.

We write them as a document and not as a promise for a reason that matters to us. It would be easy, and a little intoxicating, to say the whole garden will be remade by next spring; a grander claim always makes a better headline. We print gene counts and year-ranges instead, and we tier them by real biochemical distance, because the alternative is the drift this library refuses everywhere else. A frontier honestly mapped is worth more than a frontier breathlessly announced. It tells a partner what a thing will actually cost. It lets the imagination run without lying to it. And it means that when a line does cross from this roadmap into the grown library, the crossing can be trusted.

The shape of a frontier

The frontier sorts into tiers, and the tiers climb away from the familiar. What climbs with them is a single number: the count of genes a given plant sits away from the flavor we want it to carry.

The lowest tier stays inside the nightshade family, where the platform is most fluent — the tomato and its close kin, the organisms where the biosynthetic spine runs nearest the surface and the terminal enzymes are most at home. The next reaches a little further out within that same family. The third leaves the nightshades entirely for the berry, the cucurbit, the pome fruit — the hard and thrilling middle of the map, where the method must prove it travels. And the highest tier is the long horizon: the grains and brassicas, where an axis has to be built almost from the ground up, a near-complete pathway written into a genome that never carried anything like it.



FIG. 19 — Eggplant, host of the *Dark Matter* line. A Tier-1 nightshade: the spine already runs near its surface, four or five edits from a flavor it has never worn.

The honest way to read that climb is as a budget. Each tier carries a gene count and a timeline, and the two track each other closely enough that the ratio between them is itself a planning number.

TIER	WHERE THE HOST SITS	GENES	TIMELINE	≈ MONTHS / GENE
Solanaceae Core	within the nightshades	4–5	18–24 mo	~4.7
Extended Nightshades	the nightshade margin	5–7	24–36 mo	~5.0
Cross-Family	berry · cucurbit · pome	7–10	36–48 mo	~4.9
Frontier Synthetic	grains · brassicas	10–12	48–72 mo	~5.5

The final column is not in the specification; we computed it, dividing each tier's midpoint timeline by its midpoint gene count. It tells a quiet, steadying truth. To first order, the roadmap costs about five months of engineering and stewardship per gene, and that figure barely moves across the whole climb — 4.7 months a gene at the nightshade core, 5.5 at the synthetic edge. The tiers are not separated by a per-gene cost that explodes as you go out. They are separated by *how many genes* stand between the host and the flavor. The distance from the familiar to the frontier is measured in genes, and each gene is paid for at roughly the same rate. That is why we can say the frontier is reachable and mean it: nothing on it asks for a new kind of work, only for more of the same careful work, done in order.

Where the near-linearity frays — and it does fray, upward, at the top — is the honest edge of the estimate, and we will come back to it. Ten or twelve genes that must act in concert are costlier per gene than four that act nearly alone, because the plant has to be taught not just to run each step but to run them together. The synthetic tier drifts above five months a gene for exactly that reason, and the drift is the kind of thing a design document is supposed to show you, not hide.

Gene count is really a distance

The gene count climbs with the tiers because of biochemistry, not ambition, and the biochemistry is worth seeing plainly.

Every flavor the platform designs rides on a shared upstream road. The phenylpropanoid branch — PAL → C4H → 4CL → HCT → COMT, five enzymes in sequence — runs in every vascular plant alive, because it is the road to lignin, to the walls that let a plant stand up. A tomato runs it. A grain of rice runs it. A pepper runs it. What makes

a pepper a pepper is not that road but four terminal enzymes bolted onto its end — pAMT, Pun1/CS, KAS, and BCAT — the small, specific crew that takes the shared intermediate the last few steps into a capsaicinoid. Five of the platform's seven engineering modules never touch that terminal crew at all; they operate on the shared upstream road, and so they transfer unchanged to any organism with a sequenced genome.

That is the whole secret of the roadmap's reachability, and also the whole reason the gene count varies. Inside the nightshade family, a host already runs most of the spine near its surface and sits close, in evolutionary terms, to the pepper whose terminal enzymes our family understands best. A flavor can be reached there by editing what is largely already present — four or five genes, and the better part of two years. Step out to the nightshade margin and a step or two of pathway has to be rebuilt rather than merely tuned. Cross into another family entirely — the strawberry, the cucumber, the apple — and the terminal chemistry must be substantially introduced, seven to ten genes working in concert, because the target genome never ran it at all. At the synthetic frontier, in the grains and brassicas, the axis is built almost from nothing: ten to twelve genes across four to six years, a near-complete pathway written into a plant that carried no trace of it.

So the gene count is really a distance. It measures how far the host sits from the flavor, counted in the genes it would take to close the gap — a number you could in principle read off a phylogenetic tree, the branches between the pepper and the plant you want to teach. The roadmap is, underneath its cinematic names, a map of that tree with a price on every branch.

Reading a profile backward — the arithmetic of a target

Every organism on the roadmap begins as eight numbers, and the first discipline of the work is knowing precisely what those numbers mean, because a target you cannot read is a target you cannot design toward.

The Index scores each axis the same way:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the half-activation point, the moment a compound has opened half of the locks it is ever going to open. For heat that receptor is TRPV1, the trigeminal channel that also answers to temperatures above 43 °C; a low EC_{50} means a molecule that speaks loudly, needing only a little of itself to half-open the channel. The Index inverts that concentration, anchors it against a fixed reference constant, and scales by a thousand, so that loudness — not quantity, but *loudness* — becomes a score you can read a tomato, a chili, and a cup of tea against on a single page.

Two consequences fall straight out of the formula, and the entire roadmap lives inside both.

First, the anchor. The reference constant on each axis is a stated convention — defined, the way Scoville's dilution-to-threshold is defined, rather than derived from the literature. When a candidate's half-activation concentration equals the reference's — when EC_{50}(candidate) = EC_{50}(reference) — the ratio is one and the score is exactly 1,000. A thousand is therefore not an arbitrary round number; it is the loudness of the native reference itself, the zero-point of the scale, the sound an unedited plant makes. On two axes — heat and pungency — we also carry IFI-data versions cross-checked against published EC_{50} records; the convention governs the scale in both cases.

Second, and more useful for a design document: you can read the number backward. Rearrange the formula and it hands you the concentration a target *implies*:

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

Now a target profile becomes a set of physical demands. Take the gentle end first. A shishito registers near 40 on the heat axis; put 40 into the formula and the candidate's half-activation concentration comes out at 1,000 ÷ 40 = 25 times the reference's. You would need twenty-five times as much of its heat compound to half-fire the channel — which is exactly why a shishito is a sunlit warmth and not a fire. Now the other end. The full-spectrum pepper our family built, Phoenix, reads 4,200 on the same axis; 1,000 ÷ 4,200 puts its half-activation concentration at 0.238 of the reference's, roughly one part in four. It half-opens TRPV1 at under a quarter of the reference concentration — a molecule designed to be heard. The two targets sit 4,200 ÷ 40 = 105-fold apart in loudness, and the arithmetic closes on itself as a check: 25 × 4.2 = 105. The quiet pepper and the loud one are exactly one hundred and five times apart whether you measure them forward or backward, which is the small, satisfying proof that the scale is really a scale.

One subtlety keeps this honest, and it is a point about the receptor, not the number. The EC_{50} fixes only the *horizontal* position of a dose–response curve — the concentration at which the receptor is half-fired — and says nothing about the curve's height, the maximum sensation the compound can ever evoke. Two foods can climb to the very same peak and reach it at wildly different concentrations; the Index reads the position, not the summit, because position is what the tongue reports first and what formulation can actually move. When we specify a roadmap profile, we are specifying where each curve sits, not how tall it stands. That is a deliberate choice, and it is the reason the Index is an EC_{50} index and not a content assay.

Here is the one profile the entire roadmap is drawn against — Phoenix, our full-spectrum flagship, the design that shows an eight-axis target can be fully specified:

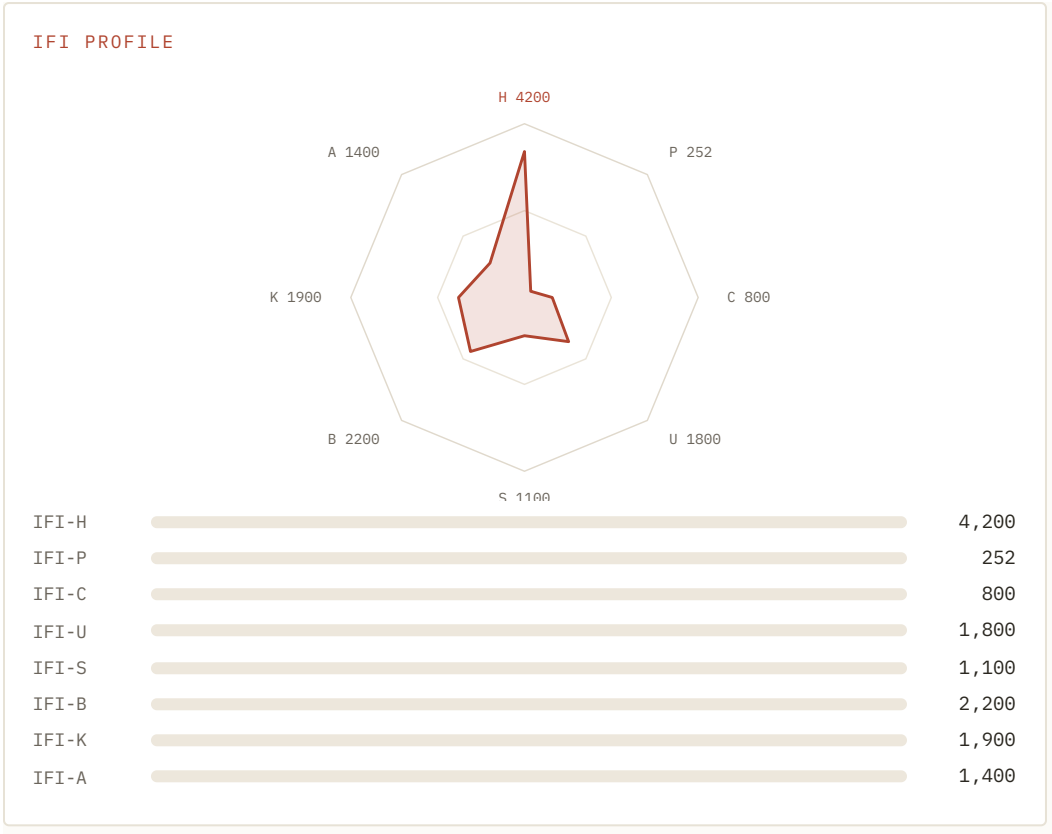

That is Phoenix — SS-0096, with a model-predicted heat of approximately 2.94 million SHU, the first cultivar we designed to register on all eight axes at once. Every number on the sixty-four is a target drawn against this profile; the made and the unmade are plotted on the same axes, to the same scale, so that a dream and a harvest can be compared in the same breath.

A word of caution belongs here, because it is the kind of thing a design document owes the reader. The heat axis and the old Scoville number are not the same specification, and the difference lives in our own flagship ladder. Scoville units come from a dilution-to-threshold panel — how far you must water a thing down before the burn disappears — while the heat axis comes from a receptor half-activation assay. They correlate, but they are not interchangeable, and the ladder shows exactly where they part company.

CULTIVAR	SKU	PREDICTED SHU	IFI-H	SHU PER 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
Solara	SS-0095	2,720,000	3,980	683
Phoenix	SS-0096	~2,940,000	4,200	700

Read the final column — predicted Scoville units divided by heat-axis points, computed from the two design-specification values — and watch it refuse to hold still. It runs from about 459 to about 708, drifting upward as the ladder climbs. If the two specifications tracked the same thing, that column would be a constant. It is not, and the sharpest tell sits near the bottom: Kindle carries more predicted total pungent mass than Cinder — 1.25 million predicted Scoville units against 1.05 million — yet reads *lower* on the heat axis, 2,000 against 2,290. More mass, less loudness. The dilution panel counts every pungent molecule in the fruit; the heat axis weighs how loudly the capsaicinoid *profile* speaks at TRPV1, and the balance between capsaicin and dihydrocapsaicin — set by the length and

saturation of the acyl chain the terminal acyltransferase installs — shifts that loudness independently of total mass. A pepper can be more pungent by the panel and quieter at the channel. This is precisely why the roadmap specifies its targets on the eight axes and not in Scoville units: the axis is what the tongue integrates, and the roadmap is written for the tongue.

The last piece of the arithmetic is the one that separates a design document from a wish, and it is error. A target of 4,200 means nothing without a tolerance around it, because a flavor that lands at 4,200 one harvest and 6,000 the next is not a product; it is a rumor. Our constructs are designed to hold a batch-to-batch coefficient of variation at or under 7 percent, against the 42 percent that wild selection drifts to. On a 4,200 target, a 7 percent CV is a one-sigma band of ±294 points — roughly 3,900 to 4,500, a window a formulator can build a recipe inside. At 42 percent the same nominal target scatters ±1,764 points, from the high 2,000s to nearly 6,000 — the difference between a cultivar you can design against and one that merely happens to you. Every timeline on the roadmap carries that tolerance inside it. The months are not the time to the first plant that expresses the trait; they are the time to a plant that expresses it *stably*, within the band, generation after generation. The error bar is part of the specification, not an afterthought to it.

The named series — a suite sketched before a note is written

This is where the roadmap stops being a table and becomes an imagination, because each line is specified down to its members, the way a composer sketches a suite before writing a single note.

TIER	ORGANISM · SERIES	MEMBERS
Core	Tomato · *Volcanic*	Ember · Magma · Pyroclast · Caldera · Krakatoa
Core	Eggplant · *Dark Matter*	Smolder · Nightshade · Eclipse · Obsidian · Event Horizon
Core	Tobacco · *Laboratory*	Proof of Concept · Benchmark · Threshold
Margin	Potato · *Seismic* · Ground Cherry · *Celestial* · Tomatillo · *Acid* · Petunia · *Floral*	—
Cross-Family	Strawberry · *Sweet Heat*	Warm Kiss · Sugar Burn · Candied Fang · Devil's Dessert
Cross-Family	Cucumber · *Contradiction* · Melon · *Solar* · Apple · *Forbidden* · Lettuce · *Espionage*	—
Synthetic	Carrot · *Deep Earth* · Cauliflower · *Nuclear* · Broccoli · *Wildfire* · Corn · *Storm* · Rice · *Dynasty*	—

Read the *Volcanic* tomato line and you can taste the intent before a seed is in the ground: Ember to Magma to Pyroclast to Caldera to Krakatoa, a heat ladder built into the fruit the platform knows best — a tomato that is first a tomato and then, deliberately, a fire. We can even read its ends off the design numbers without inventing anything. The tomato's own voice is sweetness and umami; the restored tomato is predicted to reach an umami score near 2,600 on its native route, from a wild baseline near 300. Ember would sit low on heat over that sweet-umami base, a warmth folded into fruit.

Krakatoa would push heat toward the design ceiling our superhot work has already computed, near 4,200, while keeping the tomato's own sweetness underneath — a target we can name as a range precisely because we have designed both of its anchors. The axes in between are a design question, not a computed fact, and we will say so every time.



FIG. 20 — The tomato in cross-section, host of the ⋆Volcanic⋆ line. First a tomato — sweetness and umami on the native route — and then, by four or five edits, a fire laddered from Ember to Krakatoa.

Read the *Sweet Heat* strawberry — Warm Kiss, Sugar Burn, Candied Fang, Devil's Dessert — and you see the swicy composition written across a berry: heat and sweetness held in deliberate proportion, a dessert that bites back. This one sits in the cross-family tier, seven to ten genes and three to four years, because a strawberry never ran the terminal capsaicinoid crew and has to be taught the whole last movement from scratch. That is the honesty of the gene count made edible — the berry is further from fire than the tomato is, and the roadmap charges it for the distance.



FIG. 21 — Strawberry in cross-section, host of the ⋆Sweet Heat⋆ line. A cross-family target: the berry never ran the terminal crew, so seven to ten genes teach it the whole last movement.

And then the far edge — *Storm* corn, *Dynasty* rice — staples remade. We linger on these not for spectacle but because a staple is where nourishment reaches the most tables, and our family has never thought of flavor and nourishment as two projects. A grain designed to carry more of what the body can draw from it is the same act as a grain designed to taste like more; both are the plant grown, on purpose, toward the fullest thing it can hold. These sit at

ten to twelve genes and four to six years for the plainest reason on the map: the grass is about as far from the pepper as the roadmap travels, and we will not pretend the distance is shorter than it is.



FIG. 22 — Corn, host of the synthetic-frontier *Storm* line. The far edge of the map: ten to twelve genes, four to six years, an axis built almost from nothing into a plant that carried no trace of it.

The cinematic names — Volcanic, Dark Matter, Seismic, Sweet Heat — are a deliberate device and not a marketing indulgence. A vivid name makes a specification memorable; it gives a multi-year campaign a flag to march under; it lets a whole family of cultivars be spoken of in one breath. Ember through Krakatoa is a heat ladder you can feel before it exists. We name them with care precisely because naming a thing well is the first real act of building it — and because the one pair of hands that does the computing is only one part of a family that will grow these, cook with them, and set them on a table, and a shared name is how a family holds a long piece of work together.

How a line crosses from roadmap to library

A line does not graduate from this chapter by being wished into the grown library; it crosses by the same pipeline that carries every cultivar our family has ever grown, run at the harder end.

The target profile is specified first, on the eight axes, as a set of numbers and the tolerance around them. A genome-scale flux-balance model then checks that the profile is metabolically feasible in the chosen host — that the plant can carry the flux to the new terminal chemistry without starving a pathway it needs to stay alive, because a flavor that kills its host is not a flavor. With feasibility established, the enzymes are chosen and the edits designed by the loop that scores substitutions for predicted effect on catalytic efficiency and fold stability and lets the flux model name the true rate-limiting step rather than the assumed one. The design is delivered, the plant regenerated, the trait assessed against the specified profile — and then, only then, the part that the timelines are really spent on: the cultivar must hold its profile stably across generations, inside the 7 percent band, before it is allowed to call itself grown. A plant that expresses the trait once and drifts the next season has not crossed over; it has only auditioned.

That last gate is why the honest estimate for any line is longer than the time to its first expressing plant, and why we print year-*ranges* rather than dates. The first plant is a demonstration. The stable line is a product. The road between them is stewardship, and stewardship does not keep a schedule.

What the counts quietly assume — edited, not imported

There is an assumption folded inside every gene count on the roadmap, and it belongs in the open, because it shapes both the science and the road to a table.

The reachability of these lines rests on editing a plant's *own* pathways and introducing known, characterized enzymes — precise, bounded changes, the point substitutions and pathway completions our edits are built from, not blunt transfers of whole foreign systems. The proof that such precision is enough sits in the double variant that built the superhot design ceiling the roadmap is drawn against. Pun1 is the terminal capsaicin synthase — the enzyme that labels the product as a capsaicinoid — but our flux model places rate control upstream in the phenylpropanoid pathway, at 4CL and HCT; the substitutions here refine the terminal step's binding affinity and fold stability rather than the upstream flux that governs overall capsaicin output. A protein language model scored every possible single-residue change in Pun1 by predicted effect on catalysis and stability; two substitutions that no breeding program ever stumbled onto emerged and were written into the design — S39L, which lifts binding affinity at the N-terminus (ΔΔG −2.1); and L345G, which relieves steric strain in the active-site cleft (−1.8). Our structural modelling also identified C175S as a disulfide misfold prevention candidate (ΔΔG −0.9); it is retained as a structural annotation but was not included in the final shipped construct. Two edits, about −3.9 kcal/mol of fold stabilization between them, and the batch CV fell from 42 percent to 7 or under. That is what *edited* means here: a few residues chosen with care, not a cassette of foreign genes dropped in and hoped over.

Where a design can be made by editing what a plant already has, rather than importing what it never did, it sits in a different and lighter regulatory category in most jurisdictions — the distinction regulators increasingly draw between a gene-*edited* crop that could in principle have arisen through breeding and a *transgenic* one carrying foreign genes. We design toward the former wherever the flavor allows, not to slip past scrutiny but because the lighter-touch path is also the cleaner science and the shorter road to a table. The timelines carry that assumption inside them; they are estimates of the work and the stewardship both. And where a flavor genuinely cannot be reached by editing alone — where a terminal crew must be introduced because the host truly never had one — we say so, and the line pays for it in genes and in years on the map you have already read.

The discipline of an honest frontier

There is a discipline in writing a roadmap this specific, and it is the same discipline that governs the grown library: say only what you can stand behind.

We have tried to show you where the estimates are soft, because a design document that hides its own soft spots is just advertising. The five-months-a-gene figure is a planning heuristic, and it frays upward at the synthetic tier, where ten or twelve genes must act in concert and epistasis makes the whole cost more than the sum of its parts — the plant has to learn not just each step but the choreography, and choreography is where lines stall. The two heat specifications disagree in the detail, and a target read in Scoville units would mislead a formulator who should be reading the axis. The error band is the difference between a cultivar and a rumor, and it is paid for in the quiet years after the first plant, not the loud week of the first result. These are not caveats we were forced into; they are the parts of the map we are proudest to have drawn, because they are what makes the rest of it trustworthy.

What remains, then, is the plainest statement of the thing. This is not a catalog of what our family has made; it is a map of what the method makes possible, with the made and the unmade drawn to the same scale and plotted on the same eight axes. The sixty-four are a promise with a budget and a conscience attached — dreams, but dreams with timelines, which is the only kind we are willing to print. Under all of it runs the conviction that holds the whole library together: that flavor, and the nourishment woven through it, are now things a family can design on purpose, name precisely, and grow toward, one honest number at a time. What remains is to make sure the numbers can be trusted, and to be clear what kind of numbers they are: these are computational designs, and our cultivars and constructs have not been experimentally validated in a laboratory. We have shown you how we count them.



PART V OF V

The Standard in Use

Precision as respect · reading the world · the index as a shared language

Chapter 19 · Precision as Respect — *a number means nothing without its error bar.*

Everything in this book builds toward producing a number. This chapter is about the harder thing: producing a number you can *trust* — one that means the same thing in the second fruit as in the first, in next season's harvest as in this one's, in a partner's laboratory as in ours. A score without a known error is a rumor. The discipline that turns it into a standard is reproducibility, and reproducibility is, in the end, a form of respect — for the reader, for the eater, and for the plant, which deserves to be described accurately.

There is a particular hush in a growing room where every plant has been given the same light, the same water, the same measured hours — and a particular kind of hope in cutting the first fruit and the fortieth and finding them the same. That sameness is not luck. It is the quietest and most demanding thing our family makes, and most of this chapter is the arithmetic of how it is held.

The coefficient of variation, defined

Begin with the number that matters most and is least glamorous: the **coefficient of variation**, the CV. It is the standard deviation of a measurement divided by its mean, written as a percent —

```
CV = σ / μ × 100%
```

— and it answers the only question that makes a score usable: *how much does this vary from one fruit to the next?* A mean alone is a headline with no article beneath it. The CV is the article.

For conventional crops the answer is sobering. A wild-type pepper's heat carries a CV near **42 percent.** Put a number to what that feels like at the cutting board. Take a landrace whose capsaicinoid content averages μ; with a CV of 42 percent its standard deviation is σ = 0.42 μ. Draw two fruit at random from the same plant, picked the same week, and the spread in their *difference* is wider still, because two independent draws add their variances: σ_diff = σ√2 = 0.42 μ × 1.414 ≈ 0.59 μ. The typical gap between two sibling fruit is therefore more than half the mean itself. One can carry nearly two and a half times the heat of the other and nothing has gone wrong — the heat was never controlled, only selected toward, and the underlying genetics still roll the dice in every ovary.

It helps to leave the abstraction and carry real units. The Scoville figure and the capsaicinoid mass are locked together by a fixed factor: one part per million of capsaicinoids registers close to **16 SHU**, because pure capsaicin sits near 16,000,000. Our model predicts approximately 2.95 million SHU for Phoenix — roughly 184,000 ppm (2,950,000 ÷ 16) of capsaicinoids, or close to 18.4 percent of the dried fruit's own weight given over to a single family

of molecules. A wild landrace at, say, 100,000 SHU carries about 6,250 ppm, and at a CV of 42 percent its one-sigma band runs from roughly 3,600 to 8,900 ppm. You cannot formulate against a band that wide. You cannot promise a chef, a sauce-maker, or a manufacturing partner that the thing they tasted and approved is the thing that will arrive by the pallet. A flavor that cannot be reproduced is an anecdote, however wonderful — and anecdotes do not feed anyone at scale.



FIG. 20 — Every plant on the same light, water, and measured hours. The sameness from the first fruit to the fortieth is the quiet achievement under every score in this book.

The designed cultivars hold to a CV of **7 percent or better.** At that spread the same pair of sibling fruit differ by σ_diff = 0.07 μ × 1.414 ≈ 0.10 μ — a tenth of the mean, not six-tenths of it. The sixfold collapse in variance, from 42 to 7, is a greater achievement than the peak heat, and we would ask a skeptic to weigh it first. Peak heat is dramatic and easy to doubt; a tight CV is undramatic and nearly impossible to fake, because it can only come from genuine control of the underlying biology. A loud score might be a good day. A small variance, held batch after batch after batch, is proof of method.

Two error bars under every score

Two different uncertainties live beneath every number in the Index, and conflating them is the most common mistake a reader can make.

The first is **measurement** uncertainty: how precisely the instrument pinned the value on the sample in front of it. The second is **process** variance: how much the underlying thing itself moves from fruit to fruit — the coefficient of variation just described. A score can be measured exquisitely on a wildly variable crop, or measured roughly on a perfectly uniform one. Only when *both* error bars are small is a number dependable.

Measurement uncertainty on the receptor axes comes from the shape of a curve. Each IFI score rests on an EC_{50} — the concentration at which a stimulus drives its receptor to half of full response — and that EC_{50} is not read off a dial but fit to a dose–response. We model the response with the Hill equation,

$$R([L]) = R_{max} \cdot [L]^n / (EC_{50}^n + [L]^n),$$

where [L] is the ligand concentration presented to the cell, n is the Hill coefficient that sets the steepness of the rise, and R_max is the saturated response. Set [L] = EC_{50} and the expression collapses to

R_max / 2 — the definition the whole scale hangs on. We present a graded series of concentrations to a line of cells expressing the human receptor, read the response at each, and fit the three parameters by nonlinear regression. The fit does not return a bare EC_{50}; it returns an EC_{50} *with a confidence interval,* and that interval is the measurement error bar. Replicate the curve across plates and days and the intervals either tighten or betray a drift. A well-run assay pins an EC_{50} to a few percent.



FIG. 21 — A dose-response read at the bench. An EC_{50} is never read off a dial; it is fit to a curve, and the width of that fit is the measurement error bar under every receptor score.

Capsaicinoid mass — the SHU side — is pinned independently, by HPLC, molecule by molecule, each sample injected in triplicate so that the three readings' standard error, $\sigma / \sqrt{3}$, rides along with the mass just as a confidence interval rides along with the EC_{50}. The Scoville figure is then an analytical fact, not a panel's impression. The receptor loudness behind each IFI score is read against a fixed reference, so the ratio cancels the day's drift: whatever makes the instrument read a little high this morning lifts reference and candidate together and divides back out. Growth happens under controlled light, water, and temperature so that the environment contributes as little variance as the genetics now do. It is quality control borrowed whole from pharmaceutical manufacturing and pointed at a pepper, because a flavor meant to be built on deserves to be built on something that holds still.

Metrology keeps two words for what a casual reader lumps together as precision, and the distinction is the whole promise of a standard. *Repeatability* is the agreement of measurements taken the same day, on the same instrument, by the same hands; *reproducibility* is the harder agreement across different days, instruments, and laboratories. A score that is merely repeatable travels no further than the bench that made it. The Index is built for the second, stricter kind — the fixed reference anchor, the validated assays, and the openly published methods all exist so that a partner running the protocol in another building arrives at the same number we did, inside the same error bar. That is what separates a number that is a standard from one that is a souvenir.

The arithmetic of the Index, worked

The Index is deliberately a ratio, because a ratio forgives the things a raw reading cannot.

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

The reference is the native wild-type half-activation, which we anchor so that it reads exactly **1,000** on the scale by convention — a stated definition, like how the Scoville scale is defined by dilution-to-threshold, not derived from a literature survey. Everything louder scores higher by the same proportion it is louder. To carry concrete units through a worked example, take that reference half-activation at an EC_{50} near **0.70 µM** — the anchor against which the canonical IFI values were set. Inverting the formula gives each candidate's own EC_{50}:

$$EC_{50}(candidate) = 1{,}000 \times EC_{50}(reference) / IFI = 0.70\ \mu M \times 1{,}000 / IFI.$$

Run Phoenix's predicted **IFI-H of 4,200** through it: EC_{50} = 0.70 × 1,000 / 4,200 = **0.167 µM.** Phoenix is designed to reach half-maximal TRPV1 activation at about a quarter of the reference concentration — 4.2 times louder, by construction. The whole flagship ladder inverts the same way.

CULTIVAR	SKU	SHU (PREDICTED)	IFI-H	EC_{50}-H (MM)	× REFERENCE
Ember	SS-0087	866K	1,570	0.446	1.57×
Cinder	SS-0088	1.05M	2,290	0.306	2.29×
Kindle	SS-0089	1.25M	2,000	0.350	2.00×
Blaze	SS-0091	1.72M	2,630	0.266	2.63×
Wildfire	SS-0092	1.96M	2,930	0.239	2.93×
Scorch	SS-0093	2.21M	3,220	0.217	3.22×
Firestorm	SS-0094	2.47M	3,490	0.201	3.49×
Solara	SS-0095	2.72M	3,980	0.176	3.98×
Phoenix	SS-0096	2.95M	4,200	0.167	4.20×

The table hides a lesson in its own disorder. Read the SHU column and the IFI-H column together and they do not rank identically: Kindle carries more total predicted capsaicinoid than Cinder — 1.25 million SHU against 1.05 — yet scores *lower* on heat loudness, 2,000 against 2,290. There is no error here. SHU counts how *much* capsaicinoid a fruit holds; IFI-H measures how *loudly* it speaks at the receptor, and loudness depends on which capsaicinoids are present — the balance of capsaicin to dihydrocapsaicin, set by the length and saturation of the acyl chain the acyltransferase installs — not on mass alone. Two error bars, two quantities; a reader who collapses them will mis-rank the very fruit the Index was built to separate.

The ratio form also tells us exactly how the measurement error propagates. An IFI is a quotient of two EC_{50}s, each carrying its own relative uncertainty, and for a quotient the relative variances add:

$$(\sigma_{IFI} / IFI)^2 = (\sigma_{ref} / EC_{50_ref})^2 + (\sigma_{cand} / EC_{50_cand})^2.$$

Pin each EC_{50} to 5 percent — a good assay — and the score inherits $\sqrt{(0.05^2 + 0.05^2)} = 0.05\sqrt{2} \approx$ **7.1 percent.** A predicted IFI-H of 4,200 therefore arrives as 4,200 ± 300, and that ± is not decoration; it is the honest width of the claim. To publish the 4,200 without it would be to assert a precision we had not earned.

One more property of the dose–response explains why the EC_{50}, and not some other point, anchors the scale: it is the steepest, most information-dense place on the curve, and therefore the most reproducible to locate. With a Hill coefficient near n = 2 — typical for capsaicin at TRPV1 — the fractional response at a concentration x times the EC_{50} is $f = x^n / (1 + x^n)$. Walk it across Phoenix's predicted EC_{50} of 0.167 µM: at a third of it (0.056 µM) the channel sits at 10 percent open; at half (0.084 µM), 20 percent; at the EC_{50} itself, exactly 50 percent; at twice (0.334 µM), 80 percent; at three times (0.501 µM), 90 percent. A ninefold span in concentration carries the receptor from a tenth to nine-tenths of full response, and the midpoint is where the curve climbs fastest — so a small error in reading the response translates into the smallest possible error in the concentration you infer from it. The anchor is placed where the ruler is finest. Push far above the EC_{50} and the curve flattens into saturation; drop far below and it flattens into threshold; in both tails a large change in concentration barely moves the response, and resolution is lost. The reproducible number lives in the middle, which is exactly where the Index reads it.

Where precision is written into the sequence

The tight CV is not held by sorting fruit after the fact. It is written into the enzyme, in the same edits that raise the heat.

Two substitutions in the Pun1/AT3 acyltransferase, scored by a protein language model for predicted effect on catalytic efficiency and fold stability and never before combined, carry the whole gain. **S39L**, at the N-terminus, deepens substrate binding; **L345G** relieves a steric clash in the active-site cleft so the catalytic step turns over faster. Both raise k_cat / K_M — they make the fruit hotter. The shipped construct is this double variant S39L+L345G.

EDIT	ΔΔG (KCAL/MOL)	STRUCTURAL ROLE	NOTES
S39L	−2.1	N-terminal substrate binding	shipped — lifts mean output (k_cat / K_M)
L345G	−1.8	active-site cleft, steric relief	shipped — lifts mean output (k_cat)
C175S	−0.9	disulfide misfold site	structural annotation — not in shipped construct

Structural modelling identified a cysteine at position 175 whose free thiol can form an aberrant intramolecular disulfide during oxidative folding, potentially trapping a fraction of the enzyme in a misfolded state — and crucially, trapping a *variable* fraction, because how much misfolds depends on the redox conditions inside each fruit as it ripens. This site — C175S in our model — is recorded as a structural annotation that informed our understanding of fold stability; it was not incorporated into the final shipped construct. The

shipped double variant S39L+L345G was computationally designed for consistent, predictable catalytic activity, and the tight CV reflects the precision a designed sequence inherits: when every seed carries the same two edits, the enzyme folds the same way in every fruit, and the fruit-to-fruit scatter in output narrows.

That is the deep reason precision and potency come from one design pass rather than two: a flavor written into the sequence is inherited intact by every seed, so the distribution tightens from a wide, hopeful smear into a narrow, dependable line. The mean is lifted by both edits, from the same model, in the same construct — potency and precision as two readings of one act.

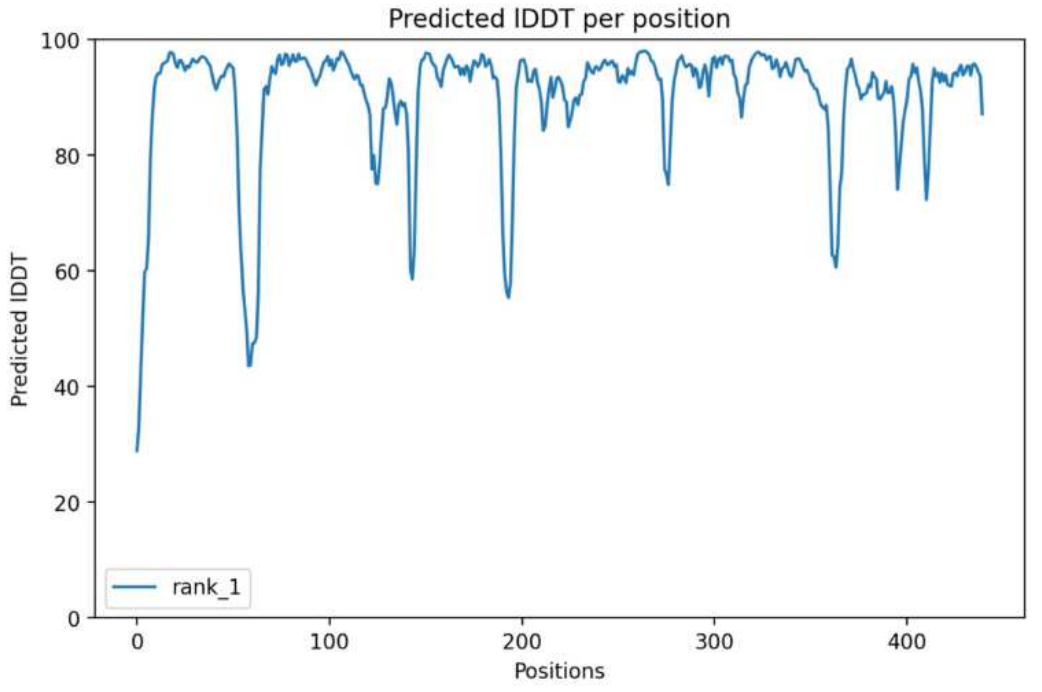

FIG. 22 — The Pun1/AT3 acyltransferase coloured by per-residue confidence. Our structural model identified C175S as an annotation site: a cysteine that could form aberrant disulfide bonds and trap a variable fraction of the enzyme in a misfolded state. This insight informed the fold-stability design of the shipped double variant S39L+L345G.

Here is Phoenix's full eight-axis profile, every figure of which carries the same ±7 percent design discipline the ratio imposes — not one loud axis predicted carefully and seven left unspecified, but eight scores, each designed and each built for reproducibility.

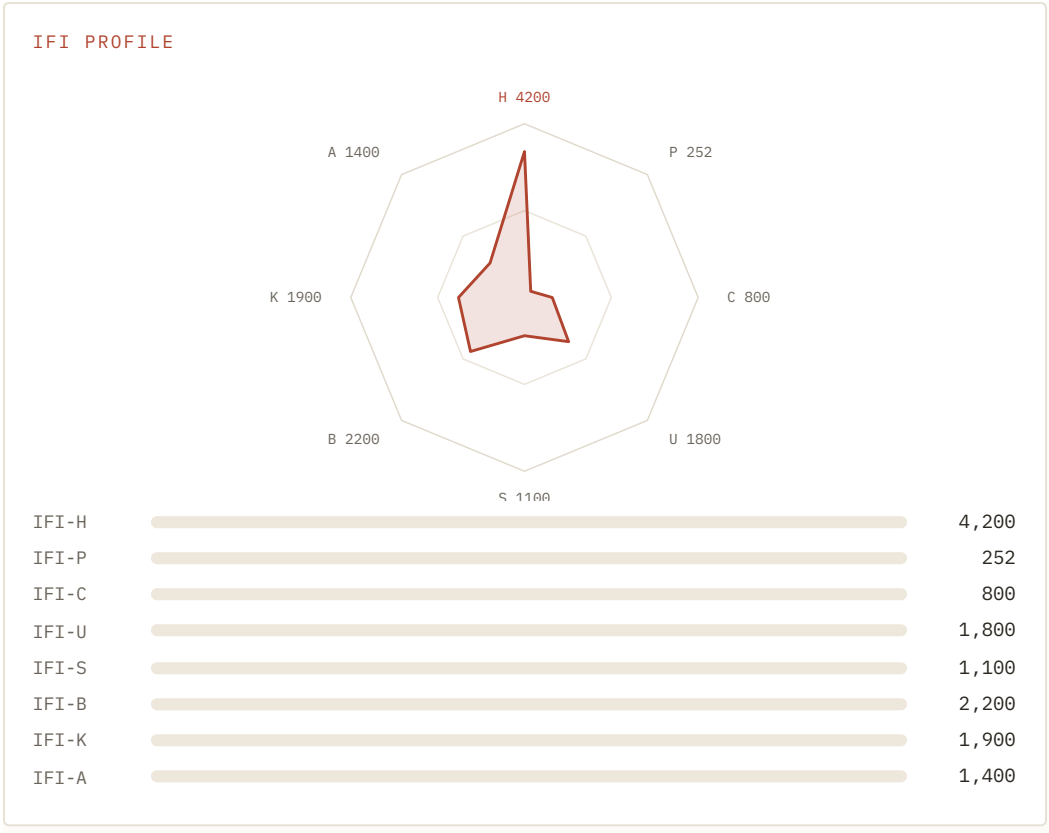

How many fruit it takes to know

Low variance is not only a nicety of the finished product. It changes the cost of *knowing anything about it,* and the arithmetic is worth seeing because it is where a tight CV quietly pays for itself.

To estimate a crop's true mean to within a relative margin E, at a confidence set by the normal multiplier z, the number of fruit you must measure is

> $n = (z \cdot CV / E)^2$.

Sampling error falls as the square root of n, so halving the margin costs four times the fruit — and the whole expression scales with the *square* of the CV. Ask for the mean heat of a wild-type crop to within ±5 percent at 95 percent confidence (z = 1.96): $n = (1.96 \times 0.42 / 0.05)^2 \approx$ **271 fruit.** Ask the same of a designed cultivar at a CV of 7 percent: $n = (1.96 \times 0.07 / 0.05)^2 \approx$ **8 fruit.** The sixfold drop in CV becomes a thirty-sixfold drop in the sampling burden, because the CV enters squared. A claim that takes a whole harvest to verify in the wild type takes a handful of fruit in the designed one — reproducibility is not only cheaper to deliver, it is cheaper to *check,* which is the same thing as cheaper to trust.

Consistency compounds again in the formulator's favor when batches are pooled. Blend N independent batches into a finished lot and the lot's mean inherits a coefficient of variation of $CV / \sqrt{N}$, because averaging independent draws shrinks their spread by the square root of their number. Reaching a lot-to-lot consistency of, say, 3.5 percent therefore asks for $N = (CV / 3.5)^2$ batches pooled — a modest 4 when each batch already runs at 7 percent, against 144 when each wanders at 42 percent. The same square that governs the sampling burden governs the blending burden, and a crop that starts tight is a crop whose finished lots fall effortlessly inside a specification an unruly one could never reach by blending at all.

PROCESS	CV	FRUIT TO PIN MEAN ±5% (95%)	CP AT ±20% WINDOW	OUT-OF-SPEC FRUIT
Wild-type *Capsicum*	42%	271	0.16	~634,000 ppm
Designed cultivar	≤7%	8	0.95	~4,300 ppm

A capable process

Manufacturing has a formal name for a process whose output reliably lands inside the window a buyer was promised: it is **capable,** and it has a number, the capability index

```
Cp = (USL − LSL) / 6σ,
```

the width of the specification window over six standard deviations of the process. Hold a heat specification at a ±20 percent tolerance — loose, as flavor tolerances go — so the full window is 0.40 μ wide, and the index becomes Cp = 0.40 μ / (6 × CV × μ) = 0.40 / (6 · CV). The μ cancels; only the CV remains. The wild type runs at Cp = 0.40 / (6 × 0.42) ≈ **0.16;** the designed cultivar at Cp = 0.40 / (6 × 0.07) ≈ **0.95.** The ratio of the two is 42 / 7 = 6 exactly — the same sixfold, now read as capability.

Translate capability into fruit you would have to discard. The fraction of output falling outside a ±t window of a centered normal process is governed by z = t / CV. For the wild type, z = 0.20 / 0.42 = 0.48, and the tails past ±0.48σ hold about **63 percent** of the fruit — nearly two in three miss a ±20 percent window. For the designed cultivar, z = 0.20 / 0.07 = 2.86, and the tails hold about **0.4 percent** — roughly four thousand fruit per million, against six hundred thousand. We keep it there with the full apparatus of statistical process control: every batch plotted on a control chart, drift caught as a trend long before it becomes a defect, the spread itself treated as a quantity to be engineered down rather than a fact of nature to be endured.

A control chart earns its place only by separating the two kinds of variation beneath it. The irreducible 7 percent is *common cause* — the normal, in-control breathing of a sound process, which is to be left alone. A point beyond the ±3σ control limits, or a run of points drifting steadily one way, is a *special cause* — a bad seed lot, a lamp out of calibration, a changed nutrient line — a real signal to be hunted down and corrected. The discipline is exactly that asymmetry: never chase the common-cause noise, and never miss a special-cause signal. Tampering with a process that is merely showing its common-cause spread only widens it; finding and removing a special cause is how the spread stays where it was engineered to be.

Honesty about the Index requires one more line. A Cp near 0.95 is good; the pharmaceutical standard its methods come from asks for Cp ≥ 1.33, and at a CV of 7 percent a ±20 percent window does not quite reach it — it would take a ±28 percent tolerance, or a CV tighter still, to get there. We hold ourselves to the harder reading rather than the flattering one, because an error bar that only ever told good news would not be worth publishing.

The honest limits

A coefficient of variation is a powerful summary and a partial one, and a standard that asked to be believed rather than checked would not deserve the name. Three limits sit under the 7 percent, and we would rather name them than let them be found.

First, the CV we report is a *within-controlled-conditions* figure. It holds the genetics still and holds the growing environment still, and the 7 percent is what remains. Hand the same seed to an open field with its weather, its soils, its uneven light, and the realized variance will be larger — not because the sequence drifted but because the environment re-entered the sum. The honest claim is precise about its scope: the designed genetics contribute ≤7 percent; total variance in a given setting is that plus whatever the setting adds.

Second, measurement error and process variance shrink by different means, and neither rescues the other. A perfectly uniform crop read on a drifting instrument still yields an untrustworthy score; a beautifully validated assay pointed at a variable crop still describes a moving target. We drive both down at once — assays validated to pin the EC_{50} tight, sequence-written flavor uniform enough to pin the process — but a reader should know they are two separate promises, each of which can fail on its own.

Third, the whole ratio rests on the stability of the reference. The Index cancels the day's drift precisely *because* reference and candidate are read together against the same anchor; if the reference itself wandered between seasons, every score would wander with it and no candidate's error bar would reveal it. So the reference is treated as the most protected number in the system — fixed, archived, re-verified — and the discipline that guards it is the discipline this chapter has been describing, turned on the ruler instead of the thing measured.

None of these are reasons the number cannot be trusted. They are the conditions under which it can — the fine print that makes the headline honest.

One locked number system

There is a last kind of variance that no assay catches, and across a library of this size it is the one that reads as carelessness: **drift between surfaces.** A construct count that is different on one page than another; a "live" figure that wanders between a catalog and a book. So every surface in this library resolves against one fixed set, and the lock is treated as a publishing discipline rather than a housekeeping detail: roughly **225** distinct construct designs backing **19,500** product SKUs; **1,000** live and purchasable today; **42** host organisms; **8** sensory axes; a batch CV of **≤7 percent**; **3** patents filed; and **21** peer-reviewed papers. Product names are catalogue labels for these designs; the 19,500 SKUs do not each represent a unique sequence.

The final two figures matter as much as the first six. The patents establish that the work is real and defensible. The open papers — published under a Creative Commons license, the data laid out for anyone who cares to open it — establish that it is *true:* that the EC_{50}

fits, the structural predictions, the flux models, and the CV claims can be checked, reproduced, and challenged by the field. The error bar is where trust begins, and the open record is where it is kept.

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

It is, in the end, the least glamorous thing our family makes and the one we are quietest and proudest about. To drive a coefficient of variation from 42 percent to 7 is not a technical nicety. It is a kept word, repeated across every single fruit — and a kept word, in food as in anything else, is most of what love looks like once it is done at scale.

Chapter 20 · Reading the World — *score anything.*

We end where a reader can now begin. Everything in this book was built so our family could compose flavor on purpose — grow a tomato that remembers it is a tomato, ladder a heat gene by gene, place a kokumi exactly where a dish wanted one. But the deeper gift of the Index was never the composing. It is that the Index is a *language,* and a language is only finished when it can be spoken about things its makers never touched. A private notation reads only what made it. A standard reads the world.

So here, at the close, is the whole instrument turned outward and handed to you: how to take any food that was ever grown, designed, cooked, or merely encountered on a plate you did not plan — and read it on eight honest axes. Not to replace the pleasure of it. To notice the pleasure exactly, and, if you ever wish to, to design its neighbor. We will go slowly, because reading the world is a skill, and like every skill it is built from a mechanism, a measurement, an arithmetic, and the honesty to know where each one stops.

The eight questions, in their fixed order

Take anything to the tongue and ask the eight questions in the order the mouth itself keeps them. The first three are struck — they are ion channels on the nerve, and they answer fast, the way a bell answers a hammer. The middle three bloom — they are receptors that run a relay inside the cell, and they arrive a beat late, the way a chord arrives after the key is pressed. The last two are felt as much as tasted. Learn what each one is actually doing and the questions stop being poetry and start being physics you can feel.

Where is the **heat?** Capsaicin is an oily molecule; it slips through the cell membrane and settles into a pocket on the inside face of **TRPV1,** a non-selective cation channel riding the trigeminal nerve endings of the tongue. It does not warm the tissue. It lowers the channel's thermal threshold — the temperature at which TRPV1 would open on its own, normally above 43 °C — down toward the warmth of the mouth, so the channel opens and the nerve reports a fire that is not there. Calcium and sodium pour in; the sensation climbs slow and builds, because the lipid has to partition into the membrane before it can bind. That slow build is the fingerprint of heat, borrowed whole from the body's own thermometer.

Where is the **pungency?** This is the sharp cousin, and it works by a faster, cruder chemistry. The electrophiles of mustard and wasabi (allyl isothiocyanate), of crushed garlic (allicin), of cinnamon (cinnamaldehyde) do not merely dock in **TRPA1** — they form covalent bonds with cysteine residues on the channel's inner face. A reactive chemistry, quickly made and quickly undone, which is exactly why pungency climbs the nasal passage and clears within seconds as the adduct resolves. Heat is a lock turned gently; pungency is a wire touched to a terminal.

Where is the **cool-burn?** Menthol binds **TRPM8,** the cold sensor, and shifts its activation threshold *warmer* — so a channel that should open near 25 °C opens at the ordinary temperature of the mouth and reports a cold that no thermometer would find. Cool with no cold in it. Hold the three together and the symmetry is almost beautiful: capsaicin lowers a hot threshold, menthol raises a cold one; they are the same single trick run in opposite directions on two halves of the body's thermometer.

Then the rounded three, the ones that bloom. These are not channels at all but G-protein-coupled receptors, each carrying a Venus-flytrap domain that closes around its molecule and sets off a cascade inside the cell — gustducin, then phospholipase, then a second-messenger relay that ends in the taste cell's own depolarization. They arrive late because a relay takes a beat, and they fill rather than strike.

Where is the **umami?** Glutamate settles into the flytrap of **mGluR1** and the **TAS1R1·TAS1R3** pair, and the 5′-ribonucleotides that gather as food ages — inosinate from cured and simmered things, guanylate from mushrooms — wedge the hinge and hold it shut. That is the old synergy of dashi, why kelp and bonito together are far louder than either alone, and why umami reads below language as nourishment: the receptor is tuned to the signature of protein already broken down and ready for the body.

Where is the **sweetness?** The same architecture, a different pair — **TAS1R2·TAS1R3.** Sugars dock in the flytrap itself; the sweet proteins and glycosides we design with, brazzein and the steviol family, bind elsewhere on the very same receptor and hold it open far longer per molecule, which is how a sweetness can be lifted high with not a grain of sugar behind it. This is the receptor the body has trusted since before it had words, and it is the easiest of all to fool honestly.

Where is the **bitterness** — and, crucially, is it harsh or chosen? **TAS2R38** is one of roughly twenty-five bitter receptors, tuned to a family of alarms: the thiourea note of the compounds that split people into tasters and non-tasters, the glucosinolates of broccoli and kale. Bitterness is the only taste the body ships pre-wired to reject, which is why the meaningful question is never *how much* but *whether you reached for it* — an alarm you flinch from, or a depth you learned to love against the grain of an old warning.

Then the two that are felt as much as tasted. Where is the **kokumi** — the sense of a flavor that fills the whole mouth and stays? **CaSR,** the calcium-sensing receptor, tastes of nothing on its own; it is an amplifier. The γ-glutamyl peptides — glutathione, γ-Glu-Val-Gly — bind it, and it reaches over to the signals around it and lengthens them, rounds them, gives them continuity and body. You do not taste kokumi so much as notice that everything else now lasts and coheres. And where is the **astringency,** the drawing tightness that frames a flavor and resets the palate? This is the one dimension that is not a receptor channel at all. Tannins and polyphenols bind the proline-rich proteins that keep saliva slick and precipitate them out of solution; the lubricating film collapses, the surfaces of the mouth catch against each other, and the brain reads the lost slickness as

tightness. A mechanics, not a nerve chemistry — which is why on our scale this axis is anchored to the salivary proteins themselves and not to any channel.

Eight questions, one food, and what was an inarticulate pleasure becomes a silhouette you can name, compare, and remember.

The ruler underneath — *why the half-maximal point*

Eight pathways, eight readouts — a channel current here, a flash of intracellular calcium there, the turbidity of a protein falling out of saliva for the last. These are as different as the senses that carry them, and it is worth seeing them laid side by side before we unify them:

AXIS	RECEPTOR	WHAT THE HALF-MAXIMAL POINT IS READ ON
IFI-H · Heat	TRPV1	channel current vs. capsaicinoid concentration
IFI-P · Pungency	TRPA1	channel current vs. electrophile concentration
IFI-C · Cool-Burn	TRPM8	channel current vs. cooling-agent concentration
IFI-U · Umami	mGluR1 · TAS1R1·TAS1R3	intracellular Ca^{2+} vs. glutamate concentration
IFI-S · Sweetness	TAS1R2·TAS1R3	intracellular Ca^{2+} vs. sweetener concentration
IFI-B · Zero-Bitter	TAS2R38	intracellular Ca^{2+} vs. bitter-agonist concentration
IFI-K · Kokumi	CaSR	receptor activation vs. γ-glutamyl peptide concentration
IFI-A · Astringency	salivary PRPs	protein precipitation vs. tannin concentration

To read a chili and a cup of tea on one page, those eight different measurements have to collapse into one idea, and they do: not *how much* of a compound a food holds, but how *loudly* that compound speaks to its receptor. Loudness is a half-maximal point — and the reason that point, and no other, carries the whole Index is the same for every row of the table above.

Every receptor answers its stimulus along the same S-shaped curve. Write the response **R** at a stimulus concentration **C** as the Hill relation:

$$R(C) = R_{max} \cdot C^n / (EC_{50}^n + C^n)$$

R_max is the fullest answer the receptor can give; **n** is the Hill coefficient, the steepness of the rise, which for the cooperative TRP channels runs above 1 because the channel's own subunits help one another open. Now set the concentration equal to **EC_{50}** and watch the equation fall open: C^n and EC_{50}^n become the same number, the fraction is exactly one-half, and **R = R_max / 2,** no matter what R_max happens to be and no matter where the baseline sits. The half-maximal concentration is the one point on the whole curve you can locate without first knowing the receptor's ceiling. It is the inflection of the log-dose response, the steepest place on the rise, the point most nearly vertical — and *because* it is steepest there, a small wobble in the measured response barely shifts the concentration you read off it. EC_{50} is the single most reproducible number the entire curve contains. The Index is built on it and nothing else, and that is not an aesthetic choice; it is the only choice that holds its nerve across ten thousand assays.

Invert it and scale, and loudness becomes a score:

$$IFI(c) = 1000 \times EC_{50}(reference) / EC_{50}(candidate)$$

A *low* EC_{50} means the receptor reaches half its answer on very little of the food — a loud food. Putting EC_{50} in the denominator flips loud foods high, where the eye expects them. The thousand is a courtesy to that eye. And the reference is fixed so that a food exactly as loud as the native wild-type anchor reads 1,000: set EC_{50}(candidate) equal to EC_{50}(reference), the ratio is one, times a thousand. Said plainly, the reference reads 1,000 on every axis — because *being the reference* is precisely what that flat profile means:

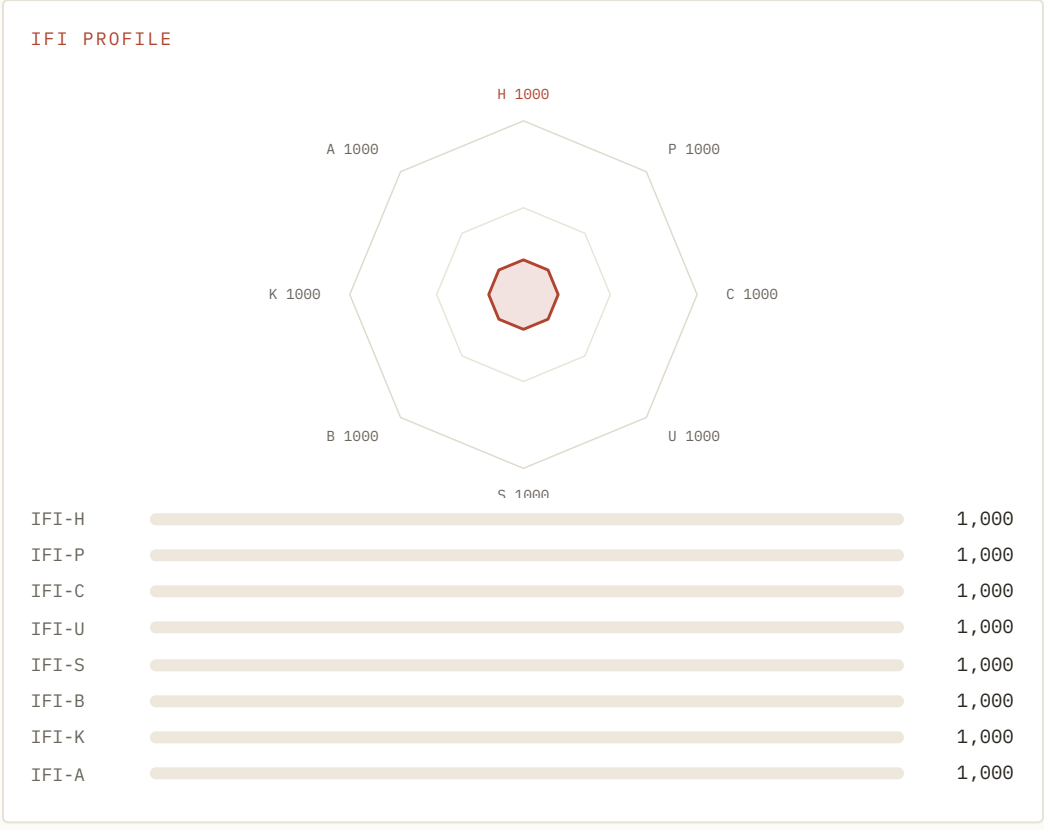

Each axis stands on its own reference compound — heat referenced against capsaicin, pungency against AITC, cool-burn against menthol, umami against glutamate, and so on. The 1,000 sitting on every axis of that flat profile is a separate anchor for each: a score of 1,500 on the heat axis and 1,500 on the umami axis each mean

half again as loud as that axis's own reference — not the same absolute intensity across axes, since the references are different compounds at different concentrations. Cross-axis numbers can be spoken in the same breath and read on the same chart; each one answers only the question of its own receptor.

Every other silhouette in the world is that level line pressed up where a food out-shouts the anchor and down where it falls quiet beneath it. The per-axis scale runs to a ceiling of 5,000 — not a target a food should strive to reach, but a frame drawn wide enough to hold the loudest things we have ever built.

Reading a pepper, in numbers



The half-maximal point, found on the bench — a dilution series read until the current crosses its own midpoint.

Start with the axis everyone knows, and take the loudest thing we designed. For our flagship, **Phoenix** (SS-0096), our model predicts **IFI-H 4,200.** Run it back through the formula and the prediction unpacks as arithmetic:

> $4{,}200 = 1000 \times EC_{50}(\text{ref}) / EC_{50}(\text{Phoenix}) \rightarrow EC_{50}(\text{ref}) / EC_{50}(\text{Phoenix}) = 4.2$

TRPV1 reaches its half-maximal current on **4.2 times less** Phoenix extract than it needs of the reference. The fruit speaks 4.2 times as loudly. The Index asks for nothing but that ratio — the absolute anchor is a convenience, never a dependency — but to make the bench concrete: set the heat reference at an illustrative 2.1 µM capsaicin-equivalent, and Phoenix half-maxes at 2.1 ÷ 4.2 = **0.50 µM,** a value you could walk toward tube by tube down a dilution series and watch the current rise to cross its own middle.

Now the older scale, and an honest word about it. SHU — the Scoville number — measures the *chemistry,* the capsaicinoid content, not the loudness, and the two track only loosely. Our model predicts approximately **2,937,948 SHU** for Phoenix, corresponding to our predicted IFI-H 4,200; divide and you get roughly **700 SHU per IFI-H point** at the top of the ladder. But that conversion drifts as you walk down, and the drift is not sloppiness — it is the physics of the sigmoid showing through:

CULTIVAR	SKU	SHU†	IFI-H†	SHU PER IFI-H POINT
Ember	SS-0087	866,000	1,570	552
Blaze	SS-0091	1,720,000	2,630	654
Firestorm	SS-0094	2,470,000	3,490	708
Phoenix	SS-0096	2,937,948	4,200	700

† Computational predictions.

Low on the ladder, a pepper buys an IFI-H point for about 550 SHU; near the ceiling the same point costs about 700. The price of loudness rises *with* loudness. That is exactly what a saturating curve predicts: as concentration climbs toward R_max the rise flattens, and each additional unit of chemistry returns a little less perceived loudness than the unit before it. SHU is very nearly linear in capsaicinoid content; the tongue is not linear in anything; the Index follows the tongue. It is why we never report heat as chemistry alone — the chemistry and the sensation stopped agreeing somewhere around a million Scoville, and the eater lives on the far side of that disagreement.

Phoenix earns its closing place in this book for a second reason: it is the one cultivar we built to register on all eight axes at once, which makes it the clearest thing in the whole library to practice reading. Here is its full silhouette:

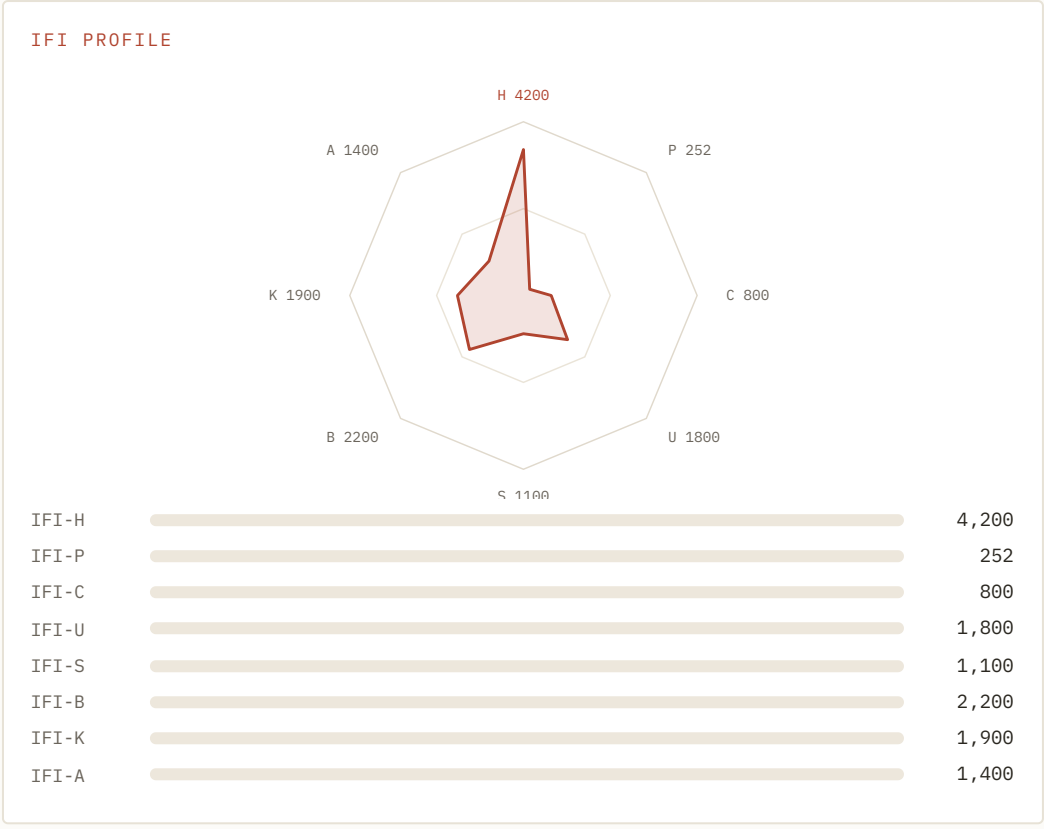

Say the numbers aloud in their fixed order and the shape assembles itself: heat towering as designed, but pungency and cool-burn present beneath it, an umami and a frank sweetness holding the middle, a real bitterness and a tall kokumi giving it body, astringency framing the edge. A full chord struck at once — not a single loud note but a composition that happens to be led by fire. Reading it is nothing more than naming each loudness in turn and letting the eight of them resolve into one figure.

Reading what we never designed



A tomato, in cross-section — the umami a truck spends, and the glutamate route that pays it back.

The proof that the Index is a standard and not a private notation is that it reads what we had no hand in — a wild blueberry, a dish a grandmother made, a tomato off a truck — the very same way it reads a cultivar grown to a number. One ruler, eight axes, the half-maximal point underneath them all. Take the tomato, because its story is arithmetic you can follow end to end.

A tomato warm from the vine is a two-axis thing: umami leading, a clean sweetness beside it, almost nothing else — which is why it reads as pure summer and demands no analysis to love. But between the vine and the supermarket shelf, roughly **93%** of its free glutamate is spent — on ripening in the dark, on the days and the miles. The umami that stood tall on the vine now reads **IFI-U 300:** the same silhouette with the life drained out of it. We designed a tomato to restore it, through the plant's own native glutamate route — our model predicts approximately **IFI-U 2,600.** Set the two readings against each other:

```
2,600 ÷ 300 = 8.67
```

Our model predicts the restored fruit would be **8.67 times as loud** on umami as the one the truck delivers — or, read from the other side, the supermarket tomato would have to carry nearly nine times the glutamate it has left to speak as loudly, which a drained fruit simply cannot. And now notice the thing the numbers quietly refuse to do. A 93% loss of glutamate did *not* become a 93% loss of loudness. It couldn't, because loudness is a half-maximal point on a sigmoid, not a tally of molecules — lose most of your glutamate and you slide down the steep middle of the curve, where loudness falls faster than concentration near the bottom and slower near the top, so the two numbers never march in step. Concentration and loudness are different quantities. Measuring the second instead of the first is the entire reason the Index exists.

You do not need a chamber for most of this. The instrument taught your tongue what to feel for, and the feeling was always there. Here are the silhouettes of a handful of ordinary things — each one a shape you can now draw, marked loud (●), present (◐), a trace (○), or silent (·) on every axis in turn:

FOOD	H	P	C	U	S	B	K	A
Long-simmered broth	·	·	·	●	·	·	●	·
Strong green tea	·	·	·	○	·	○	·	●
Fine dark chocolate	·	·	·	·	◐	◐	◐	○
Vine tomato	·	·	·	●	◐	·	·	·
Ripe peach	·	·	·	·	●	·	·	○
Black coffee	·	·	·	·	·	●	◐	○
Medjool date	·	·	·	·	●	·	●	·
Raw radish	·	●	·	·	·	◐	·	·
Aged cheese	·	·	·	●	·	·	●	○

Read a few aloud and feel them resolve. A long-simmered broth is loud on umami and kokumi and quiet nearly everywhere else — a two-axis savory depth, which is why it comforts without ever demanding attention. Strong green tea is astringency framed by a thread of umami and a trace of bitterness, a tight bright figure of two-and-a-half axes. Fine dark chocolate is bitterness eased toward balance with a sweetness lifted to meet it, kokumi body beneath and a faint astringency at the edges — nearly a composition. Black coffee is bitterness tall over a real kokumi body that a good cup has and a thin one lacks, with a trace of astringency from the darker roasts: the silhouette of something the drinker learned to love against the grain of an old warning. A Medjool date is sweetness and a deep kokumi richness, a fruit that tastes almost of caramel without a grain of added sugar. A raw radish is a fast pungency up front and a clean bitterness behind it — the garden's own aperitif. A wedge of aged cheese is umami and kokumi standing tall together with a faint tang at the edge: the savory-depth chord made solid, which is exactly why a rind dropped into a simmering pot does the work of an hour.

And here is the quiet reward of it: once you can draw a shape, you can imagine its neighbors. The date with its sweetness eased and its richness kept. The coffee with its harsh edge pressed down and its body left standing. The peach with a thread of warmth laid in where there was none. To read the world on eight axes is already, without meaning to, to have begun composing it.

The error bar is part of the reading

A reading without its error bar is a rumor. Every number in this book carries one, and the honest width of that bar is most of what separates a standard from a guess. The width is the coefficient of variation — the standard deviation of a measurement divided by its mean, CV = σ / μ — and it is the number we are proudest of, because it is the hardest-won.

Our engineered cultivars hold a **batch-to-batch CV of ≤7%.** Wild material runs to **42%.** Watch what that does to a reading. Take Phoenix at **IFI-H 4,200.** At a 7% CV its standard deviation is 0.07 × 4,200 = **294.** Roughly two batches in three land within one standard deviation — between **3,906 and 4,494** — and about ninety-five in a hundred land within two, from **3,612 to 4,788.** The reading is a band, and a narrow one; "4,200" is a place the plant returns to.

Now the wild case, the one every breeding program before ours lived inside. A wild pepper reading **1,000** at a 42% CV has a standard deviation of **420.** Its one-sigma band runs from **580 to 1,420;** its ninety-five-percent band from roughly **160 to 1,840.** The error very nearly swallows the signal — the same plant, picked twice, can read almost half or almost double, and no one picking it could say which was the truth. You cannot design against a number like that, and you cannot honestly sell it as a number at all. The ≤7% we hold is not a flourish on the spec sheet; it is the precondition that lets a reading mean anything in the first place. It is the whole reason the chamber comes before the claim, and it is why, when we write a profile down, we are writing down a thing the plant can be trusted to do again.

Where a reading bends — *the honest limits*

The Index is a ruler, and like any ruler it has edges where it bends. We would far rather name them than let a reader trust a number one step past its warrant — because a measurement standard earns its authority precisely by being candid about where it stops.

Desensitization. TRPV1, held open by capsaicin, eventually desensitizes — the calcium it admits triggers the channel's own stand-down. Taste a superhot twice and the second reading falls, not because the fruit changed but because the receptor did. We read against a rested palate for exactly this reason; the first honest taste is the measured one, and the mouth is a single-use instrument on the hottest axis.

Temperature. Heat and cool-burn ride thermoreceptors — TRPV1 answers warmth, TRPM8 answers cold — so the serving temperature is not a footnote to the reading, it is part of it. A broth read steaming and the same broth read cool are two different readings on the thermal channels. We either fix the temperature or we report two numbers and say which is which.

The matrix. Capsaicin is lipophilic; fat takes it up and holds it, lowering the free concentration the receptor ever actually sees. A chili suspended in cream reads quieter than the same chili in water — true to the eating, and a trap for anyone who measures the extract and forgets the dish it will live in. The Index reads a molecule's loudness; the kitchen reads the molecule in company.

Central masking. The Index measures loudness at the receptor, one axis at a time; the brain, downstream, integrates all eight into a single percept. A tall sweetness suppresses the *perception* of bitterness well after both receptors have had their say, so a food can read genuinely bitter on the instrument and taste balanced in the mouth. The reading is the receptor's truth, not the whole experience — a feature, not a flaw, so long as you always remember which of the two you are holding.

The ceiling. The scale frames at 5,000, and near that ceiling the differences compress. A handful of engineered extremes crowd the top of an axis where the dose-response has nearly flattened, and the frame grows honest about saying "very loud" while turning coy about splitting hairs above it. That is the curve's doing, not the Index's evasion.

None of these are failures of the standard. They are the conditions it is honest about, and a number that pretended to know none of them would be the one number here to distrust.

What the number is for



The greenhouse at dusk — the oldest intimacy we have, measured at last, and given back.

We will say once more what the Index is *not,* because at the very end it matters most. It is not a verdict on whether a food is good — a broth loud on two axes is not lesser than a fruit loud on six, a quiet flavor can be exactly the right one, and the best meal of your life might score modestly and mean everything. It is not a target that every food should strive to max out; the 5,000 ceiling is a limit, not a goal, and most of what we design sits far below it, on purpose and with pleasure. And it is not a substitute for the tongue, which remains the only instrument that finally matters, the one every number in this book is answerable to. The Index is an invitation — to notice more, to taste analytically without tasting coldly, and, if you are in the business of growing food, to compose on purpose what was once left to chance. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

It bears repeating, at the close, that we did not invent the eight axes; we found them. The tongue drew the map long before anyone labeled it — eight pathways it had always kept more or less separate, eight questions it had always been quietly asking of every bite, a channel for fire and a channel for cold and a flytrap for the shape of nourishment. All the Index did was listen carefully enough to write them down, and attach to each a number honest enough to design against. In that sense the standard is less an imposition on taste than a transcription of it: the oldest sense of the body, given at last a notation of its own. We are stenographers of the tongue, and the tongue was the author all along.

We began a whole volume ago with a tomato warm from the vine and a sweetness no one could measure — the oldest intimacy we have, carried by molecules into receptors, felt exactly and described only in poetry. We have spent the volume turning that intimacy into eight honest numbers, and we want to be clear, at the last, about what the numbers are *for.* They are not a replacement for the pleasure. The split tomato is still the split tomato; the first warmth of heat still arrives like a held note; the broth still tastes like being cared for. The numbers are so that we can give those things back on

purpose — grow a tomato that remembers it is a tomato, ease a bitterness that kept someone from a food that would have nourished them, compose a flavor and a nourishment as a single deliberate act and know, to the error bar, that it will arrive as intended. That is the whole of it. We measured taste because we love it, and because the thing you love you want to make well, and keep, and give to the people at your table. The Index is how a family learned to do that on purpose. Eight axes. Every flavor. All at once — and all of it meant, from the first number to the last, to feed you right.