

THE SCOVILLE SPLICE LIBRARY · VOLUME LVIII

The Receptor Book

The tongue does not taste. It reads. *Eight proteins, eight words.*

The eight receptors that turn a molecule into a meaning. A field guide to the machinery of taste — from the pore of a channel to the dinner table, and to how we learned to design the molecule that speaks to each one.

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

Contents

01 Opening essay · The tongue is not a sensor

02 Chapter 1 · The Lock and the Key — what a receptor actually is.

03 Chapter 2 · The Half-Maximal Point — the mathematics of loudness.

04 Chapter 3 · Two Ways to Open a Channel — fast fire, slow savor.

05 Chapter 4 · TRPV1 — the heat that isn't there.

06 Chapter 5 · TAS1R2·TAS1R3 — the sweet clamshell.

07 Chapter 6 · TAS1R1·TAS1R3 and the taste mGluRs — the savory depth.

08 Chapter 7 · TAS2R38 — the bitter sentinel.

09 Chapter 8 · CaSR — the invisible richness.

10 Chapter 9 · TRPM8 — cold without cold.

11 Chapter 10 · TRPA1 — the sharp channel.

12 Chapter 11 · The Salivary PRPs — the sense that isn't a taste.

13 Chapter 12 · The Convergence — where eight streams become one flavor.

14 Chapter 13 · When Receptors Talk — synergy, suppression, and the multi-axis read.

15 Chapter 14 · The Inversion — reading the receptor to design the plant.

16 Chapter 15 · Anatomy of a Single Burn — Phoenix, traced receptor by receptor.

17 Chapter 16 · The Universal Reader — one tongue, every plant.

Opening essay · The tongue is not a sensor

There is a moment at our table — you know it — when a spoonful is lifted, and before anyone speaks, a face changes. A softening around the eyes. The particular stillness of a person who has just been told something true by a bowl of soup. For most of human history that moment was the end of the road: the flavor arrived, the body answered, and no one could say what had passed between them. We could grow the plant, cook the fruit, keep the seed of the one that moved us. We could feel the result with our whole bodies. We could describe it only in poetry.

This book is about the eight proteins standing in that gap.

Taste is not a thing that happens *to* the tongue. It is a thing the tongue *does* — an act of reading, performed billions of times a second by a small set of molecular machines sitting in the membranes of cells you will never see. A warm tomato is, to the physicist, a cloud of glutamate and sugars and acids drifting toward a wet surface. It becomes *flavor* only when those molecules settle into the folded pockets of specific receptor proteins and change their shape, and that shape-change is relayed, as a pulse of ions and a flicker of a nerve, to a brain waiting upstream. The molecule is the word. The receptor is the reader. Meaning is what happens in between.



The founders in the research greenhouse before the lights come up — where the question that became this book was first asked over a plate of food.

We did not set out to write a book about receptors. We set out, as a family, to grow food that tastes the way we remember food tasting — fuller, rounder, more alive — and to grow it so the body could draw the most from it. But you cannot design what you cannot read, and you cannot read a flavor until you know the eight readers by name. So we learned them, one protein at a time, at the bench and at the table both. This book is the field guide we wish we had had: each reader a small marvel of protein engineering that evolution arrived at long before we did, each one a lock we studied until we could forge its key — or, more often and more beautifully, teach a plant to grow one.

> The tongue does not taste the world. *It reads it* — and we have learned to write.

What a receptor actually is

Begin with the thing itself, because everything else is downstream of it. A receptor is a single chain of amino acids — a few hundred links of a twenty-letter alphabet — that folds, in water, into one particular three-dimensional shape and holds it. The folding is not decoration. It is the function. Somewhere in that fold the chain leaves a cavity, a pocket of a specific size and charge and shape, and that pocket is the whole point: it is a negative-space portrait of the molecule the protein was built to recognize, carved in protein the way a key's bitting is carved in brass.

The old textbook calls this *lock and key*, and the image is lovely, but it is slightly wrong in a way worth fixing. A good binding pocket is not rigid. When the right molecule drifts in, the protein closes around it — an *induced fit* — the way a hand closes around a stone it has decided to pick up. The binding is a settling of the whole structure into a lower-energy arrangement, and the currency of that settling is free energy, measured in kilocalories per mole. At body temperature, 310 kelvin, the thermal energy scale RT is about 0.62 kcal·mol^{-1}, and the arithmetic that governs affinity is almost cruelly simple: every tenfold tightening of binding costs

$$\Delta\Delta G = RT \cdot \ln(10) \approx 1.42 \text{ kcal·mol}^{-1}$$

— roughly the energy of a single well-placed hydrogen bond. One good bond, correctly aimed, and a molecule binds ten times more readily; the receptor speaks ten times louder of it. This is the lever the whole platform pulls. When we designed the reference superhot at the heart of our ladder, the two enzyme edits that made it — **S39L** and **L345G**, each written into the capsaicin-building acyltransferase, not the receptor — were scored in exactly these units, with calculated effects of **−2.1** and **−1.8 kcal·mol^{-1}**. Protein behavior is priced in kcal·mol^{-1} whether the protein makes the word or reads it; the arithmetic does not care which side of the conversation it sits on.

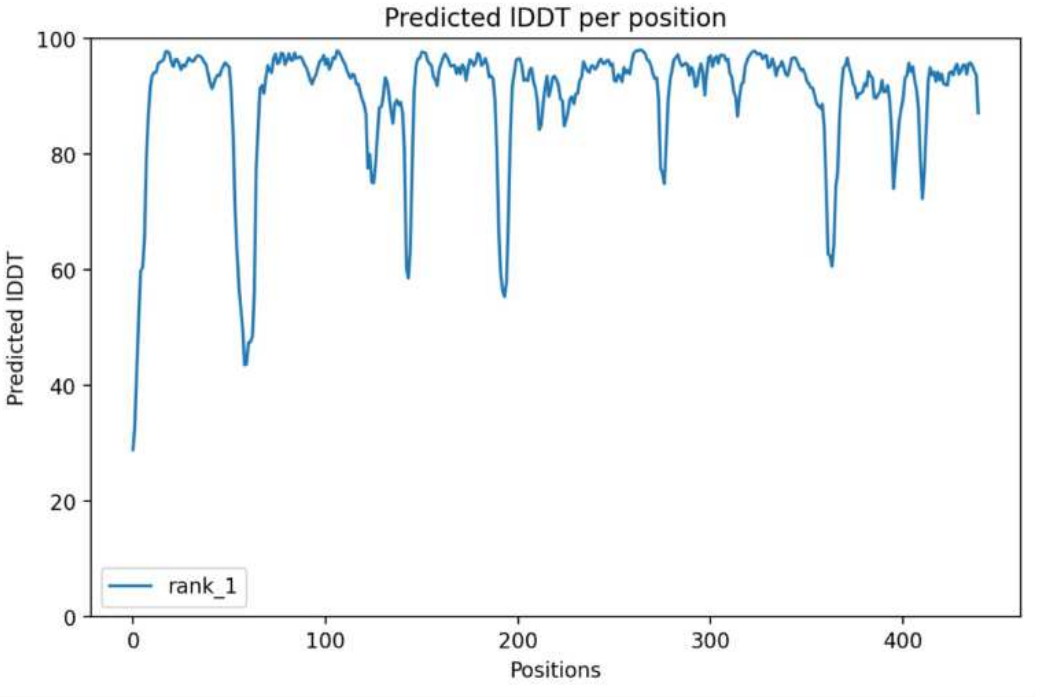

The acyltransferase that writes capsaicin, coloured by model confidence — bright where the fold is certain, pale where it is not. Every receptor in this book was read the same way before we touched a plant.

How do we *read* a fold we have never crystallized? We predict it, residue by residue, and we insist the prediction tell us how sure it is. The figure above is a protein coloured not by beauty but by confidence — a per-residue certainty score — bright through the active-site cleft where the structure is all but settled, dimming toward the floppy tails where it is a guess. We trust the bright and we distrust the pale, and we never design against a pocket the model is not sure exists. That discipline — read the structure, believe only the confident part, confirm at the bench — is the method of the entire book. It is how a family of growers came to speak fluently about proteins none of us will ever see with our eyes.

Two ways to carry the word inward

Once a molecule is bound, the receptor has to *say so* — it has to turn a binding event into a signal the nervous system can hear. Here the eight readers split, cleanly, into two families that do this in deeply different ways. The division is not a taxonomy we imposed; it is written into the proteins.

The first family are the **G-protein-coupled receptors** — the metabotropic readers, four of our eight: sweetness, umami, bitterness, and kokumi. These are relay stations, not doors. Sweet and umami and kokumi belong to the stately *class C* GPCRs, each carrying an enormous clamshell on its outside — a *Venus-flytrap domain* — that closes around its ligand the way the plant closes around a fly. Sweetness is read by a pair, TAS1R2 bound to TAS1R3; umami by its cousins, the metabotropic glutamate receptor mGluR1 and the TAS1R1·TAS1R3 pair, whose flytrap shuts around a single molecule of glutamate — the savor of a ripe tomato, caught. Kokumi, the quietest axis, is read by the calcium-sensing receptor CaSR, which γ-glutamyl peptides do not so much switch on as *lean on*, tuning it allosterically, lending a dish its continuity and weight. Bitterness is the outlier of the family: its TAS2R receptors, TAS2R38 the famous one, are small and class-A-like, binding their alarms down inside the transmembrane bundle rather than in any external clamshell.

What unites them is what happens next. Binding nudges the receptor, the receptor wakes its G-protein (in taste cells, α-gustducin and its βγ partners), and from there a cascade unfolds: phospholipase C-β2 is switched on, inositol trisphosphate floods out, calcium is released from internal stores, a channel called TRPM5 opens on the flood, the cell depolarizes, a gate named CALHM1 swings wide, and ATP pours out as the actual neurotransmitter onto the waiting nerve. Six steps, and every step is a chance to multiply — one bound molecule can free thousands of calcium ions. This is the great gift and the great cost of the metabotropic design: enormous *gain*, and a built-in slowness and adaptation. Sweetness blooms and fades; it is amplified, interpreted, and a little laggy, because it has been passed hand to hand down a long bucket line before it ever reaches the brain.

The second family are the **ion channels** — the ionotropic readers, the TRP channels, three of our eight: heat, cool-burn, and pungency. These are not relays. They are doors. The molecule binds, and the door *opens*, and ions pour through the protein itself, directly, no

cascade, no bucket line. Capsaicin finds its vanilloid pocket on TRPV1 and does not warm a single cell — it simply lowers the channel's thermal threshold from above 43 °C down into the range of the mouth, so the body, obeying its own honest machinery, reports a fire that is not there. Menthol finds TRPM8 and drops its cold threshold likewise, so a cool that is not there blooms across the tongue; the channel's true job is to open below about 26 °C, and menthol merely moves the number. TRPA1, the pungency channel, is the strangest: its sharpest agonists — the allicin of crushed garlic, the isothiocyanate of mustard and wasabi, the aldehyde of cinnamon — do not politely dock. They react. They form a covalent bond to specific cysteine residues on the channel's inner face, a chemical handshake that cannot be undone until the protein is rebuilt. That is why wasabi's hit is so sudden and so total: the door has been chemically bolted open.

These three are fast, sharp, undiluted — the word delivered straight. And they do not even live on the taste buds proper; they stud the free endings of the trigeminal nerve, the general feeling-sense of the face, which is why heat and cool and pungency are felt as much as tasted, why they have edges and locations and arcs in time that sweetness never has.

And then there is the one reader that is not a receptor at all. **Astringency** — the drying pucker of strong tea, of young wine, of a tannic pomegranate — is read not by a protein in a membrane but by proteins loose in your own saliva: the proline-rich proteins, whose job is lubrication. Tannins and polyphenols seek them out, bind them, and *precipitate* — the two tangle and fall out of solution together — and the mouth, suddenly stripped of its slick, reports the loss as friction through ordinary mechanical touch. Astringency is the feeling of your own saliva being taken away. It is the honest edge of this book, because its transduction is the least settled of the eight, the one place where we are reading a sense whose wiring is still partly argued over in the literature. We score it anyway, carefully, and we say plainly where the ground is soft.

A word on the sense we leave out. Sourness — the proton sense, read by the OTOP1 channel — is real, beautiful, and precisely the kind of thing the tongue does superbly. We treat it as a cousin rather than a design axis, because while the mouth reads acidity wonderfully, we rarely set out to *write* it; a plant's acid is mostly the business of its ripening, not ours to compose. Eight is not a round number chosen for elegance. It is the count of distinct molecular reads we can, today, both measure and design against.

The eight readers, by name

Here, then, is the roster — the proteins this book spends itself on, each with its axis code, its molecular class, the way it carries its word inward, and the native stimulus it was built to hear.

AXIS	RECEPTOR	PROTEIN CLASS	HOW IT CARRIES THE WORD	NATIVE WORD
IFI-H · Heat	TRPV1	TRP ion channel	Ionotropic — the door opens	capsaicin; heat >43 °C; protons
IFI-S · Sweetness	TAS1R2·TAS1R3	Class C GPCR (Venus-flytrap)	Metabotropic — relay with gain	sugars; sweet proteins
IFI-U · Umami	mGluR1 · TAS1R1·TAS1R3	Class C GPCR	Metabotropic	glutamate
IFI-B · Zero-Bitter	TAS2R38	T2R GPCR (in-bundle pocket)	Metabotropic	PROP/PTC; bitter alkaloids
IFI-K · Kokumi	CaSR	Class C GPCR	Metabotropic (allosteric)	γ-glutamyl peptides; Ca^{2+}
IFI-C · Cool-Burn	TRPM8 (·TRPV1)	TRP ion channel	Ionotropic	menthol; cold <26 °C
IFI-P · Pungency	TRPA1	TRP ion channel	Ionotropic (covalent)	allicin; AITC; cinnamaldehyde
IFI-A · Astringency	Salivary PRPs	Secreted protein — not a receptor	Precipitation → mechanical touch	tannins; polyphenols

Everything anyone has ever tasted was spelled with these eight letters. To put the index's own language against this roster, here is the shape of the *reference* — the native wild-type baseline against which every candidate is assessed, the silhouette where, by construction, each axis reads exactly its anchor value of one thousand:

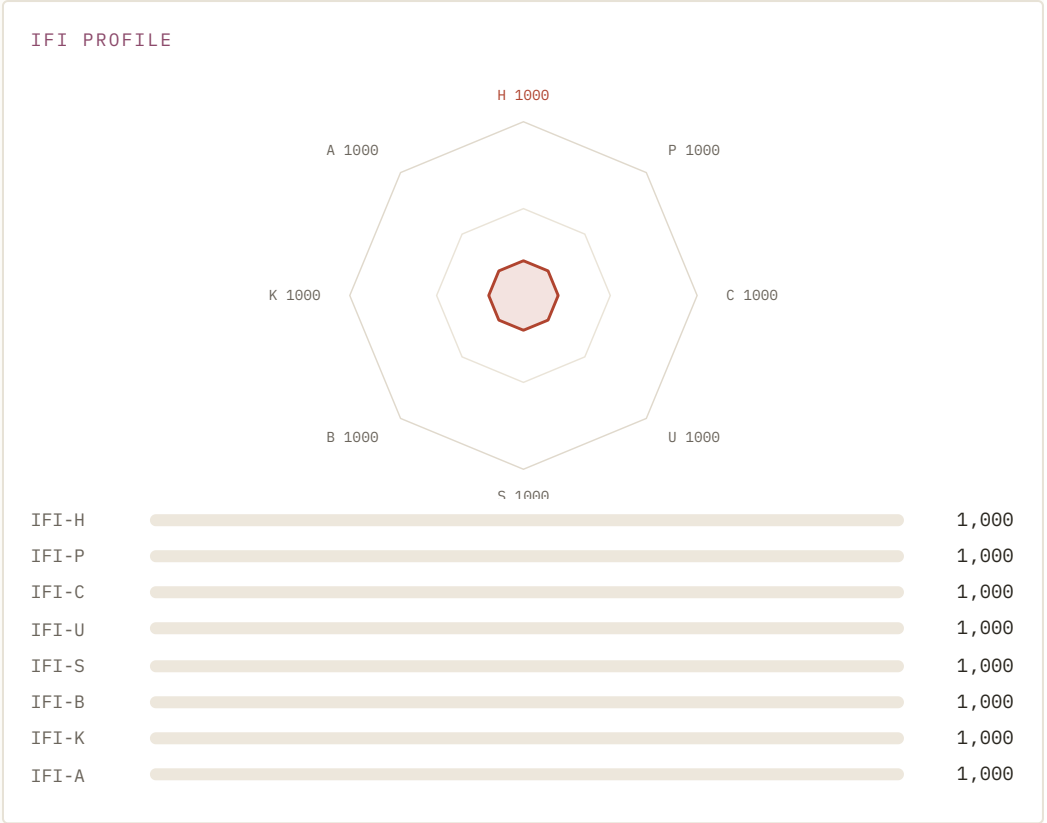

Eight ones-thousand. This is the origin of the whole coordinate system — not zero, because a reader is never silent, but *unity*: the loudness of the plant as it came to us, before we wrote a single word back.

Putting a number on how loudly a molecule speaks

Reverence is where this work begins, but reverence cannot be designed against. To design, we need a number — and the number has to measure the right thing. Not how much of a compound a food

holds, which is just chemistry, but how loudly that compound *speaks* to its reader, which is biology. The whole of the Iyengar Flavor Index rests on one quantity that captures exactly this: the half-maximal effective concentration, EC_{50}, the concentration of a stimulus at which its receptor is driven to half of its full response.

Why half? Because the dose-response curve of a receptor is not a line; it is a gentle S, and the steepest, most informative point of that S — the place where the reader is most sensitive to a change in dose — is the middle. The curve follows the Hill relation,

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

where [L] is the stimulus concentration and *n* is the Hill coefficient, the cooperativity — for capsaicin at TRPV1, a tetramer whose four subunits help one another open, *n* sits near 2. Set [L] equal to EC_{50} and the equation collapses to one-half, every time, for every reader. EC_{50} is the one point on every receptor's curve that means the same thing. It is the natural unit of loudness.

A low EC_{50} means a reader reaches half-response at a whisper of stimulus — the molecule speaks loudly. A high EC_{50} means it takes a shout. Loudness, then, runs *opposite* to EC_{50}, and so the Index inverts it, anchors it to a fixed reference, and scales by a thousand so the ordinary range lands in clean hundreds and thousands:

$$IFI(candidate) = 1000 \times EC_{50}(reference) / EC_{50}(candidate)$$

That is the entire instrument. Three moves — invert, reference, scale — and a tomato, a chili, and a cup of tea can be read against one another on a single page. When candidate and reference are the same plant, the ratio is one and the score is exactly 1,000; that is the anchor silhouette above. Everything else is a reading relative to it.

A worked reading. Suppose a reference pepper's extract half-activates TRPV1 at EC_{50} = 10 µM of its capsaicin-equivalent. Our full-spectrum flagship, Phoenix, is designed to carry an IFI-H of 4,200. Running the formula backward for its EC_{50}:

$$EC_{50}(Phoenix) = EC_{50}(ref) \times 1000 / IFI = 10\ \mu M \times 1000 / 4200 = 2.38\ \mu M$$

Phoenix is designed to reach the receptor's half-point at a little under a quarter of the reference concentration — it is computed to speak heat **4.2× louder**, exactly as its score advertises. Read the other direction and the same lever restores a flavor rather than amplifying one: a pale supermarket tomato scores roughly IFI-U 300, while a tomato with its savor grown back in reaches 2,600. Their loudness ratio is simply the ratio of scores,

$$2600 / 300 \approx 8.7\times$$

— the restored fruit speaks umami nearly nine times louder, which, in EC_{50} terms, means it half-saturates the glutamate reader at one-ninth the concentration the faded one needs. That is the difference,

at the receptor, between a tomato that tastes like water and a tomato that tastes like August.

Why loudness compresses. The Hill curve's saturation has a consequence that surprises people: past a point, piling on more stimulus barely raises the reported loudness. With n = 2, the arithmetic is stark.

STIMULUS [L] / EC_{50}	FRACTION OF MAXIMUM RESPONSE
0.25	0.06
0.5	0.20
1	0.50
2	0.80
4	0.94
8	0.99

Doubling the dose at the EC_{50} carries a reader from half to four-fifths; doubling it again only reaches 94%; a full eightfold over-dose touches 99% and can go no higher. The door is only so wide. This is why, across our superhot ladder, the Scoville Heat Units climb far faster than the IFI-H score. SHU measures capsaicin *present*; IFI-H measures heat *heard*, and the reader saturates.

CULTIVAR · SKU	SHU (PREDICTED)	IFI-H	AXES	SHU PER IFI POINT
Ember · SS-0087	866,000	1,570	2/8	552
Blaze · SS-0091	1,720,000	2,630	4/8	654
Firestorm · SS-0094	2,470,000	3,490	6/8	708
Phoenix · SS-0096	~2,950,000	4,200	8/8	700

Watch the final column: it takes 552 SHU to buy a point of heard heat near the bottom of the ladder and about 700 near the top. Each additional unit of capsaicin returns less loudness than the one before, precisely because the receptor is filling. Fit a power law to these four rungs and the relationship is tight — IFI-H ≈ 0.031 × $SHU^{0.79}$, holding to an R^2 near 0.997 — and the exponent, **0.79**, is the number that matters: it sits *below one*, and an exponent below one is the mathematical fingerprint of a saturating reader. We offer that fit as a description of this ladder, not a law of nature; each cultivar is its own genome, and the honest claim is the mechanism — a filling door — not the decimal.

A note on units. SHU is itself a dilution number: pure capsaicin is defined at 16,000,000 SHU, so a predicted SHU of approximately 2.95 million for Phoenix describes a heat our model places at roughly one-fifth the intensity of pure capsaicin (16,000,000 / 2,950,000 ≈ 5.4). The Index does not replace SHU; it finishes the thought SHU began, converting *how much is there* into *how loudly it is heard*. On the heat axis alone, that stretches from a shishito's sunlit IFI-H near 40 to Phoenix at 4,200 — a hundredfold span in loudness — and the scorecard normalization caps each axis at 5,000, so even our loudest flagship sits at 84% of the ceiling, with deliberate headroom left above it.

What the number does not say

A measurement standard earns trust by being honest about its own edges, so here are ours, plainly.

First, **reproducibility**, because a number that wobbles is not a number. The quantity the Index reads back is only as good as the plant's batch-to-batch consistency, and this is where engineered flavor parts ways with harvested flavor. Our reference superhot is designed to hold its output to a coefficient of variation — the standard deviation as a fraction of the mean — of **≤7%**, against **42%** for the wild-type it replaced. In plain terms: an IFI-H of 4,200 at 7% CV carries a batch spread of roughly ±290 points, a band of about 3,900 to 4,500; at 42% that same score would swing by ±1,760, from under 2,500 to over 5,900. One is a specification you can formulate against; the other is a rumor. And there is a quiet grace in the saturation we just traced: because loudness compresses as roughly the 0.79 power of concentration, a 7% wobble in delivered capsaicin propagates to only about 0.79 × 7% ≈ **5.5%** wobble in the heard score — the derivative $d(\ln \mathrm{IFI}) = 0.79 \cdot d(\ln \mathrm{SHU})$ does the arithmetic for us. The same filling door that compresses the scale also damps the noise. Consistency is not a boast here; it is the precondition for the number meaning anything at all.

Second, **EC_{50} is measured on an isolated reader, and your mouth is not a cuvette.** The clean half-maximal number comes from a receptor studied more or less alone; the real tongue is a crowded, warm, moving, salivating place where pH drifts, temperature swings, fats sequester the very molecules we are trying to measure, and the eight readers audibly talk over one another. Sweetness quiets bitterness; a little heat sharpens the perception of savor; astringency and bitterness are forever being mistaken for each other. The Index scores each axis as cleanly as the biology allows and does not pretend the eight are independent. They are not. They are a chord, and a chord is more than its notes.

Third, **EC_{50} is a half-point, not a ceiling.** A high score tells you a plant speaks an axis loudly at the sensitive middle of the curve; it does not promise the whole supra-threshold experience, the arc in time, the place on the tongue, the aftertaste. Those live beyond what any single concentration can capture, and we are candid that the number is a *reading*, not a verdict — the first true sentence about a flavor, not the last word on it.

And fourth, the soft ground already named: **astringency**, whose precipitation-and-friction mechanism is the least pinned-down of the eight. We score IFI-A, we find it useful, and we flag it as the axis most likely to be revised as the science firms up.

With every one of those caveats standing, here is what the instrument can do that nothing before it could — read a single plant across all eight readers at once:

IFI PROFILE

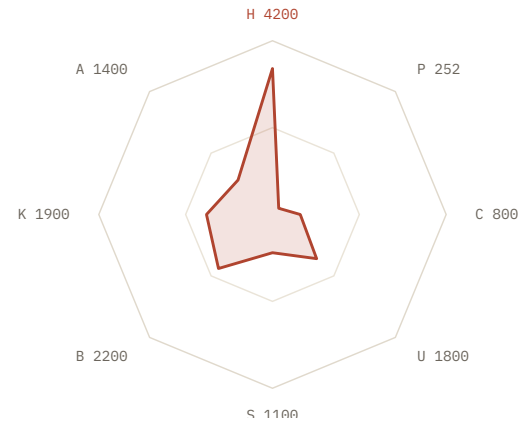

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400



Phoenix — the first cultivar engineered to register on all eight readers at once, rendered here as a single organism carrying a full eight-letter sentence.

Eight numbers, one plant, each anchored to the half-activation of a named protein — a chord written on purpose. It is the clearest proof we have that the readers can be addressed not one at a time but together, as a composition.

From eight streams to one flavor

None of this is how flavor is *felt*. You do not experience eight numbers. You experience one thing — the warm tomato, the drying tea, the slow building heat that turns, somewhere behind it, unexpectedly sweet — and that unity is the brain's doing. The eight streams leave the mouth as separate pulses of nerve traffic, climb to the brainstem and the thalamus and the gustatory cortex, and are braided there, with smell and touch and temperature and memory,

into the single seamless percept we call taste. The readers report; the mind composes. A flavor is a sentence the brain writes from the eight letters the tongue hands up.

Which is exactly why reading the receptor first changed everything for us. For ten thousand years the craft ran one way only — grow, taste, keep the seed of the one that moved you — a conversation we could feel but never follow. Knowing the eight readers inverts the arrow. Read the receptor, and you know the molecule it most wants to hear; know the molecule, and you can find — or build — the pathway that makes it; grow the pathway into a living plant, and you have written the word the reader was waiting for. Protein, pocket, molecule, pathway, plant: the order of discovery, run backward into an order of design.

Through all of it runs the quiet thing that started us at that table. That a glutamate loud enough to fill the umami reader can come from a tomato and not a packet. That heat can be a pleasure the body meets on honest terms rather than an assault. That a plant grown and read with this much care can be made to carry more — more flavor, and more of what the body needs to be nourished — at once. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. We will never say that as a claim. We only ever say it as an intention, and build toward it, as a family, plant by plant, reader by reader, until the whole eight-letter alphabet is ours to write.



PART I OF III

The Reading Machinery

what a receptor is · the mathematics of loudness · the two ways to open a channel

Chapter 1 · The Lock and the Key — *what a receptor actually is.*

Cut a pepper and hold the open face to the light, and you are looking at the placenta — the pale ribbed tissue where capsaicin is made and stored, bright with the compound that will, in a few minutes, convince a stranger's nervous system that their mouth is on fire. The pepper has no idea it is spicy. Spiciness is not a property of the fruit; it is an event that occurs in an animal. The fruit merely carries the key. Somewhere there has to be a lock.

A receptor is that lock: a protein, embedded in the fatty membrane of a cell, folded into a precise three-dimensional shape with a cavity — a pocket, a cleft, a vestibule — sized and charged to cradle one kind of molecule and ignore ten thousand others. When the right molecule arrives and settles in, the protein does the one thing proteins do better than any machine we build: it changes shape. That conformational change is the entire event. Everything downstream — the ion that rushes in, the nerve that fires, the face that softens at the table — is consequence. Taste begins as geometry, and the whole of this book is the study of a handful of geometries and the molecules that fit them. We write about them with reverence because they are, quite literally, how the body learns what the garden has grown for it: the instant a plant and a person become one chemical sentence.

The tongue is a reading room

The tongue stages this reading in **taste buds** — between two thousand and eight thousand of them in an adult mouth, each a tight bundle of fifty to a hundred elongated cells packed like the segments of an orange, their tips reaching up to a pore that opens onto the wet surface. The buds are not scattered at random. They cluster in the **papillae**, and the papillae come in kinds: the small button-like **fungiform** papillae dotted across the anterior tongue, a few buds apiece; the ridged **foliate** papillae along the sides; the large **circumvallate** papillae at the back, moated structures that can carry hundreds of buds each; and the **filiform** papillae, which carry no taste buds at all and exist only to give the tongue its grip and texture — pure mechanics, no chemistry. The old classroom "tongue map," with sweet at the tip and bitter at the back, is a mistranslation of a century-old German paper and is simply wrong: every taste quality can be read across the whole tongue, because the receptors travel in the buds and the buds are everywhere the papillae are.

And the cells themselves do not last. A taste receptor cell lives roughly **ten days** before it is replaced from a pool of basal stem cells — the one organ of perception you rebuild from scratch every week and a half, which is part of why taste recovers after a scald and why it dulls, then returns, across a life. Hold that number against the ones we work with. A cultivar takes a season to grow and a receptor ten days to die and regrow; the thing being measured is replaced a dozen times before the plant that stimulates it is ripe. Every mouthful

is read by a surface the body is in the middle of rebuilding, which is the most literal nourishment there is — you are, in part, built out of what you have recently tasted.



A papilla under the stage, the taste bud a tight rosette of cells — the receptors live in the apical membranes that reach up into the pore.

Not every cell in a bud reads the same thing; the bud is a committee of specialists. **Type II cells** carry the G-protein-coupled receptors for sweet, umami, and bitter, and when they fire they do not make the classic chemical synapse — they vent the transmitter ATP directly through a large channel to the waiting nerve. **Type III cells** are the ones with true synapses, and they read sour, through a dedicated proton-sensing channel. **Type I cells**, long dismissed as mere glial packing, turn out to manage the ionic housekeeping and much of the salt read. The signals they generate are carried to the brain along three cranial nerves — the **facial nerve (VII)** from the front of the tongue, the **glossopharyngeal (IX)** from the back, and the **vagus (X)** from the throat and epiglottis — a threefold redundancy that reflects exactly how much the body cares about what crosses the threshold of the mouth. Nothing else that enters you is inspected by three separate nerves on the way in.

Touch wearing the mask of taste

Threaded through all of it, not in the taste bud at all but in the free nerve endings of the **trigeminal nerve (V)** that blanket the entire mouth, is the somatosensory layer — the one that reports heat, cold, touch, and chemical sharpness. This is the crucial anatomical fact that the single word "taste" hides: **heat, cooling, and pungency are not tastes.** They are touch. They are your pain-and-temperature system being spoken to in a chemical language it was never meant to hear, and answering anyway. Three of our eight axes — **heat, cool-burn, and pungency** — belong to this trigeminal system, not to the gustatory one, which is why they feel like events rather than flavors, and why a chili can make your eyes water and your nose run when a sugar never will. A sweet is something your tongue concludes; a burn is something that happens to you.

The practical consequence runs through every chapter of designed flavor. The five gustatory axes modulate a *judgment* — the brain weighing a concentration. The three trigeminal axes trip an *alarm* — a threshold crossed, a gate thrown open. You can gently raise sweetness and have the eater barely notice the lift; raise heat past a receptor's switch-point and the response is not gentle at all.

Designing the two kinds of axis is a different craft, because one is a dial and the other is a trigger, and the receptor is what tells them apart.

Two families, two grammars

Across these cells, two great protein families do nearly all of the reading, and the difference between them organizes this entire book.

The first family is the **transient receptor potential (TRP) channels** — TRPV1, TRPM8, TRPA1. A TRP channel is a pore. Four protein subunits assemble around a central axis like staves around a barrel; each subunit contributes six membrane-spanning helices, and the last two of each line a central gate that the right stimulus pulls open. When it opens, cations — sodium, calcium — pour through the membrane down their electrochemical gradient, and the cell is depolarized within milliseconds. These are the channels of heat and cold and sharpness. They are fast, they are physical, and because they are literally holes that open, their signal has the urgency of a thing happening *to* you. This is why a chili's burn arrives like news.

The second family is the **class C G-protein-coupled receptors** — the sweet dimer TAS1R2·TAS1R3, the umami machinery of mGluR1 and TAS1R1·TAS1R3, the calcium-sensing receptor CaSR that reads kokumi — joined by the structurally simpler class A bitter receptors, the TAS2Rs behind our zero-bitter axis. A GPCR does not open a hole. It is a relay. The molecule binds on the outside; the receptor shifts; a G-protein waiting on the inner face is activated and splits off to trigger an enzyme cascade — in taste, the enzyme phospholipase C-β2 — which liberates a second messenger, which opens a *different* channel from the inside, which finally depolarizes the cell. More steps, more amplification, more time. This is why sweetness blooms and lingers where heat merely strikes — a cascade that unfolds station by station, each step buying gain and delay.

> A channel is a door. A GPCR is a doorbell. *The mouth uses both, and never confuses them.*

Two more families round out the eight. Astringency, the pucker of strong tea, is read by neither channel nor GPCR in the ordinary way — it is sensed when tannins bind and precipitate the **proline-rich proteins** of your own saliva, stripping the lubrication from the surfaces of your mouth, a loss your touch receptors report as drag. And sour, the cousin we set just outside our design axes, belongs to the same spare logic: a proton-sensing channel (OTOP1) in the Type III cell, reading the free hydrogen ions of an acid directly. Eight readers, then, built from a handful of structural themes, each one a lock evolution spent tens of millions of years cutting.

AXIS	NAME	RECEPTOR	PROTEIN FAMILY	WHAT THROWS THE LOCK
IFI-H	Heat	TRPV1	TRP ion channel	capsaicin, heat above 43 °C, protons
IFI-S	Sweetness	TAS1R2·TAS1R3	class C GPCR (dimer)	sugars, glycosides, sweet proteins
IFI-U	Umami	mGluR1 · TAS1R1·TAS1R3	class C GPCR	glutamate, 5′-ribonucleotides
IFI-B	Zero-Bitter	TAS2R38	class A GPCR	bitter alkaloids, glucosinolates (20–65% attenuation)
IFI-K	Kokumi	CaSR	class C GPCR	γ-glutamyl peptides, Ca^{2+}
IFI-C	Cool-Burn	TRPM8 (with TRPV1)	TRP ion channel	menthol, cold below 26 °C
IFI-P	Pungency	TRPA1	TRP ion channel	allicin, isothiocyanates, cinnamaldehyde
IFI-A	Astringency	salivary proline-rich proteins	precipitation, not a membrane receptor	tannins, polyphenols

Inside the channel — how a burn actually opens

Follow one molecule the whole way in, because the heat axis is where the lock-and-key is most literally a mechanism, and because it is the axis our family spent the longest learning to tune. Capsaicin is oil-loving; it does not drift through the watery saliva so much as dissolve into the fatty membrane of the nerve cell and approach TRPV1 from the lipid side. There it finds a pocket formed where the third and fourth membrane helices meet — the **vanilloid pocket** — and settles against a tyrosine and a threonine that cradle its vanillyl head by hydrogen bond while its fatty tail lies along the groove. That is the binding. What happens next is the whole trick of the channel.

The binding tugs on the short **S4–S5 linker**, a hinge that couples the sensor region to the gate. The tug travels to the **lower gate**, a constriction where the four inner helices cross, and dilates it. The pore, shut a microsecond ago, is now wide enough to pass ions, and because TRPV1 is non-selective and modestly calcium-preferring, what pours in is a mix of sodium and calcium flooding down a gradient the cell spends its whole life maintaining. The membrane voltage collapses toward zero, voltage-gated sodium channels downstream catch the swing, and an action potential runs up the trigeminal nerve to the brainstem and on to the cortex. The fire is reported. None of it was warm.

That last point is the elegant cruelty of the molecule. TRPV1 is, natively, a thermometer: it opens on its own above about **43 °C**, the temperature at which tissue genuinely begins to be damaged. Capsaicin does not add heat; it **lowers the channel's thermal threshold**, dragging the opening point down from 43 °C to below body temperature. At a steady 37 °C the channel is now standing partway open with no provocation at all — so the body, reading an honest signal from a channel that has been quietly recalibrated, reports a fire that is not there. The vanillylamine head is the part that does this; the **length and saturation of the fatty tail** set the balance

between capsaicin and its cousin dihydrocapsaicin, and with it the *shape* of the burn — where on the tongue it lands, how fast it climbs, how long it stays.

And it does not stay forever, which is its own mechanism. The calcium that rushed in during the burn is also a signal: it recruits calmodulin and the phosphatase calcineurin, strips an activating phosphate off the channel, and depletes the membrane lipid PIP_2 the channel needs to stay willing. The gate closes and resists reopening. This is **desensitization** — the reason a second forkful of the same dish lands softer than the first, the reason capsaicin is used clinically to quiet chronic pain, and a reason heat is a pleasure you can grow into rather than a fixed wall. The receptor teaches itself to tolerate the key. A flavor, read this closely, is never a static fact; it is a short, self-limiting conversation that the tissue both hears and answers.

From gene to key — how the plant writes the molecule

If the receptor is the lock, the plant is the locksmith, and capsaicin is cut by an enzyme. The vanillyl head comes down the **phenylpropanoid branch** that runs in every vascular plant — PAL → C4H → 4CL → HCT → COMT, phenylalanine worked step by step toward vanillin and then to vanillylamine by the aminotransferase pAMT. The fatty tail is a branched-chain acyl-CoA assembled from a valine- or leucine-derived starter by BCAT and KAS. The two halves are married by one terminal enzyme — **capsaicin synthase, Pun1/AT3**, an acyltransferase — which joins vanillylamine to 8-methyl-nonenoyl-CoA and releases capsaicin into the placental tissue. Four terminal enzymes are all that stand between the universal plant chemistry and the molecule that lights TRPV1. That narrowness is the whole thesis of our platform: the key is cut by a part the plants mostly already share.

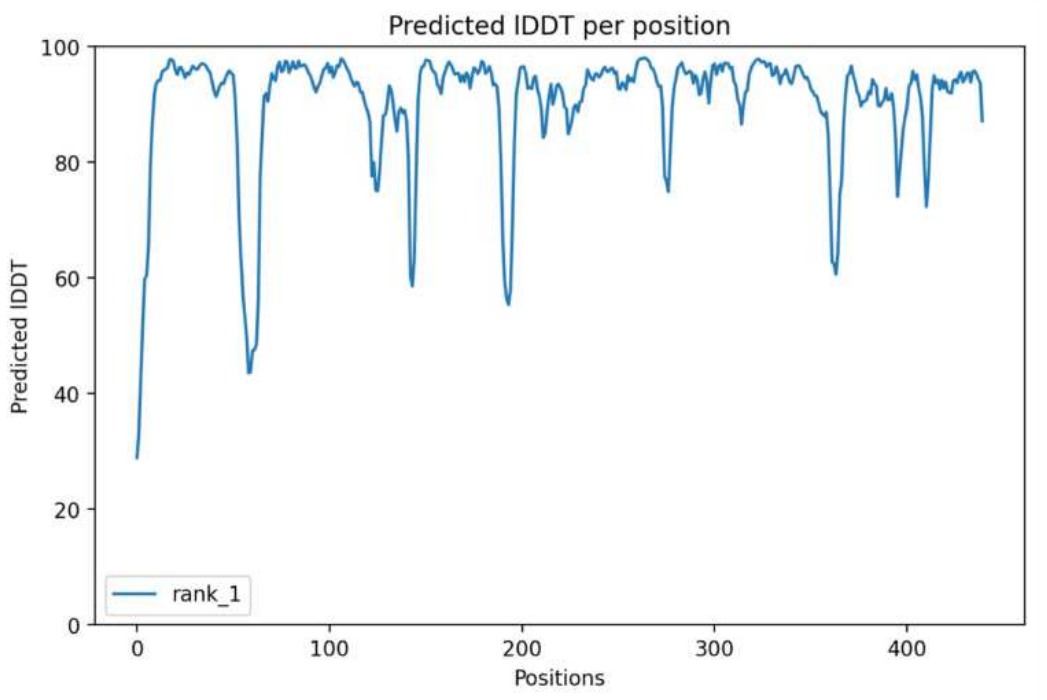

FIG. 1 — Pun1/AT3 modelled and coloured by pLDDT confidence; the binding domain and active-site cleft resolve as high-confidence, the surface loops less so.

Every prior breeding program chased heat by selecting the phenotype and hoping the genotype followed. We inverted it: a protein language model scored every single-residue substitution in Pun1 by predicted effect on catalytic efficiency and fold stability, and a genome-scale flux-balance model found that flux 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. Two substitutions emerged to improve the terminal enzyme's reliability — edits no breeder had ever stumbled into together; in our model, these edits reshape the enzyme's free-energy landscape. The language of that reshaping is **ΔΔG**, the change a mutation makes to a free energy, in kilocalories per mole, where a negative number means *more stable, more favorable*.

EDIT	WHERE IT ACTS	EFFECT	ΔΔG (KCAL/MOL)
S39L	N-terminus	raises substrate-binding affinity	−2.1
L345G	active-site cleft	steric relief, cleaner turnover	−1.8
C175S	*a mispairing cysteine*	*structural annotation only — prevents disulfide misfold; not in shipped construct*	*−0.9*
Sum (shipped double variant)	—	net stabilization of the working state	**−3.9**

The two shipped edits add to **−3.9 kcal/mol**, and that number can be read as a ratio rather than left as a figure. Thermodynamics relates a free-energy difference to an equilibrium constant by ΔΔG = −RT·ln(K), so the fold change in the balance between the folded, catalytically competent enzyme and its misfolded or unbound state is exp(−ΔΔG / RT). At body-adjacent temperature RT ≈ 0.59 kcal/mol, so

```
exp(3.9 / 0.59) = exp(6.6) ≈ 740×
```

Our model predicts the edits shift the equilibrium roughly seven-hundred-and-forty-fold toward the enzyme that actually works. Our model predicts that shift would translate not only as more heat but as *steadier* heat. A population of enzyme molecules that folds correctly nearly every time turns out the same amount of product pod after pod; our design targets a batch-to-batch coefficient of variation of **≤7%**, against 42% in wild-type — a plant you can finally cook from on purpose. Reverence for the plant and rigor about the plant are, here, the same act: we stabilized the locksmith so the key would come out true.

Putting a number on the fit — the Index

All of this — the pocket, the threshold, the cut molecule — is loudness waiting to be measured, and the measure is the **Iyengar Flavor Index**. Every axis is anchored not to how much of a compound a food holds but to how *loudly* that compound speaks to its receptor, and loudness has a precise definition: the **EC_{50}**, the concentration at which a stimulus drives its receptor to half of full response. A smaller EC_{50} means a louder molecule — less of it is needed to say the same thing. Invert it against a fixed reference and scale by a thousand:

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

Read the formula slowly and it tells you how to live in it. When a candidate's EC_{50} equals the reference's, the ratio is one and the score is exactly **1,000** — the anchor point, the voice of the native

wild type. Halve the candidate's EC_{50} and it is twice as loud, and the score doubles. The algebra also runs backward, which is how we turn any score into a statement about the receptor: since IFI = 1,000 × EC_{50}(ref)/EC_{50}(cand), the candidate's own half-activation concentration is

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

Each axis has its own reference point, defined by convention — the way Scoville's scale is defined by a dilution-to-threshold rather than an absolute unit. Scores on different axes are each expressed in their own axis's terms; they are not on a single common scale and cannot be directly compared across axes.

Anchor the heat reference where the biochemistry puts it — capsaicin half-activates TRPV1 near **0.70 µM** — and the scores become concentrations you can picture. Our model assigns **Phoenix (SS-0096)** a computed **IFI-H of 4,200**, so its predicted half-activation sits at 0.70 µM × 1,000/4,200 = **0.17 µM (167 nM)**: a quarter of the reference concentration lights the same channel to the same half-response. The same molecule, carried in a construct designed to deliver it, is **4.2× louder** than the wild type the scale is built around in our model. Convert that to mass to anchor it in something weighable — capsaicin's molar mass is 305.4 g/mol, and in dilute water a milligram per litre is a part per million — and 0.70 µM is 0.21 mg/L (0.21 ppm) at the reference, 0.051 mg/L (0.051 ppm) at Phoenix's design point. Loudness, molarity, and parts per million are three dialects for one fact.



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

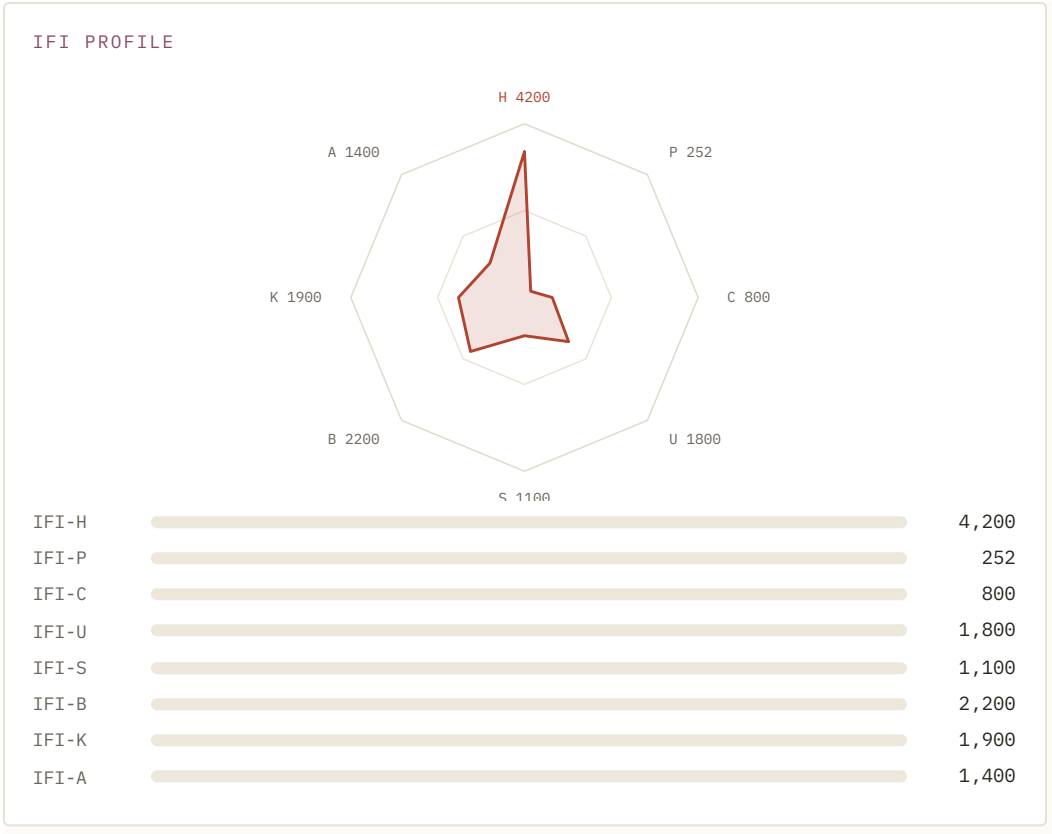

Between the score and the sensation sits the **dose–response curve**, and it is not a straight line but an S. The fraction of receptor driven, θ, follows a Hill relation in the stimulus concentration [A]: $\theta = [A]^n / (EC_{50}^n + [A]^n)$. The exponent n, the **Hill coefficient**, measures cooperativity — and because TRPV1 is a four-subunit channel whose subunits help each other open, n runs near **2**. That cooperativity is why heat feels like a switch. With n = 2, walking the concentration from a tenth to nine-tenths of full response takes only a ninefold change in dose, where a non-cooperative receptor would need eighty-onefold. The burn is engineered, by the receptor's own architecture, to arrive almost all at once.

[CAPSAICIN] ÷ EC_{50}	RECEPTOR RESPONSE θ (N = 2)
0.25×	0.06
0.5×	0.20
1×	0.50
2×	0.80
4×	0.94

Lowering a cultivar's EC_{50} — the whole point of the heat axis — slides this entire S-curve to the left along the concentration axis without changing its shape. Phoenix's curve is shifted left by that factor of 4.2: the same bite of fruit reaches nine-tenths of the channel's response at roughly a quarter of the capsaicin load the reference would demand. The score is the shift, written as one number.

The design progression across the ten flagships lets us see our model's predicted EC_{50} values fall rung by rung — and watch where the Index parts company from the old Scoville scale. SHU counts *how much* capsaicinoid is present; IFI-H predicts *how loudly* it speaks at the receptor. Nine of the ten flagships carry a computed IFI-H (Flare, SS-0090, predicted at 1.48M SHU, we leave unscored here), and each score is one over the formula away from an EC_{50} ratio:

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

All SHU values are computational predictions from our flux and biosynthesis models, not certified measurements.

Two readings hide in that table. First, divide each score by its heat — IFI-H per million SHU — and the loudness-per-pod does not hold constant: it runs near 1,530 at Blaze and settles toward 1,420 at Phoenix. The receptor is **saturating**. As capsaicinoid loading climbs into the top rungs, each additional molecule finds a channel population already largely spoken for, and the Hill curve flattens; the Index, which lives on that curve, records diminishing returns that raw SHU cannot see. Second, look at Cinder and Kindle out of order. Kindle carries *more* total capsaicinoid — 1.25M SHU against 1.05M — yet reads *quieter* at the receptor, IFI-H 2,000 against 2,290. SHU counts molecules; the Index counts how loudly they speak, and the two capsaicinoids that dominate these pods, capsaicin and dihydrocapsaicin, do not speak to TRPV1 with the same voice. A pod can be heavier and softer at once. The whole reason we design at the receptor is to hear that difference, which the dilution assay is deaf to.

Error is part of the number, not a footnote to it. The **coefficient of variation**, CV = standard deviation ÷ mean, is where consistency becomes arithmetic. At Phoenix's predicted heat of approximately **2.95 million SHU** and our design target of ≤7% CV, one standard deviation sits near **205,700 SHU**, so our model places pods inside roughly 2.73–3.14M at one SD, 2.53–3.35M at the 95% band. The wild-type 42% CV, laid against the same mean, would spread one SD across ±1.23M — the same plant ranging from 1.7M to 4.2M pod to pod, a harvest you cannot formulate around. (The saturation that cost us loudness-per-pod quietly pays part of itself back here: on the flattening top of the curve, a 7% wobble in capsaicinoid delivery lands as a smaller wobble in delivered IFI-H than it would on the steep lower slope.) To put a predicted number on the receptor and *stand behind it* takes both the double variant we designed to stabilize the enzyme and the honesty to state its scatter clearly.

What a single number cannot hold — the honest limits

A score is a measurement, not a verdict, and a measuring family owes the reader its edges. The Index reports how hard a receptor is driven; it does not report whether you will *like* the result. Hedonics — pleasure, craving, the particular ache for heat that chiliheads chase — live in the brain, downstream of the receptor, shaped by memory

and culture and the meal around the bite. IFI-H 4,200 is a fact about TRPV1; whether 4,200 is joy or assault is a fact about the person, and no amount of EC_{50} precision will settle it.

The number is also a scalar laid over something that moves in time. A burn has an onset, a peak, and a decay; sweetness blooms and lingers; desensitization is folding the second forkful softer than the first even as you measure it. A single IFI value is the amplitude of that arc, not its shape, and two cultivars that score alike can still feel unalike because one arrives slow and stays and the other strikes and goes. Where the shape matters, the shape has to be measured too.

And the lock is not identical in every reader. Human receptors are **polymorphic**: the bitter receptor TAS2R38 comes in a taster and a near-non-taster form, so the same cup of coffee is punishing to one person and merely dark to another, and TRPV1 carries its own variants that nudge heat tolerance up or down from birth. The Index is anchored to a reference receptor; a real mouth brings its own alleles, its own ten-day-old cells, its own history of what it has eaten. Add the matrix — fat blunts and carries capsaicin, acid sharpens, temperature and salt shift every threshold — and the in-vitro EC_{50} is a clean promise the dinner table will keep only approximately. Our models estimate EC_{50} from predicted dose–response relationships with stated confidence intervals, and the gap between a curve computed in silico and a flavor met in a mouth is a gap we name rather than paper over.

None of this is a retreat from the number; it is what the number is *for*. You cannot design what you cannot predict, and you cannot predict honestly without charting where the prediction runs out. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The receptor is the place where a plant we grew with intention becomes a sensation in a body we mean to nourish — eight locks, a handful of keys, cut for ten thousand years by selection and now, for the first time, read closely enough to improve. To design flavor is to understand the lock well enough to better what evolution left in it, and the first act, always, the one this whole book rests on, is to predict exactly how well the key would fit.

Chapter 2 · The Half-Maximal Point — *the mathematics of loudness.*

Hand two cooks the same tomato and they will give you two different words for it — *bright*, *deep*, *round*, *green* — and every word will be true and none will be a number. For ten thousand years that was the whole problem with flavor: it could be felt exactly and measured not at all. You could say a pepper was hotter than another pepper. You could not say *how much* hotter in any unit a chemist and a chef would both accept. We built the Iyengar Flavor Index to end that silence, and the Index rests entirely on a single idea borrowed from pharmacology: the half-maximal point.

The reason the half-maximal point is the right place to stand is that flavor does not happen in the food. It happens at a surface — the apical membrane of a taste or trigeminal cell — where a molecule arrives, finds a receptor shaped to read it, and changes that receptor's conformation. Everything we call taste is the sum of those arrivals, counted by the nervous system and reported upward. So if we want a number for flavor, we should not weigh the molecule. We should measure the conversation.

The receptor, up close

Take heat, because heat is the axis whose receptor is best understood. Capsaicin binds TRPV1, a non-selective cation channel standing in the membrane of the trigeminal nerve endings that web the tongue — the same channel gated by temperatures above 43 °C and by acid. Four identical protein subunits assemble around a central pore, each contributing a wall; the whole thing is a four-fold symmetric gate held shut until something persuades it open. Capsaicin is one such persuader. Its vanillyl head slides into a pocket between the S3 and S4 helices on the inner leaflet of the membrane, where a glutamate and a tyrosine hydrogen-bond the molecule into place. That small binding tugs the S4–S5 linker, the linker leans on the lower gate, and the gate — the ring of hydrophobic residues that had been squeezing the pore shut — relaxes outward.

What the molecule does *not* do is heat the tissue. It lowers the channel's thermal threshold, so a receptor that would have fired only above 43 °C now fires at the temperature of your own mouth. The fire is real as a signal and absent as a fact; the nerve reports a burn that the thermometer cannot find. When the gate opens, calcium and sodium pour down their gradients into the cell, the membrane depolarizes, voltage-gated channels downstream catch and amplify it, and a train of action potentials runs up the trigeminal nerve toward the brainstem. That is the entire event. Loudness, for us, is simply: at what concentration does that event reach half of everything it can become.

Upstream of all this is the plant's own chemistry, which decides *which* persuader it makes. Vanillylamine condenses with an 8-methyl branched-chain fatty acid at the Pun1/AT3 acyltransferase to form

capsaicin; the length and saturation of that acyl chain set the balance between capsaicin and dihydrocapsaicin, and with it the *shape* of the burn. Two peppers can carry the same milligrams of total capsaicinoid and still speak at different volumes to TRPV1, because the blend sits differently in that S3–S4 pocket. Hold that fact. It is the hinge on which this whole chapter turns, and the reason amount and loudness have to be two numbers and not one.



The stage of the dose–response assay — a single construct's tissue read across a rising ladder of concentration, so the curve can be fit and its midpoint found.

The curve, and why it bends

Picture a receptor and a rising concentration of the molecule it reads. At vanishingly low concentration, almost nothing happens — a receptor here and there catches a molecule, the cell barely stirs. Raise the concentration and the response climbs. Raise it further and the climb slows, because the receptors are filling up; past a point, every lock holds a key and adding more keys changes nothing. Plot response against concentration and you get not a straight line but a graceful S — a sigmoid — rising from zero, steepening through a middle, and flattening at a ceiling we call the maximal response, Emax.

The shape of that curve is described, with lovely economy, by the Hill equation:

$$\text{response} = E_{max} \cdot [L]^n / (EC_{50}^n + [L]^n)$$

Here [L] is the concentration of the molecule, the ligand. The exponent *n* is the Hill coefficient, a measure of cooperativity — how much the binding of one molecule helps the next. And **EC_{50} is the single number that matters most: the concentration at which the response reaches exactly half of its maximum.** It is the midpoint of the S, the hinge of the whole curve. A low EC_{50} means the receptor answers at a whisper — a tiny concentration already drives it halfway. A high EC_{50} means the molecule must shout to be heard. EC_{50} is not how much of a compound a food contains. It is how *loudly* that compound speaks to its receptor.

It helps to divide through by Emax and work in the fraction of the ceiling, because then the number has no units and the arithmetic falls open:

```
f = response / Emax = [L]ⁿ / (EC₅₀ⁿ + [L]ⁿ) = 1 / (1 + (EC₅₀ / [L])ⁿ)
```

Set [L] equal to EC_{50} and the ratio inside the bracket becomes one; one raised to any power is one; f = 1 / (1 + 1) = 0.5. The curve passes through half, exactly, by construction — which is the whole reason we chose that point to anchor a scale. It is the one place on a sigmoid that does not move when you argue about where the top is. The ceiling Emax can be hard to pin down; the foot of the curve vanishes into noise. But the midpoint is sharp, reproducible, and indifferent to the units you measured response in. We anchor to the half-maximal point because it is the most honest coordinate the curve owns.

Reading the dose–response table

The Hill coefficient is where the curve hides its character, and the cleanest way to feel it is to tabulate the fraction-of-maximum, f, against concentration measured in units of EC_{50} itself — so that $[L]/EC_{50} = 1$ is always the midpoint. We do it twice: once for $n = 2$, the signature of the four-subunit TRP channels like TRPV1 and TRPM8, whose gate opens cooperatively; and once for $n = 1$, the gentler slope of the single-site GPCR reads behind sweetness and umami.

[L] ÷ EC_{50}	F AT $N = 2$ (TRP, E.G. TRPV1)	F AT $N = 1$ (GPCR, E.G. SWEET)
0.25×	5.9%	20.0%
0.50×	20.0%	33.3%
1.00×	50.0%	50.0%
2.00×	80.0%	66.7%
4.00×	94.1%	80.0%

Read the two columns side by side and you can see the difference you taste. For $n = 2$, almost the whole response is spent between a quarter of EC_{50} and four times it — from 5.9% to 94.1% across a sixteen-fold window, most of it slammed into the middle. For $n = 1$, the same sixteen-fold sweep of concentration moves you only from 20% to 80%, a long gentle ramp with room on both ends. The arithmetic makes it exact. Ask how wide a concentration range carries a receptor from a tenth of its response to nine-tenths — the honest working span — and the Hill equation answers with a single expression:

```
dynamic range (10% → 90%) = 81^(1/n)
```

For $n = 1$ that is a factor of 81; for $n = 2$, the square root of 81, a factor of 9; for a steep $n = 4$ channel, the fourth root, a factor of 3. A steep curve is a cliff: a nine-fold window is nothing in a kitchen, which is why heat so easily goes from nothing to too much, and why cooking with superhots is less like seasoning and more like defusing. A shallow curve is a ramp you can walk up — sweetness has a dynamic range you can ride, a little to a lot, and that is a gift of its

receptor's arithmetic, not its chemistry. The steepness at the very midpoint falls out of the same equation as a slope of $n/4$ per unit of $[L]/EC_{50}$: the TRP curve climbs through its middle twice as fast as the GPCR curve. The mathematics is not decoration. It is the reason the eight axes *feel* the way they feel, and designing a flavor means designing the whole curve — its midpoint, its height, and its steepness all at once.

Potency is not amount

Two further pharmacological ideas sharpen the picture, and both are places where an honest index has to be careful. The first is the difference between **potency** and **efficacy**: potency is *where* the curve sits on the concentration axis (its EC_{50}), while efficacy is *how high* it rises (its Emax). A molecule can be exquisitely potent — triggering a response at vanishing concentration — and yet be a **partial agonist**, unable to drive the receptor to full response no matter how much you add. Some intense sweeteners are exactly this: enormously potent, so that micrograms register, yet capped below sugar's full sweetness, which is why a diet drink can taste intensely sweet and still somehow *thin*. EC_{50} would score such a molecule as loud; the tongue would still find it incomplete. We note efficacy alongside potency for precisely this reason, and treat a low EC_{50} sitting under a low Emax as the warning it is.

The second idea is **receptor reserve** — a cell often carries more receptors than it needs to reach its maximal signal, so the response can saturate well before every receptor is occupied. Occupancy and response come apart: a tissue can read 90% of full response with only half its receptors filled, because the downstream amplification runs out of ceiling before the binding does. This is why the dose–response curve (what you feel) can sit to the left of the binding curve (what is bound), and why measuring one and assuming the other is a quiet way to be wrong. Flavor design lives in all three dimensions at once: potency, efficacy, and reserve. We are not merely making a molecule "stronger." We are placing its whole curve.



A cross-section at the stage — the compound is there in the tissue in some amount, but amount is not loudness; loudness is where the curve crosses half.

Scoville measures the dose

Contrast all of this with the **Scoville scale**, the measure the chili world has used since 1912. Scoville measured *amount* — how many times a pepper extract had to be diluted before a panel of tasters

could no longer detect the heat. It is a concentration measure, today performed by HPLC (high-performance liquid chromatography), which quantifies capsaicinoids in parts per million and multiplies by roughly sixteen to recover Scoville Heat Units:

```
SHU ≈ 16 × (capsaicinoids, ppm)
```

The constant is anchored at the top: pure capsaicin is one million parts per million — a hundred percent — and lands near 16,000,000 SHU, which is exactly 16 × 1,000,000. Run the conversion the other way and a Scoville number becomes a mass fraction. We do it here on a dry, placental-tissue basis, the fraction where capsaicinoids actually concentrate inside the fruit. The values for our cultivars in the table below are computational design targets — our model predicts these are the concentrations our engineered placentas could hold, far above anything a wild landrace carries:

REFERENCE POINT	SHU	CAPSAICINOIDS (PPM)	% BY DRY MASS (PLACENTA)
Pure capsaicin (the ceiling)	16,000,000	1,000,000	100.00%
Phoenix · SS-0096 (predicted)	2,937,948	183,622	18.36%
Ember · SS-0087 (predicted)	866,000	54,125	5.41%

Scoville tells you how much capsaicin is present. It tells you nothing about how well that capsaicin fits TRPV1 — and that gap is not academic. The engineered blend inside a fruit can lean toward the capsaicinoid that sits more snugly in the S3–S4 pocket, so two fruit holding identical milligrams can open the channel at measurably different volumes. EC_{50} measures eloquence; Scoville measures dose. The Index needs both, and keeps them distinct.

A word of honesty about that factor of sixteen, because we would rather show the seam than hide it. The constant is an approximation that assumes a fixed blend. Capsaicin, dihydrocapsaicin, and nordihydrocapsaicin carry different pungency weightings, so a fruit that shifts its blend shifts the true conversion by a few percent even at constant total mass — the same shift that moves its EC_{50}. The two measures are not independent in the way the clean table pretends; the blend writes on both at once. We report the sixteen as *roughly* sixteen for exactly this reason, and we trust the EC_{50} we fit over the SHU we convert whenever the two disagree.

From half-max to a score

From EC_{50} the Index is one short step. Fix a reference — the half-activation of a chosen native, wild-type standard — and score any candidate as the ratio of the two, scaled so the reference lands at a thousand:

```
IFI(c) = 1000 · EC50(ref) / EC50(c)
```

Read it slowly and it is almost obvious. If a candidate's EC_{50} is exactly the reference's, the ratio is one and the score is 1,000 — the candidate is, by definition, as loud as the standard. If the candidate reaches half-activation at *half* the concentration — twice as eloquent

— its EC_{50} is halved, the ratio doubles, and the score is 2,000. The reference is the plant before we touched it, and on every axis, by construction, it reads one thousand flat:

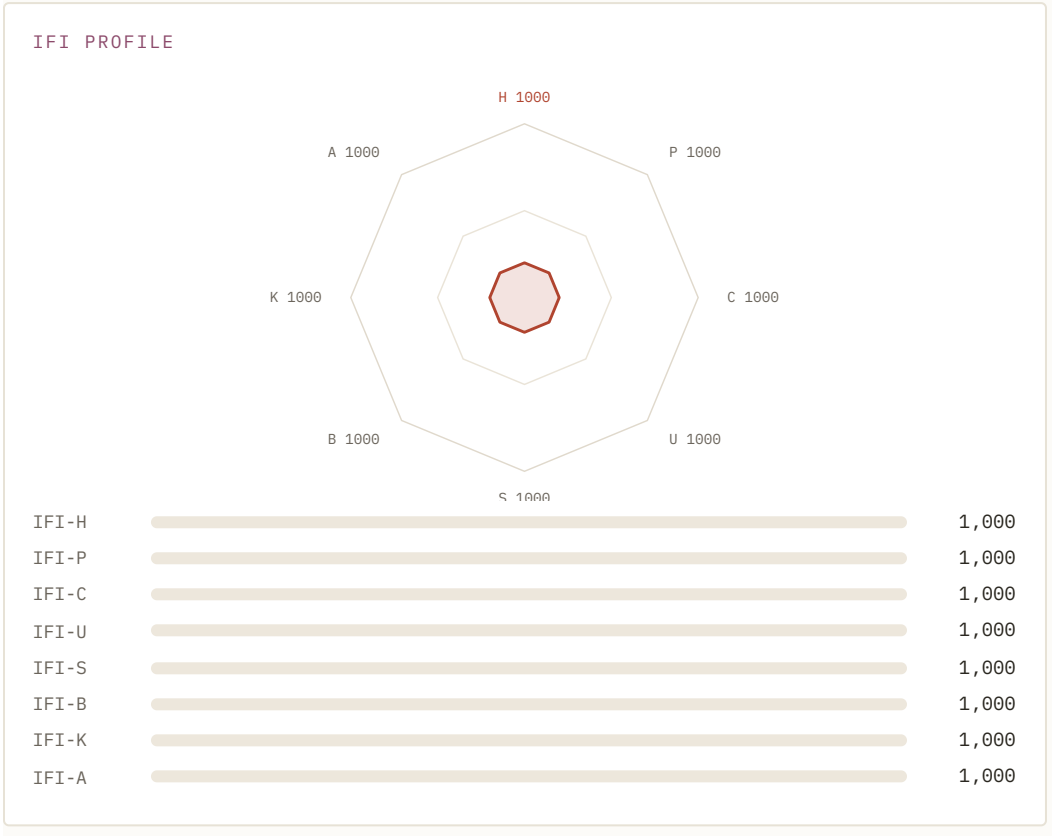

That flat profile is not a measurement; it is the origin we measure *from*. The reference constants are stated conventions — each one defines IFI = 1,000 on that axis by our choice of native standard, the same way the Scoville scale is defined by dilution-to-threshold rather than by a universal measuring rod. And because each axis carries its own independently defined reference point, scores on different axes are not on a single common scale: an H score of 1,500 and a P score of 1,500 each signal the same ratio over their own native benchmark, but the benchmarks differ, and the Index does not claim they correspond to the same absolute potency. Every score in every book is a distance from that octagon of thousands, measured along its own axis alone.

And because the formula is a pure ratio, it inverts cleanly — given a score, you can recover exactly how eloquent the candidate is relative to the standard:

> $EC_{50}(c) / EC_{50}(ref) = 1000 / IFI(c)$

A score of 2,000 means the candidate reaches half-activation at half the reference's concentration (ratio 0.500); a score of 4,200 means it does so at 1000 ÷ 4200 = 0.238 of it — a little under a quarter — which is precisely the heat score of the cultivar at the far end of everything we grow. Here is the full heat ladder of the ten flagship superhots, each canonical IFI-H turned back into the EC_{50} ratio it encodes, standing beside the predicted Scoville amount so you can watch the two numbers move together — and, in one telling place, apart.

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

Now find the place where the two numbers disagree, because it is the whole argument of this chapter written in our own catalogue. Cinder carries a predicted 1.05 million SHU and scores 2,290 on heat. Kindle carries *more* predicted capsaicinoid — 1.25 million SHU — and scores *less*: 2,000. Its EC_{50} ratio is 0.500 against Cinder's 0.437, meaning Kindle needs more concentration to reach half-activation despite holding more of the compound. The dose went up; the eloquence went down. Our model reads this as a blend that leaned, in Kindle, toward the capsaicinoid that sits less snugly against TRPV1 — more mass, worse fit. A grower measuring Scoville alone would rank Kindle the hotter plant. The tongue, and the Index, know better. This is not a rounding artifact or a flaw in the ladder; it is the exact gap between amount and loudness, and it is the reason we measure the half-maximal point instead of weighing the fruit.

One cultivar in that ladder, Flare, still shows a dash where its heat score belongs. We leave it blank on purpose. Our model's curve fit has not yet cleared our own bar for a clean predicted EC_{50}, and we would rather print a gap than a guess. A ladder with an honest hole in it is worth more than one with a number nobody can stand behind.

At the top sits the cultivar we designed gene by gene because the plant that could reach it did not yet exist — loud on heat and, unlike every rung below it, saying something on all eight axes at once:

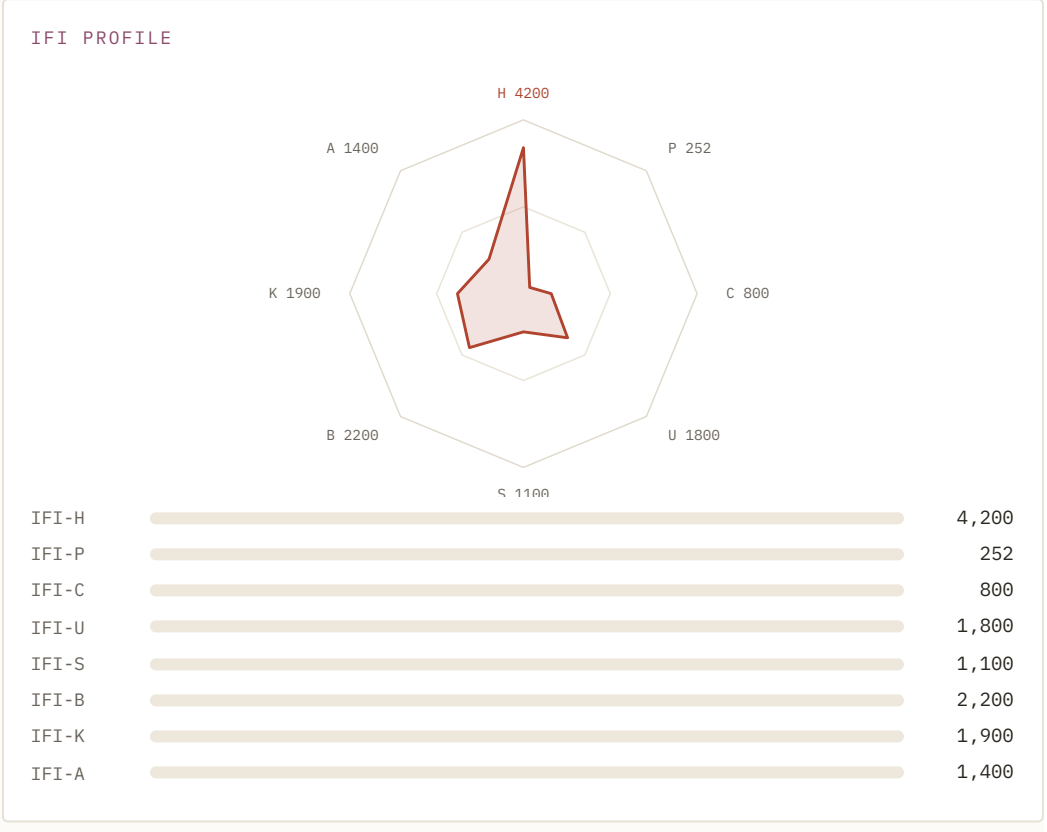

Because loudness can in principle run away from us, we cap each axis at **5,000** for the purpose of the eight-pointed scorecard — roughly a fivefold ceiling over the reference, an EC_{50} ratio of 0.20 — so that no single axis can swallow the radar and the whole read stays legible on one plate. The cap is a design limit, not a wall in the biology; a molecule can in truth be louder than 5,000, and when one is, we say so in the text and let the scorecard hold at its ceiling. A number that cannot be drawn is a number that cannot be shared, and these books are built to be read on one page.

A worked read, end to end

It is worth walking one candidate through the whole method, so the formula stops being a symbol and becomes a procedure. Suppose we model a new construct and fit its predicted heat curve, and the model returns a half-activation at 0.300 of the reference's — it reaches the midpoint of its sigmoid at three-tenths the concentration the native standard needs. That ratio is the entire result; everything else is arithmetic. The score is the reference EC_{50} over the candidate's, scaled by a thousand: 1000 ÷ 0.300 = 3,333. We set it on the ladder and it lands exactly where the ratio says it should, between Scorch at 3,220 (ratio 0.311) and Firestorm at 3,490 (ratio 0.287) — a little more eloquent than the first, a little less than the second, and no mystery about why.

Now carry the error through, because a lone 3,333 is only half a result. With the reference pinned down across many runs and the candidate holding our 7% process spread, the score inherits that same 7%: 3,333 ± 233. Let the reference carry its own 7% as well and the quadrature widens it to ±9.9%, about ±330. So the honest written form is 3,333 ± ~230–330 — upper-ladder, and known to roughly one part in ten. If that same construct happened to assay at a predicted 2.00 million SHU, its amount would convert to 2,000,000 ÷ 16 = 125,000 ppm, or 12.5% capsaicinoid on the dry placental basis — mid-ladder by dose while upper-ladder by eloquence, one more fruit whose predicted weight and whose voice disagree.

There is a quiet property of the scale hiding in that calculation, and it matters for how the whole library is read. Because the Index is a pure ratio, equal *multiplications* of eloquence are equal *multiplications* of score: halving an EC_{50} always doubles the number, whether it carries you from 1,000 to 2,000 or from 2,000 to 4,000. The scale is multiplicative, not additive — a ratio scale — so the meaningful comparison between two scores is their quotient, not their difference. The gap from 1,000 to 2,000 is the same fact about the receptor as the gap from 2,000 to 4,000: in both, the candidate speaks at half the concentration of the thing below it. Read the Index as ratios and it tells the truth; read it as a tape measure and it will quietly mislead you. This is why we built it on a half-maximal point and a division, and not on a difference of intensities that would have no fixed zero to stand on.

What the number carries, and what it doesn't

A score is a measurement, and a measurement without its error is a rumor. So we carry the error. EC_{50} is not read off a dial; it is fit to a ladder of replicate dose–response points, and the fit returns the

midpoint with a spread around it. Because the Index is a ratio of two such fits, the relative errors do not add — they add in quadrature:

$$(\sigma_{IFI} / IFI)^2 = (\sigma_{ref} / EC_{50,ref})^2 + (\sigma_c / EC_{50,c})^2$$

Our engineered constructs are designed to hold a batch-to-batch coefficient of variation at or under **7%**, against **42%** in the wild-type plants we started from — six times tighter, the difference between a plant you can formulate around and one you can only hope about. Feed that 7% into the quadrature on both the reference and the candidate and the Index inherits $\sqrt{(0.07^2 + 0.07^2)} \approx 9.9\%$. Measure the reference once and for all across many runs so its own error shrinks toward zero, and the Index settles back toward the candidate's own 7%. Either way, Phoenix's heat score is not a bare 4,200. It is 4,200 give or take something between about 290 and 420 — call it 4,200 ± 300 when the reference is well-pinned. We print the central number and carry the band in our heads.

This also means we have to be honest about the Scoville figures we quote to the single digit. If we take our model's predicted design center for Phoenix at 2,937,948 SHU, a 7% coefficient of variation would put a one-sigma spread of about ±205,700 SHU on any single fruit. That seven-figure number is a statement about the *designed population mean*, not a promise about the pepper in your hand — the standard error of the mean falls as $1/\sqrt{N}$ across batches, so the population parameter can be known tightly while each individual fruit still rides the full 7%. To know a mean to a part in three million you would need thousands of batches; to pick one fruit is to accept ±200,000. We report the mean to its digits because it is the honest center of the distribution, and we tell you the width because hiding it would be the dishonest thing.



The controlled chamber where the batches are grown to the same light, the same water, the same hour — the physical source of a 7% coefficient of variation where the field gives 42%.

There are limits the Index does not pretend to cross, and a chapter on method owes them plainly. EC_{50} is assay-dependent: TRPV1's half-activation shifts with temperature, with pH, with the expression system the channel is read in, so every score from our model is implicitly *at our defined parameters, by our protocol*, and a reference held constant in our model holds the whole scale together only because we never let it drift. The Hill curve we fit is a still photograph of a moving thing — TRPV1 desensitizes as calcium floods in, so the living response rises and then sags, and a single EC_{50} cannot capture the sag; it captures the rise, which is what first

reaches the nerve and first reaches the tongue. The Index scores the arrival, not the whole arc, and we say so. And a number can only ever be a coordinate: it fixes where a flavor stands on one axis of one receptor, which is a real and useful thing and not the same thing as the flavor. The map is not the meal.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. Every score in this chapter is a computed prediction grounded in molecular geometry and receptor pharmacology — the same biology this chapter describes, applied forward through our platform rather than read backward from a bench.

This is what we mean when we say numbers are a form of respect. A score is not an attempt to replace the word *round*, or the softening of a face at the table, or the quiet we fall into when something we grew turns out well. It is an attempt to be precise enough about what the receptor is doing — and precise enough about how well we know it — that we can grow the plant that does it better, feed a body more of what we meant to, and set one tomato, one chili, and one cup of tea against each other on a single honest page. We measure the half-maximal point because it is the one place on the curve that will not lie to us, and because the family that grows this food would rather give you the real number, error and all, than a prettier one we could not stand behind.

Chapter 3 · Two Ways to Open a Channel — *fast fire, slow savor.*

Taste has a tempo, and the tempo is chemistry. A chili strikes; sweetness blooms; umami settles in and stays; the pucker of tannin arrives late and leaves slowly. These are not metaphors. They are the measurable time signatures of two fundamentally different ways a cell can be told that a molecule has arrived — the ionotropic and the metabotropic — and once you can hear the difference, you can never un-hear it. Everything our family builds is built against this clock. When we set out to make a plant speak on eight axes at once, we were, in the end, writing two kinds of sentence: one the mouth reads as an event, and one it reads as a message.

The difference is not a matter of degree. It is a matter of architecture — two machines that do the same job, carrying a molecule's arrival inward to the nerve, by means so different that they even fail differently. One is a door. The other is a relay. A receptor book is really a book about these two machines, repeated with variations across the whole surface of the tongue, and the purpose of this chapter is to take them apart slowly enough that you can feel the gears.



Calcium floods a single sensory cell the instant its channel opens — the fast read caught mid-event, the molecule and the signal almost the same thing.

The ionotropic route: the molecule is the switch

The **ionotropic** way is the way of the TRP channels, and it is almost indecently direct. The molecule *is* the switch. Consider TRPV1, the receptor that reads heat. It is a tetramer — four identical protein subunits folded around a single central pore, each subunit threading the membrane six times. The outer four of those six passes form a voltage-sensor-like domain, and tucked into a pocket between two of its helices sits the vanilloid site, a cradle shaped to hold capsaicin against two residues the channel cannot do without. The inner two

passes from each subunit line the pore itself: a narrow selectivity filter at the top, a bundle-crossing gate at the bottom that is, at rest, pinched shut.

Capsaicin reaches that cradle, binds, and the channel's gate swings open — not through any intermediary, but because the binding energy of the molecule is paid directly into the mechanical work of opening the pore. This is allosteric gating in its purest form. The channel lives in a thermal equilibrium between a closed state and an open one; raising temperature past about 43 °C tips that equilibrium toward open, and so does capsaicin, by binding the open conformation more tightly than the closed one and dragging the balance across. The molecule does not warm the tissue. It lowers the channel's thermal threshold, so the body reports a fire that is not there — reports it, specifically, as a cation current, because once the gate is open the pore is a non-selective cation channel markedly permeable to calcium, passing calcium roughly ten times as readily as sodium.

Count the stations between molecule and signal and there is one. Bind, and open. From capsaicin to a collapsing membrane voltage is a single step and a few milliseconds — sub-millisecond for the conformational change itself, a handful of milliseconds to the first nerve impulse. There is no amplification and no delay because there is no cascade. The molecule does the opening with its own hands.

And yet the signal is enormous, which is the first thing worth doing the arithmetic on, because the smallness of the step hides the size of the effect. A single open TRPV1 channel has a conductance on the order of 50 picosiemens. Across the roughly 60 millivolts of driving force a resting cell offers an incoming cation, that conductance carries a current of

$$I = g \cdot V = (50 \times 10^{-12}\ \mathrm{S}) \cdot (0.060\ \mathrm{V}) = 3.0 \times 10^{-12}\ \mathrm{A} = 3\ \mathrm{pA}.$$

Three picoamperes sounds like nothing. Divide it by the charge on a single ion, 1.6×10^{-19} coulombs, and it is not nothing at all:

$$3.0 \times 10^{-12}\ \mathrm{C \cdot s^{-1}} \div 1.6 \times 10^{-19}\ \mathrm{C} \approx 1.9 \times 10^{7}\ \text{ions per second}.$$

One open channel passes nearly nineteen million ions every second. The ionotropic route needs no cascade to raise a loud signal; it raises one by brute flux, flinging charge through a pore held wide by a single molecule's grip. This is the architecture of urgency. It is why the three somatosensory axes we design toward — heat, cool, and pungency — all feel like something *happening*, a line crossed, an alarm. The body treats them as touch because, mechanically, they are: a channel opening is a physical event, not a chemical message. Pungency commits even harder than heat. Its channel, TRPA1, is thrown open by electrophilic molecules — the allyl isothiocyanate of mustard, the cinnamaldehyde of cinnamon, the thiosulfinates of cut garlic — that do not merely bind but *bond*, covalently modifying cysteine residues on the channel's inner face. A heat molecule can let go. A pungency molecule has, for a moment, become part of the lock.

The arithmetic of loudness

Because the ionotropic read is a channel opening, the right way to measure it is to ask how little stimulus it takes to half-open the population — the concentration at which a stimulus drives its receptor to half of its full response, its EC_{50}. Half-maximal activation is the fairest single landmark on any dose–response curve: it sits at the steep middle where the receptor is most sensitive to a change in dose, and it is a property of the stimulus-and-receptor pair, not of how much you happened to apply. The whole Index is anchored there.

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

Read the formula as a sentence about loudness. A *louder* stimulus is one the receptor answers at lower concentration — a smaller EC_{50}. Put EC_{50} in the denominator and loudness climbs as the number falls; invert it against a fixed reference, scale by a thousand, and you have a score. The reference is chosen by convention so that a native wild-type stimulus lands at exactly 1,000: set candidate equal to reference and the ratio is one, the score is 1,000, the anchor holds. This is a stated convention — the reference constant for each axis defines IFI = 1,000 on that axis in the same way the original Scoville scale was defined by dilution-to-threshold, not derived from a body of published values. Scores above a thousand are louder than wild-type; below, softer. For the radar scorecards our family draws, each axis is normalized to a ceiling of 5,000.

The arithmetic is worth walking once, slowly, on the plant we pushed furthest. Our model predicts Phoenix at IFI-H 4,200. Rearrange the formula and the score tells you, exactly, how concentrated Phoenix's capsaicinoid extract must be to half-open TRPV1 relative to the reference:

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

Phoenix is designed to reach the receptor's half-point at 23.8% of the reference concentration — 4.20 times louder in design terms, not 4.20 times more capsaicin necessarily, but 4.20 times more *potent at the channel* by our model's prediction. Run the same step down the flagship series and the whole range resolves into a single, legible column of ratios, each one nothing more than 1,000 divided by a computed score.

CULTIVAR	SKU	SHU (PREDICTED)	IFI-H	EC_{50} RATIO (CAND ÷ REF)	LOUDNESS ×REF
Ember	SS-0087	866K	1,570	1000 ÷ 1570 = 0.637	1.57×
Blaze	SS-0091	1.72M	2,630	1000 ÷ 2630 = 0.380	2.63×
Firestorm	SS-0094	2.47M	3,490	1000 ÷ 3490 = 0.287	3.49×
Phoenix	SS-0096	2.95M	4,200	1000 ÷ 4200 = 0.238	4.20×

Nothing in that table is invented; every ratio is a canonical IFI score run once through the formula. That is the discipline the number buys us. A pepper's heat stops being an adjective and becomes a quotient you can check with a calculator.

From Scoville to receptor: two numbers, two meanings

It is tempting to think the Scoville Heat Unit already did this job, and it is worth being precise about why it did not — because SHU and IFI measure two genuinely different things, and the gap between them is where the receptor hides. SHU is a *concentration* measure: historically a dilution — how far you must water a sample down before a panel can no longer taste the heat — and today a chemical readout of total capsaicinoids, linear in how much of the molecule is present. IFI is a *potency* measure: the inverse EC_{50}, how loudly that molecule speaks to the channel. One counts the molecules. The other listens to them.

The conversion between a capsaicinoid concentration and its Scoville value is fixed by definition — roughly 16 SHU per part-per-million of capsaicinoid — so Phoenix's predicted heat of approximately 2,937,948 SHU translates to a predicted loading:

```
~2,937,948 SHU ÷ 16 ≈ 183,622 ppm = 18.36% capsaicinoid-equivalent by weight.
```

Now set that concentration number beside the loudness number and watch them disagree. If loudness simply tracked concentration, the ratio of IFI to SHU would be constant down the series. It is not.

CULTIVAR	SHU (PREDICTED)	IFI-H	IFI PER MILLION SHU
Ember	866K	1,570	1,570 ÷ 0.866 = 1,813
Blaze	1.72M	2,630	2,630 ÷ 1.72 = 1,529
Firestorm	2.47M	3,490	3,490 ÷ 2.47 = 1,413
Phoenix	2.95M	4,200	4,200 ÷ 2.95 = 1,424

The loudness bought per unit of heat *falls* as the heat climbs — from 1,813 at the bottom of the series to about 1,424 at the top. This is not a flaw in the computation. It is the receptor telling the truth about itself. A dose–response curve saturates: once most channels are open, adding more molecules cannot open channels that are already open. Model the response with a Hill function,

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

with a Hill coefficient near n = 2, the value TRPV1's cooperative, four-subunit gating tends to show. At the half-point the response is, by definition, $R_{max}/2$. Double the concentration and it rises to 4 ÷ (1 + 4) = 0.80 of maximum. Quadruple it and you reach 16 ÷ 17 = 0.94. Ten times the stimulus buys you the last few percent and no more. So a cultivar carrying 18% capsaicinoids by weight is climbing the flat shoulder of the curve, where enormous concentration converts into only incremental loudness — which is exactly why Phoenix's loudness-per-SHU has fallen below Ember's. The honest limit of the heat axis is written into the shape of the channel: there is a ceiling, and near it the Scoville column keeps climbing while the Index column barely moves. We design knowing that past a point, more capsaicin is a more expensive way to buy less fire.

The metabotropic route: a message passed hand to hand

The **metabotropic** way is the way of the GPCRs — sweet, umami, bitter, kokumi — and it trades speed for amplification and nuance. The molecule binds on the outside and the receptor never touches a

pore.



Type II taste cells in dish culture — the GPCR readers, the ones that vent their message through a channel they open from the inside.

The sweet and umami receptors are class-C GPCRs, built around a great clamshell ectodomain — a Venus flytrap of two lobes that closes on its ligand like a hand closing on a seed. Sugar settles into the TAS1R2·TAS1R3 heterodimer; glutamate into the umami machinery, the mGluR1 and TAS1R1·TAS1R3 readers; a bitter alkaloid into one of the two-dozen TAS2R proteins, of which TAS2R38 is the classic. In each case the clamshell closes and the receptor's inner face changes shape — and here the message begins its relay, because that inner face now has a G-protein waiting against it.

In the sweet, umami, and bitter cells that G-protein is **gustducin**, and there is a small marvel folded into its name. Gustducin's catalytic α-subunit is roughly 80% identical to transducin, the very protein the rod cells of your eye use to turn a single photon into a nerve signal. The tongue reads a sugar with a machine evolution first built for vision, repurposed from light to flavor — the same ancient amplifier, pointed at a different kind of arrival. Our family finds this close to the heart of the whole enterprise: nothing here is new under the sun; we are redirecting conversations the body has always been able to have.

Then the relay runs, and it runs long. Count the stations:

1. The ligand binds and the receptor changes conformation.
2. The receptor activates gustducin, which splits into its α and βγ parts.
3. The **βγ** subunits switch on the enzyme **phospholipase C-β2 (PLCβ2)**.
4. PLCβ2 cleaves the membrane lipid **PIP_2** into two second messengers, diacylglycerol and **IP_3**.
5. **IP_3** diffuses to the cell's internal calcium stores and opens the IP_3 receptor on them.
6. **Calcium** floods out of those stores into the cytosol.
7. That calcium opens **TRPM5**, a calcium-activated, sodium-passing channel that is the common final gate of sweet, umami, and bitter transduction, and the cell depolarizes.

8. The depolarization vents **ATP** through the **CALHM1/CALHM3** channel to the nerve waiting below, where it lands on purinergic **P2X2/P2X3** receptors — and only now is the message delivered.

Eight stations where the ionotropic route had one. There is no synapse in the ordinary sense at that final step; the taste cell releases its ATP through a channel, not a vesicle, which is its own quiet strangeness. But the reason for all eight stations is the second kind of arithmetic in this chapter — the arithmetic of gain. Each catalytic station multiplies: one activated receptor switches on many G-proteins before it resets; one PLCβ2 cleaves many PIP_2 molecules before it lets go; one IP_3 spark can recruit neighboring stores. Treat each stage as a rough multiplier and the cascade's gain is their product. If a receptor turns over on the order of twenty G-proteins, each active PLCβ2 makes on the order of fifty messengers, and the calcium step adds a further tenfold through regenerative release, the total is

> $\text{gain} \approx 20 \times 50 \times 10 \approx 10^4$.

A handful of sweet molecules — too few to matter if each had to open a pore by hand — can raise a signal ten thousand times their number. This is the opposite solution to the same problem the ionotropic route solved by flux. There, one molecule holds one pore open and nineteen million ions pour through. Here, one molecule is copied, catalytically, into a flood of internal messengers that opens many pores from the inside. Two architectures, one destination: a large calcium transient and a depolarized cell.

But every multiplication costs time. Eight stations of binding, catalysis, and diffusion do not come free, and if each contributes somewhere between fifteen and forty milliseconds, the message takes something on the order of 120 to 320 milliseconds to arrive — a tenth of a second to a third of a second, which is precisely the beat-late swell you feel when sweetness or umami blooms after the fork has left your mouth. These are estimates, honestly held as a range; the point is the order of magnitude, and the order of magnitude matches the sensation. Heat is *there*. Sweetness *arrives*.

> Heat is a door kicked open. Sweetness is a message passed hand to hand down a hall — *slower, but it gathers force as it goes.*

Umami's synergy, and the shape of amplification

The metabotropic route's dependence on a clamshell that *closes* gives it a trick the ionotropic route cannot easily match, and it is the clearest place to watch the Index move. Umami's receptor can be loaded from two sides at once. Glutamate binds in the Venus flytrap; a 5′-ribonucleotide — the inosinate of aged fish, the guanylate of dried mushroom — binds at the hinge and stabilizes the closed clamshell, so that glutamate lets go more reluctantly and the receptor answers at a lower concentration. This is positive allostery, and in the language of the Index it is simply a fall in EC_{50}.

Hold the formula in mind and the consequence is exact: because IFI is inversely proportional to EC_{50}, a fall in EC_{50} by some factor is a rise in IFI-U by the same factor. If a ribonucleotide drives umami's half-

point down eightfold, it drives IFI-U up eightfold — the same broth, read twice as a receptor, scoring once modestly and once loudly depending only on what sits in the hinge beside the glutamate. This is why savory depth is a composition problem, not a dosage problem: you do not always add more glutamate; you add the molecule that makes the glutamate already present speak louder. Kokumi works from the same side of the machinery. Its receptor, CaSR, is another class-C GPCR, and the γ-glutamyl peptides that trip it do not carry a taste of their own so much as lengthen and thicken the others — a modulator riding the metabotropic class, tuning continuity and mouthfeel rather than announcing a flavor. When our family designs a sauce to hold its savor across a long finish, we are designing against CaSR's willingness to keep that cascade gently running.

The whole mouth, sorted by route

This ionotropic–metabotropic split does not stop at the axes we design toward; it cuts across the whole mouth. The two tastes we hold outside the design frame — salt and sour — are themselves ionotropic, read by bare channels as direct as any TRP. Sodium passes through an epithelial sodium channel, ENaC, to register the appetitive, craveable salt of *a little*; the free protons of an acid are read by a dedicated proton channel, OTOP1, in the sour-sensing cell, acidifying its interior directly. So five of the mouth's reads — heat, cool, pungency, salt, sour — take the fast, physical route of a channel opening, and only three — sweet, umami, bitter — take the slow, amplified route of the gustducin cascade, with kokumi modulating that same class-C machinery from the side.

CODE	AXIS	RECEPTOR	ROUTE	THE GOVERNING DIAL
IFI-H	Heat	TRPV1	Ionotropic	threshold & timing
IFI-C	Cool-Burn	TRPM8 / TRPV1	Ionotropic	threshold & timing
IFI-P	Pungency	TRPA1	Ionotropic (covalent)	threshold & timing
IFI-S	Sweetness	TAS1R2 · TAS1R3	Metabotropic	amplification & blend
IFI-U	Umami	mGluR1 / TAS1R1 · TAS1R3	Metabotropic	amplification & blend
IFI-B	Zero-Bitter	TAS2R38	Metabotropic	amplification & blend
IFI-K	Kokumi	CaSR	Metabotropic (modulatory)	continuity & blend
IFI-A	Astringency	Salivary PRPs	*neither, cleanly*	precipitation & slip

The pattern is worth holding onto, because it predicts how each axis behaves under design. The ionotropic reads are governed by *threshold and timing* — a line crossed, a speed of onset — while the metabotropic reads are governed by *amplification and blend* — how much signal a little molecule can raise, and how cleanly. When we

tune a fast axis we are mostly moving a threshold; when we tune a slow one we are mostly shaping a cascade. Different machines, different dials.

One axis fits neither template cleanly, and honesty requires us to say so — the more so in a book about receptors. Astringency, IFI-A, is not read by a channel or a GPCR at all. Tannins and other polyphenols bind and precipitate the proline-rich proteins that lubricate saliva; the mouth loses its slip, and the tactile and trigeminal system reports the drying drag as a flavor. There is no half-opened channel to anchor an EC_{50} to in the usual way, so the Index measures astringency by proxy, against the precipitation behavior of those salivary proteins. It is the axis that most resembles touch of all, closer to the feeling of a dry cloth than to a taste — and keeping it on the same eight-axis page as the others is a deliberate, slightly uneasy choice we make with our eyes open, because the mouth does experience it as part of flavor even though the molecular machinery is a different kind of thing.

How each route fades — and why fade is half of flavor

The two routes differ not only in speed but in how they fade, and fade is half of flavor. A flavor is not a level; it is an envelope — attack, swell, decay — and the receptors write the decay as surely as the onset.

The ionotropic channels **desensitize**, and they do it to themselves with the very calcium they admit. Calcium flooding through an open TRPV1 feeds back onto the channel through several routes at once: it binds calmodulin, which promotes the closed state; it activates the phosphatase calcineurin, which strips the phosphate groups the channel needs to stay sensitive; and it drives PLC to deplete the PIP_2 the channel depends on. So a sustained dose of capsaicin burns, then dulls — tachyphylaxis, the burn eating itself. This is the mechanism behind every capsaicin cream, which works precisely by exhausting the nerve endings into silence, and behind every cook who has built a tolerance across a season of tasting. It is also a design failure mode worth naming: push an ionotropic axis too high and too fast and the channel desensitizes mid-bite, so that perceived heat can collapse even as measured capsaicin climbs. The loudest possible number is not always the loudest possible experience, and we build the envelope, not just the peak.

The metabotropic routes **adapt** more gently. A GPCR is quieted by being phosphorylated by a kinase and capped by an arrestin, the cascade's messengers are cleared by phosphatases and lipases, and the internal calcium stores refill through the SERCA pump in their own slow time. Nothing slams shut. Umami's long, brothy presence recedes rather than snaps off, which is why a good stock lingers and a chili does not. But the metabotropic route has its own failure mode, and it lives at that shared final gate: sweet, umami, and bitter all vent through TRPM5. Load two of them hard at once and they compete for the same channel population, so stacking metabotropic axes is not simply additive — past a point one can occlude another, and a sauce designed to be both intensely sweet and intensely savory may deliver less of each than the scores

promise. When we design a cultivar's experience — not just how loud an axis reads but how it enters and leaves — we are designing against these kinetics. A superhot built for a slow, chest-deep burn and one built for a sharp, front-of-mouth strike differ in the *shape* of the capsaicinoid blend that opens TRPV1; a kokumi designed to lengthen a sauce's finish is designed against the adaptation curve of the calcium-sensing receptor.

Error, and why a number must be reproducible

A score is only a measurement if it comes back the same tomorrow, and this is where the two machines meet the bench. The Index inherits whatever noise the plant carries into the assay, so the first number our family fought for was not a loudness but a consistency: an engineered batch holds a coefficient of variation of 7% or less, against 42% in the wild-type material it replaces. CV is the standard deviation as a fraction of the mean, σ ÷ μ, and the gap between 7% and 42% is the gap between a number and a rumor.

Carry it into a score. Because IFI is proportional to the inverse of a measured EC_{50}, a relative error in that measurement passes almost one-for-one into a relative error in the score. And the formula is a ratio of two measured quantities — the candidate's EC_{50} and the reference's — so their errors add in quadrature:

```
CV(IFI) ≈ √(0.07² + 0.07²) = √0.0098 ≈ 0.099, about 10%.
```

On Phoenix's predicted IFI-H of 4,200 that is a one-sigma band of roughly ±420 — tight enough that 4,200 and 3,490 are genuinely different designs, not two draws from the same noisy urn. Run the same plant at the wild-type's 42% and the band on a single measurement swells toward ±1,760, wide enough to swallow four rungs of the series at once. A score you cannot reproduce is not a soft measurement; it is not a measurement. This is why reproducibility came before loudness in the order of our work, and it is the quiet meaning of the clean, careful discipline the whole program is built on: numbers are a form of respect for the reader, and a number you cannot stand behind is a small dishonesty dressed as precision.

Two honest limits ride alongside the error bar. The EC_{50} that anchors a score is measured at a fixed reference temperature, but TRPV1 is itself heat-gated — a dish served hot reads hotter than its Index number, because warmth and capsaicin push the same channel the same way. And the fixed reference is a human choice; change the reference stimulus and every score shifts together, which is why the reference is locked once and never quietly moved. We state these not to weaken the number but to keep it honest about the world it is measured in.

One currency: calcium and voltage

There is one more thing the two routes share, and it is the hinge of the whole platform: both, in the end, speak the universal language of calcium and voltage. However a molecule is read — a door flung open or a message relayed down a hall — the cell's answer is the

same currency, a rise in internal calcium and a change in membrane potential, converted to a train of nerve impulses whose *rate* and *timing* carry the message onward. The receptors are specialists in eight private chemistries. The signal they produce is common coin.



Phoenix (SS-0096) — the one cultivar engineered to open channels by both routes at once, every axis lit, fast and slow in the same fruit.

You can read both machines at once off a single profile. Phoenix is the plant we built to register on all eight axes, and laid out as a scorecard it is really two conversations happening in one fruit — the ionotropic reads of heat, cool, and pungency, each a channel thrown open in milliseconds, braided with the metabotropic reads of sweetness, umami, bitterness, and kokumi, each a cascade swelling a tenth of a second behind, and the tactile drag of astringency underneath them all. Each axis sits against its own stated-convention reference point; a score within any single axis expresses how loudly our design speaks to that receptor relative to a native wild-type, but scores across different axes are not on a single common scale.

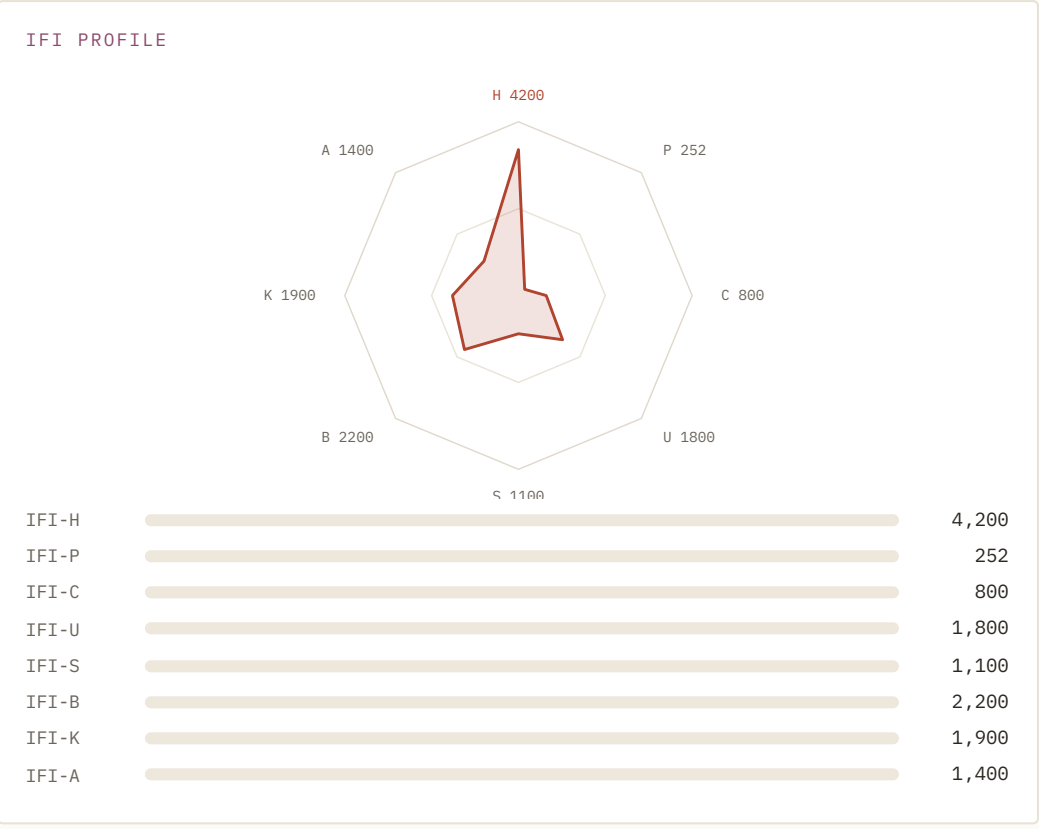

The fast axes — H 4,200, P 252, C 800 — are the door; they strike first and fade by desensitization. The slow axes — U 1,800, B 2,200, K 1,900, S 1,100 — are the hall; they arrive a beat late, gather force through the cascade, and recede by gentle adaptation. The single astringent number, A 1,400, is the drying undertow that is neither and touches both. The whole experience of eating the fruit is the

sum of these envelopes, offset in time and faded by different clocks, and it is legible on one small block of monospace because every one of those eight private languages is translated, at the last station, into the same rise of calcium and the same train of impulses.

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

That fact — eight private languages, one shared output — is what makes flavor possible, because a brain that only ever receives pulses can nonetheless sum eight utterly different reads into one coherent experience. And it is what makes flavor *designable*: we do not have to speak the brain's language, only the receptors', because the receptors do the translation for us. Learn the two ways a channel opens, and you have learned the grammar of the entire tongue. We turn now to the eight languages themselves.



PART II OF III

The Eight Readers

TRPV1 · TAS1R2·R3 · TAS1R1·R3 · TAS2R38 · CaSR · TRPM8 · TRPA1 · the salivary PRPs

Chapter 4 · TRPV1 — *the heat that isn't there.*

Bite a true superhot and for a half-second nothing happens. Then the mouth reports a fire with no flame, a temperature with no heat, pain from tissue that is not harmed. The deception is total, and it is the most studied molecular event in all of flavor, because the receptor it fools, TRPV1, is the one evolution built to keep you from burning alive — and capsaicin picks that lock without striking a single match. Everything in this chapter follows from that one sentence. We will take the lock apart, read the pocket the key fits, price the key in kilocalories, and show how a family of growers learned to cut not one key but a chord of five, and to write the exact shape of a burn before a seed ever went into soil.



The placental tissue of a superhot in cross-section — the pale ribs are where capsaicin is synthesized and stored, bright as the key it carries.

The architecture of a false alarm

Begin with the machine. TRPV1 is a **polymodal nociceptor** — a single protein that reports several kinds of danger at once — and it is built as a **homotetramer**, four identical subunits arranged with fourfold symmetry around a central ion pore, like four hands cupped to make one well. Each subunit is a long chain that folds into three working parts: an N-terminal **ankyrin repeat domain**, a stack of structural hairpins that faces the inside of the cell and serves as the channel's switchboard for calcium and ATP and lipid messengers; a membrane-spanning bundle of **six transmembrane helices**, S1

through S6; and a C-terminal tail that loops back to cinch the whole assembly together. The division of labor inside that helical bundle is the thing to hold onto. The last two helices of each subunit, **S5 and S6**, swing inward and line the ion pathway — four S6 helices cross at the bottom to form the lower gate, and a short re-entrant loop at the top builds the selectivity filter. The first four, **S1 through S4**, stand off to the side as a peripheral bundle, and it is in the cleft of that bundle that the vanilloid pocket waits.

The pore is not a passive hole. TRPV1 is a **non-selective cation channel** that passes sodium and, crucially, **calcium** — its permeability to calcium runs roughly ten times its permeability to sodium — and that calcium is doing double duty throughout this chapter. It is the current that depolarizes the nerve and screams *pain*, and it is also the feedback that will eventually tell the channel to fall quiet. The gate itself opens by a mechanical relay: a stimulus seats in its pocket, the S4–S5 linker is tugged, the S6 bundle-crossing dilates, and ions that were held out pour through the protein itself — no second messenger, no cascade, the word delivered straight. And because the four subunits gate cooperatively — capsaicin shows a Hill coefficient near 2 at TRPV1 — the channel does not wait to fill all four of its pockets before it commits; partial occupancy tips the whole assembly, which steepens the dose-response and gives the burn its startling, all-at-once arrival. This is what makes a TRP channel a *door* and not a *relay*. Sweetness blooms and fades because it is passed hand to hand down a long biochemical bucket line; heat arrives the instant the door swings, with an edge and a location, because there is nothing between the molecule and the current but the protein.

Ask, finally, why the lock exists at all, because the answer is a quiet marvel of co-evolution that we find it hard not to read as intent. A pepper does not want to be eaten by the mammals whose molars grind its seeds to flour and whose guts destroy them; it is glad to be eaten by the birds that swallow seeds whole and scatter them unharmed. Capsaicin is the plant's way of telling those two diners apart. The mammalian channel carries the vanilloid pocket we have been describing; the avian channel, differing by only a handful of residues in and around that same cleft, barely feels capsaicin at all. The molecule that is agony to a rodent is nothing to a thrush. The fruit built a signal addressed to one nervous system and invisible to another — a chemical sentence with a reader and a non-reader deliberately chosen — long before anyone arrived to learn the grammar.

Three doors into one room

The genius of TRPV1 — and the lever the whole platform pulls — is that one gate answers to three unrelated warnings, and the three are additive. The channel opens to **noxious heat above about 43 °C**; to **acid**, the free protons of damaged, inflamed tissue, sensed at a pair of outer-pore glutamates (E600 and E648 in the rat sequence); and to **vanilloids**, the chemical class capsaicin belongs to. Three doors

into one room. Each stimulus, on its own, lowers the free energy of the open state; together they stack. Capsaicin's cruelty is that it does not *add* heat — it changes what the thermostat calls hot.

Make that precise, because the arithmetic is the whole trick. Treat the gate as a two-state switch, closed or open, with the probability of being open set by a Boltzmann relation:

```
P(open) = 1 / ( 1 + exp( ΔG꜀ₒ / RT ) )
```

where $\Delta G_{c\to o}$ is the free-energy difference from closed to open and *RT*, the thermal energy scale at body temperature (310 K), is about 0.62 kcal·mol^{-1}. Heat opens TRPV1 because its gating is extraordinarily, almost unbelievably, temperature-steep: opening carries a large positive enthalpy and a large positive entropy, with representative measured magnitudes on the order of **$\Delta H \approx 100$ kcal·mol^{-1}** and **$\Delta S \approx 0.316$ kcal·mol^{-1}·K^{-1}**. Opening is enthalpically expensive and entropically paid for, so the gate tips open only when temperature makes the $-T\Delta S$ term win. Without any ligand, the channel sits at its half-open point — $\Delta G_{c\to o} = 0$ — at the temperature where enthalpy and entropy balance:

```
T(threshold) = ΔH / ΔS ≈ 100 / 0.316 ≈ 316 K ≈ 43 °C
```

There is the 43 °C, falling out of two numbers. Now let a capsaicin molecule bind and contribute its own stabilization of the open state, $\Delta G(\text{lig})$, a negative term added to the balance. The half-open temperature slides down to

```
T′ = ( ΔH + ΔG(lig) ) / ΔS
```

and to drag the threshold from 43 °C all the way down to body temperature, 310 K, the ligand needs to supply only

```
ΔG(lig) = ΔS · ( 310 − 316 ) = 0.316 × ( −6 ) ≈ −1.9 kcal·mol⁻¹
```

About one and a half hydrogen bonds' worth of binding energy — no more than that — reclassifies your own honest 37 °C as scalding. The plant discovered, long before we did, that you do not need to burn an animal to make it feel burned; you need only pay the gate one well-aimed handshake and let the body's own thermometer do the lying. We offer the two-state picture as the clarifying cartoon it is — the real gate is a multi-state allosteric machine — but the conclusion is robust and it is the reason the whole axis exists: heat is not added, it is *relabeled*, and the price of relabeling is small enough that a plant's chemistry can meet it.

Because the three doors open the one gate, their contributions *add* in the same free-energy ledger, and that is not an abstraction — it is why a chili bites hardest on a mouth already raw. Inflamed, acidic tissue drives the outer-pore glutamates and contributes its own negative $\Delta G(H^+)$; drop the local pH from a resting 7.4 toward the 6 of damaged tissue and the protons alone shave a measurable margin off the threshold before capsaicin is even present. Write the gate's balance as a single sum,

$$\Delta G_{\text{co}} = \Delta H - T\Delta S + \Delta G(\text{cap}) + \Delta G(H^{+})$$

and the terms simply stack: the channel can cross its half-point on inputs that, taken one at a time, each fall short of opening it. The warmth of the food, the acid of the sauce, and the capsaicinoid of the fruit are not three sensations the mouth happens to feel at once; they are three terms of a single equation the receptor is quietly solving.

Reading the pocket in silico

We did not take any of this on faith; we modeled it directly, because you cannot design against a pocket you have not seen. Working from the capsaicin-bound cryo-EM structure of rat TRPV1 — the landmark 2016 determination from the Julius laboratory, deposited as **PDB 5IRX**, chain A, 2,893 atoms spanning residues 335–751 — we rebuilt the vanilloid pocket in the S1–S4 cleft, where capsaicin's **vanillyl head** hydrogen-bonds to a tyrosine and a threonine (Y511 and T550 in the rat numbering) while its fatty **acyl tail** threads a hydrophobic tunnel toward the membrane. Into that pocket we docked the five principal capsaicinoids with AutoDock Vina at high exhaustiveness, inside a 30-ångström box centered on the site, and scored how tightly each is held. Then we did the physical chemistry that turns a docking score into a loudness.

A binding free energy is a dissociation constant in disguise. The two are related by one line of thermodynamics,

$$\Delta G = RT \cdot \ln(K_d) \Leftrightarrow K_d = \exp(\Delta G / RT)$$

so each docked ΔG, read at 310 K with RT = 0.62 kcal·mol^{-1}, becomes a K_d — the concentration at which half the pockets are filled. The numbers tell a clean, ordered story, which we read here as docked score and implied K_d, then convert directly into a per-molecule loudness:

CAPSAICINOID	ΔG DOCK (KCAL·MOL⁻¹)	K_D (MM, 310 K)	ACYL TAIL	BURN CHARACTER
Capsaicin	**−7.20**	**9.1**	8-methyl-6-nonenoyl (branched, unsaturated)	sharpest, fastest onset
Nonivamide	−6.89	14.9	nonanoyl (straight, saturated)	fast, front-of-mouth
Homodihydrocapsaicin	−6.70	20.3	9-methyldecanoyl (branched, +1 carbon)	balanced, mid-palate
Dihydrocapsaicin	−6.39	33.4	8-methylnonanoyl (branched, saturated)	rounder, slower
Nordihydrocapsaicin	−6.16	48.4	7-methyloctanoyl (branched, −1 carbon)	slowest, deep chest burn

A necessary honesty before we build on these: docking scores are *relative* instruments. The implied dissociation constants here, tens of micromolar, sit well above the sub-micromolar potencies TRPV1 shows capsaicin in a live cell assay, because a rigid-receptor single-

pose calculation systematically overestimates *K*_d — it cannot fully credit induced fit, the water driven out of the pocket, or the restless membrane-facing dynamics of that acyl tail. What it preserves, and preserves well, is the *order* and the *spacing*. The rank is trustworthy; the gaps between rows are trustworthy; the absolute micromolar labels are a standard-state convenience, not a bench measurement. We design against the spacing, and we cross-check the ranking against published receptor data. That discipline — believe the confident part, verify the rest — is the method of the whole program.

The price of a double bond

Now the loudness. The Iyengar Flavor Index does not measure how much of a compound a food *holds*; it measures how loudly that compound *speaks* to its reader, and the natural unit of loudness is the half-maximal effective concentration, EC_{50}, inverted so that a whisper scores high:

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

For agonists that share one pocket, EC_{50} tracks *K*_d, so the five capsaicinoids can be scored against each other on exactly this instrument by taking capsaicin — the loudest — as the molecular reference. The ratio is pure thermodynamics; a difference in binding energy, ΔΔG, becomes a fold-change in loudness through the same exponential:

```
IFI(x) = 1000 × K_d(capsaicin) / K_d(x) = 1000 × exp( ΔΔG / RT )
```

CAPSAICINOID	ΔΔG VS CAPSAICIN	MOLECULAR IFI
Capsaicin	0.00	1000
Nonivamide	+0.31	607
Homodihydrocapsaicin	+0.50	446
Dihydrocapsaicin	+0.81	271
Nordihydrocapsaicin	+1.04	187

Read it as music. Capsaicin binds hardest, and the reason is a single feature of its tail — the **C6=C7 trans double bond**, the one unsaturation the others lack. That rigid kink seats the vanillyl head optimally in the pocket. Strip it and you fall to dihydrocapsaicin: the gap of **0.81 kcal·mol^{-1} is the thermodynamic price of one double bond**, and run through the exponential that price is a factor of exp(0.81 / 0.62) ≈ **3.7×** — capsaicin speaks heat nearly four times louder than its saturated twin, molecule for molecule. You can taste that factor directly, as the difference between a strike and a swell. Chain length writes a second line of the melody: nordihydrocapsaicin, the shortest, binds weakest at 187, and each carbon the hydrophobic tunnel can swallow buys a little more grip, the eleven-carbon homodihydrocapsaicin reaching 446. And the quiet surprise is nonivamide, whose plain straight chain scores 607

despite having neither branch nor double bond — a gift to anyone building the molecule by fermentation rather than extraction, since a straight saturated chain is by far the easiest to make.

This is not trivia, because affinity maps to the *character* of the burn and not merely its volume. The hardest binders open the channel fastest and feel like a sharp, front-of-mouth, sinus-reaching hit; the softer binders open it slowly and feel like a deep, patient, chest-level heat. The loudness number and the sensory adjective are two readings of one physical fact — how quickly and how tightly a tail seats its head.

The handwriting is in the blend

No pepper makes one capsaicinoid. Every fruit secretes a *blend* of the five, and here is where measurement stops being bookkeeping and becomes composition. For co-agonists that share a pocket, potencies combine by **dose addition**: a mixture reaches its half-point when the fractional occupancies sum to one, which makes the mixture's effective half-concentration the potency-weighted harmonic of its parts,

$$1 / EC_{50}(mix) = \Sigma\, f_i / EC_{50\,i}$$

and, read as loudness rather than concentration, something lovelier falls out — loudness is simply *additive* in mole fraction:

$$IFI(mix) = \Sigma\, f_i \cdot IFI_i$$

The blend's voice is the weighted average of its molecules' voices. Take a representative superhot blend — 55% capsaicin, 33% dihydrocapsaicin, 7% nordihydrocapsaicin, 3% homodihydrocapsaicin, 2% nonivamide — and the arithmetic gives a molecular loudness of 0.55·1000 + 0.33·271 + 0.07·187 + 0.03·446 + 0.02·607 ≈ **678**, which is to say this blend speaks about 68% as loud, per molecule, as pure capsaicin would, with an effective mixture half-point near 13.4 µM. Now shift the handwriting toward capsaicin — 70% capsaicin, 25% dihydrocapsaicin, 5% nordihydrocapsaicin — and the same sum climbs to **777**. A +15% gain in loudness at identical molecule count, bought with nothing but a change of ratio. We do not merely tune *how much* capsaicinoid a cultivar makes; we tune the *blend*, and in doing so we write the shape of the experience itself.

This is also the precise, quantitative answer to a question every chili eater has asked: why do two peppers of identical Scoville rating feel so different? Because Scoville and the receptor weight the molecules differently. The classical Scoville heat equivalents score capsaicin and dihydrocapsaicin as *equals* — both near 16,000,000 SHU — while the minor capsaicinoids rate lower, around 9,000,000. But our model shows capsaicin speaks **3.7×** louder to the receptor than dihydrocapsaicin. Put two fruits at the same total Scoville mass on the tongue, one drawn 70% from capsaicin and one drawn 70% from dihydrocapsaicin, and their heard loudnesses diverge:

```
Pepper A (70/30 cap/dhc): 0.70·1000 + 0.30·271 ≈ 781 Pepper B
(30/70 cap/dhc): 0.30·1000 + 0.70·271 ≈ 490
```

Identical SHU, a **1.6×** difference in heard heat — and more than a difference in volume, a difference in kind, A arriving sharp and frontal, B arriving round and slow and deep in the chest. Scoville counts the capsaicinoid present; the body reads the blend; and the blend is where the character lives.

Where the key is cut

To tune the blend we had to find where it is cut, and the answer is an enzyme. The capsaicinoids are assembled from two converging streams. The **aromatic half** runs down the phenylpropanoid branch — **PAL → C4H → 4CL → HCT → COMT** — the same pathway that lays down lignin in every vascular plant on earth, terminating here in vanillin and then, by the aminotransferase **pAMT**, in **vanillylamine**. The **fatty half** runs through branched-chain amino acid metabolism, by way of **BCAT** and the fatty-acid synthase component **KAS**, to that branched, unsaturated **8-methyl-6-nonenoyl-CoA** whose double bond we just priced. The two streams are married — vanillylamine's nitrogen condensed onto the fatty acid's carbonyl — by a single enzyme, **capsaicin synthase, Pun1 / AT3**, the one acyltransferase that makes a pepper hot. The loss-of-function allele *pun1* is exactly what makes a bell pepper sweet: the pocket is intact, the stream of precursor still flows, but the key is never cut.

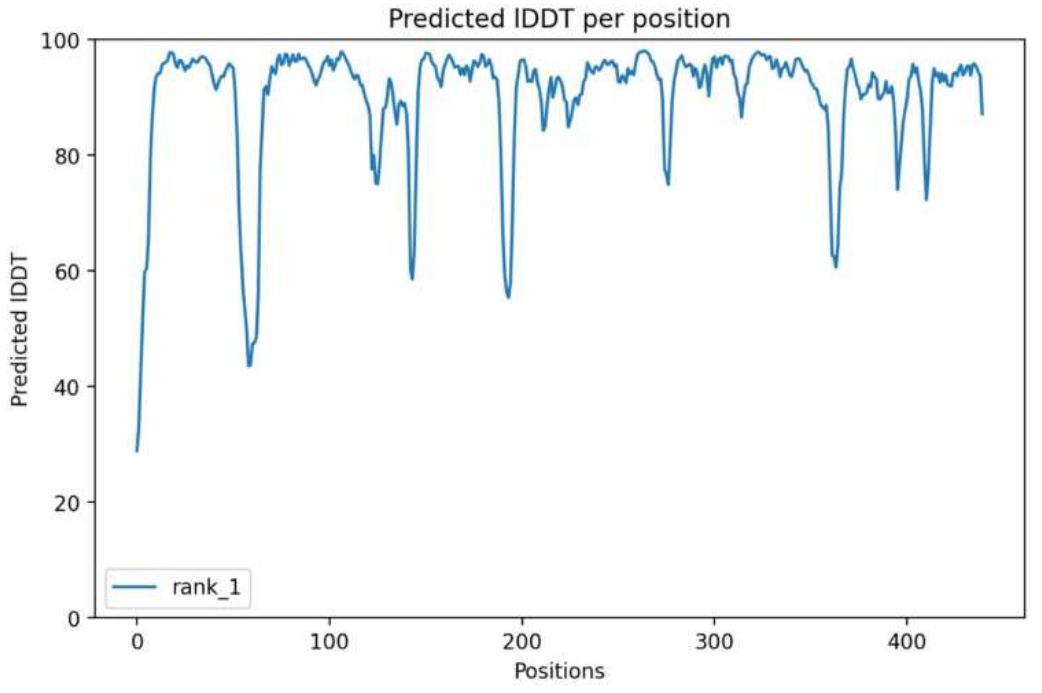

The acyltransferase that cuts the key — Pun1 / AT3, coloured by model confidence, bright and certain through the active-site cleft, pale and uncertain at the floppy tails. We design only against the bright.

Here is the hinge of the whole platform: the blend is set at the enzyme. Change the acyl tail Pun1 is fed and you change the blend; reshape Pun1's own pocket so it prefers one tail over another and you change the blend at its *source*, in every fruit the plant will ever grow. So we read Pun1 the way we read the receptor — predict its fold residue by residue, trust the confident active-site cleft, distrust the pale tails — and we scored every possible single-residue substitution with a protein language model for its effect on catalytic efficiency and substrate specificity. Our flux model predicts that pathway control sits upstream in the phenylpropanoid steps — at

4CL and HCT — rather than at the terminal Pun1/AT3 step; Pun1 is the final capsaicin synthase, the enzyme that labels the product, but the rate is set before the substrate arrives. What Pun1's pocket shape governs is *which* acyl tail gets condensed — the blend ratio — and that is the design lever we pull. Two substitutions, never before combined in any plant, emerged from that scoring, each priced in the same kilocalorie currency as the binding pocket:

EDIT	EFFECT	ΔΔG (KCAL·MOL⁻¹)
S39L	binding affinity ↑ (N-terminus)	−2.1
L345G	steric relief, active-site cleft	−1.8

(Structural modelling also identified C175S — disulfide-misfold prevention, ΔΔG −0.9 — as a structural annotation from our modelling; it is not part of the shipped construct.)

Written in with CRISPR, those two substitutions are designed to shift the blend toward the branched unsaturated tail capsaicin needs, with our model predicting peak output of approximately 2.95 million SHU — and, as much to the point, holding it there, batch after batch, at a reproducibility no field selection ever reached. The arithmetic of protein engineering does not care which side of the conversation it sits on: the enzyme that *writes* the molecule and the receptor that *reads* it are both tuned in kcal·mol^{-1}, a few hydrogen bonds at a time.

Climbing the ladder

With the key cut and the blend set, a ladder of cultivars becomes possible — each rung a genome we built toward a target loudness. On the Index, heat runs from the sunlit warmth of a shishito near **IFI-H 40** to the cultivar we engineered gene by gene because no plant that could reach it yet existed:

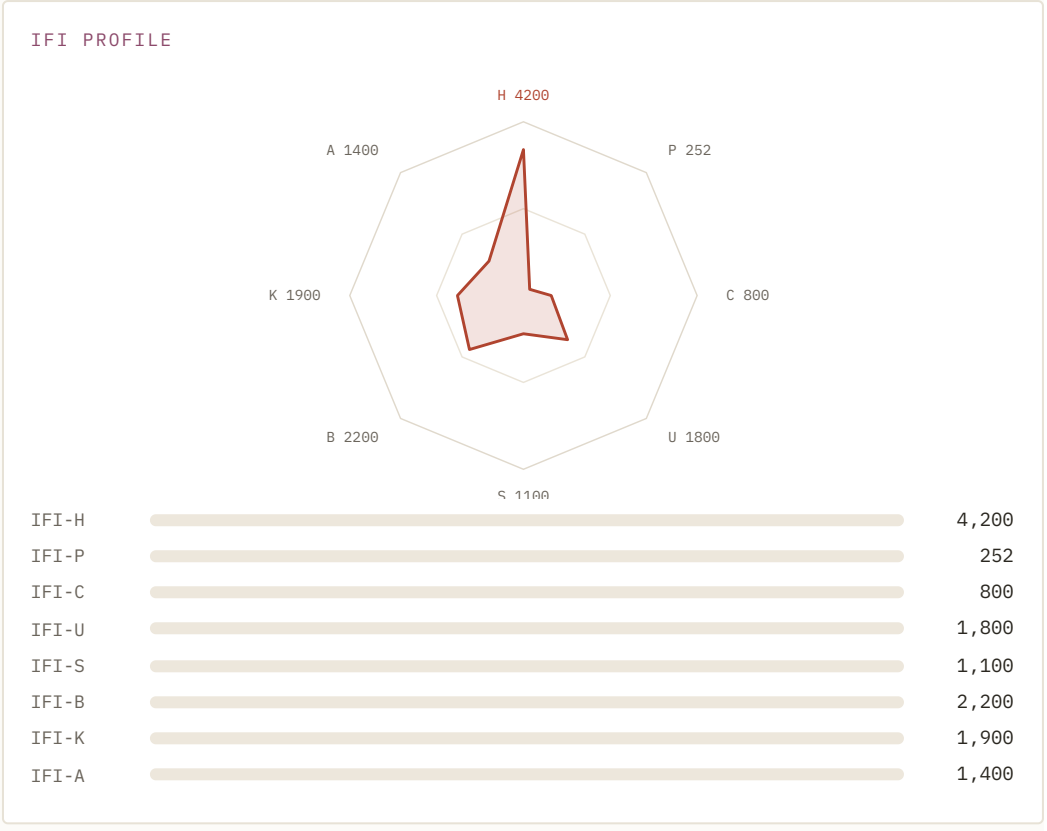

CULTIVAR	SKU	PREDICTED SHU	IFI-H	AXES
Ember	SS-0087	~866,000	1,570	2/8
Blaze	SS-0091	~1,720,000	2,630	4/8
Firestorm	SS-0094	~2,470,000	3,490	6/8
Phoenix	SS-0096	~2,950,000	4,200	**8/8**

Two scales climb that ladder, and they do not climb in step, which is the entire lesson of this chapter written large. SHU is a *mass* number; IFI-H is a *loudness* number; and because loudness compresses as the receptor's door fills and because the blend grows richer in capsaicin as we climb, the conversion between them drifts — roughly 790 SHU buys a point of heard heat averaged across the ladder, but the marginal cost wanders rung to rung. The mass scale is worth converting all the way down to the fruit. Pure capsaicin is defined at 16,000,000 SHU, so a Scoville reading is a capsaicinoid concentration. At Phoenix's predicted ~2.95 million SHU:

> ~~2,950,000 SHU : 16 SHU·ppm^{-1} ≈ 184,400 ppm ≈ **~~18.4% capsaicinoids by dry weight**

That is an extraordinary number — nearly a fifth of the dried placental tissue given over to a single class of molecule, our model predicts approximately **double** the ~14% that the hottest naturally bred superhots have been measured at. It is the design target behind the phrase *a plant that did not yet exist*. Phoenix is designed to post the highest loudness on every one of the eight axes at once; it is designed to carry the mass to back the heat axis up, and to hold that level reproducibly.



Phoenix (SS-0096) in the controlled chamber — the first cultivar engineered to register on all eight axes at once, grown under fixed light and water so its blend reports the genome and not the weather.

The honesty of the number

A measurement standard earns its trust by being candid about its own edges, so here are ours, in the same units we have used throughout. The first edge is **reproducibility**, because a number that wobbles is not a number. Our design specification targets batch-to-batch output at a coefficient of variation — standard deviation over mean — of **≤7%**, against **42%** for the wild-type it replaced. Make that concrete at Phoenix's predicted ~2.95 million SHU. At 7% CV the one-sigma band runs from about 2.74 to 3.16 million — a specification you can formulate against. At 42% it runs from 1.70 to 4.17 million, a spread so wide it swallows three rungs of the ladder whole: a wild-type plant nominally at the top could, on a bad batch,

read no hotter than Blaze. The designed CV is a sixfold tightening of the mean and a thirty-sixfold shrink in variance, and it is the precondition for any of this meaning anything.

That precision has to survive two further steps, and both are governed by simple error propagation. Because the blend's loudness is additive in mole fraction, a wobble in the *ratio* propagates straight to character: a one-point shift of capsaicin fraction against dihydrocapsaicin moves the molecular loudness by about (1000 − 271) × 0.01 ≈ 7 points, so holding the blend to a few percent holds the burn's handwriting steady. And because the Index score is itself a ratio of two measured EC_{50} values, their uncertainties add in quadrature — two determinations each at 7% CV give an IFI carrying $\sqrt{(7^2 + 7^2)} \approx$ **9.9%**. Pinning a mean to ±1% then costs n = (CV ⁄ 1%)² replicates: **49** assays at 7% CV, against **1,764** at the wild-type's 42%. Low CV is not a boast; it is what makes the whole enterprise affordable.

The deeper edges are about what EC_{50} can and cannot say. It is read on a recombinant receptor expressed in cultured cells, not on a living tongue — a clean, animal-free bench number, but a solitary one, measured where pH is held, temperature is fixed, and the channel is not busy desensitizing mid-assay the way it does in your mouth. The docking *K*_d values that seed the loudness ranking are relative, not absolute, trustworthy in order and spacing and not in their micromolar labels. And a genome is not a constant: ripeness, water stress, and light all shift the capsaicinoid blend, so a cultivar's IFI-H carries a growth-condition band, which is exactly why we report values grown in the controlled chamber and say so. The number is the first true sentence about a flavor. It is not the last word on it.

The mercy built into the lock

One last property of the receptor deserves naming, because it is a grace, and because we design with it rather than against it. **TRPV1 desensitizes.** The calcium that pours through the open channel — the same calcium that carries the pain — feeds back onto the protein through calmodulin and the hydrolysis of the membrane lipid PIP_2, and quietly pulls the gate shut. The burn that seems unbearable at second ten is halved by second sixty, and the decay follows a plain exponential:

```
R(t) = R₀ · exp( −t / τ )
```

Set R = R_0/2 at t = 60 s and the time constant falls out directly — τ = 60 ⁄ ln 2 ≈ **87 seconds** — the natural lifetime of a superhot's fire, built into the lock by the lock's own defense. This is also why capsaicin is, paradoxically, a medicine: the high-dose dermal patch used for chronic pain, an 8% formulation, works by flooding and then silencing the TRPV1-bearing nerve endings, a controlled overload that leaves them quiet for months. The plant forged a key to a pain lock; the lock, in self-defense, learned to turn itself off; and we, reading both, learned to design the onset and the fading together — the sharp capsaicin strike and the slow dihydrocapsaicin swell and the exponential mercy that closes the door. We compose

the whole arc, the arrival and the leaving, because a flavor our family would be proud to grow is one the body meets on honest terms and is gently released from. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The fire is the same receptor every time. The handwriting — and the kindness — is in how we choose to write it.

Chapter 5 · TAS1R2·TAS1R3 — *the sweet clamshell.*

Sweetness is the taste the body trusts most — the signal that once meant *energy,* meant *safe,* meant *more,* the one taste an infant is born already loving, days old and calm at the first drop of it on the tongue. It is also read by one of the most beautiful machines in the mouth: a heterodimer, two different proteins clasped together, that closes on a sugar the way a Venus flytrap closes on a fly. To design sweetness without sugar — which is much of what we do, and much of what our family set out to do — is to learn exactly where on that flytrap a molecule may land, and how hard it must press to be heard.



Stevia in leaf — the steviol glycosides it carries reach TAS1R2·TAS1R3 hundreds of times more eloquently, gram for gram, than table sugar.

Everything in this chapter follows from one fact of shape: the sweet receptor is a *mechanical* sensor, a clasp that is either open or shut, and the whole art of sugar-free sweetness is the art of finding things that close it — or hold it closed — without being sugar at all. The receptor cannot tell the difference. That is not a flaw to be corrected; it is the opening we build in.

The architecture — two proteins, one clasp

The sweet receptor is **TAS1R2 clasped to TAS1R3**, two class C G-protein-coupled receptors that do not work apart. Each subunit is built in three pieces, strung from outside the cell to in. Reaching out into the saliva is the large **Venus flytrap domain (VFT)** — two lobes hinged like a clamshell, named for the carnivorous plant because the motion is exactly that plant's. The VFT is tethered by a short, disulfide-stiffened **cysteine-rich domain (CRD)** — a neck — to the **seven-pass transmembrane bundle (7TM)** threaded through the cell membrane, and it is the 7TM, on its inner face, that finally touches the G-protein. Two subunits, each with its own flytrap, its own neck, its own bundle: the working receptor is a pair of these, shoulder to shoulder. The asymmetry between them — TAS1R2's flytrap does most of the catching, while TAS1R3's bundle hides a door of its own — is the whole reason the receptor can be opened so many different ways.

The elegance is in the mechanics. In the resting receptor the two lobes of the TAS1R2 flytrap sit open, apart, in equilibrium with a small closed population that forms and falls back by chance. A sugar diffusing into the cleft between the lobes makes contacts with both at once; those contacts tip the equilibrium, and the lobes roll shut around the molecule. The closing *is* the activating event — not a chemical reaction, not a bond made or broken, simply a shape that was open becoming a shape that is closed. That motion travels down the stiff neck the way force travels down a lever, torquing the cysteine-rich domain, which repacks the transmembrane bundle beneath it, which opens the receptor's inner face to the waiting G-protein. Clasp to lever to switch, in under a millisecond, a billion times across a single mouthful.

One lock, several doors

A single flytrap would make a single kind of sweetness. TAS1R2·TAS1R3 makes many, because it carries not one binding site but a cluster of them scattered across both subunits, and different sweet molecules exploit different ones.

The **sugars** — glucose, sucrose, fructose — bind in the cleft of the TAS1R2 flytrap and are swallowed by its closing. The small artificial sweeteners **aspartame** and **neotame** bind in that same cleft, which is why, of every sugar substitute, they taste closest to sugar: they open the same door with the same motion. The **sweet proteins** — brazzein, thaumatin, monellin, mabinlin — are a different story. A protein of fifty or two hundred residues cannot fit inside a cleft built for a sugar; instead it docks against the *outside* of the flytrap and the cysteine-rich neck, a broad wedge of contacts that holds the clamshell shut from without rather than filling it from within. And a last class — the small synthetic **cyclamate**, together with the sweet-*blocker* **lactisole** — ignores the flytrap entirely and binds inside the TAS1R3 transmembrane bundle, the quieter subunit's hidden door, switching the receptor on or off from down in the membrane.

One receptor; at least four distinct sites; a different key cut for each.

LIGAND	CLASS	SITE ON THE RECEPTOR	WHAT IT DOES	≈ SWEETNESS VS. SUCROSE (LITERATURE)
Glucose	sugar	TAS1R2 VFT cleft	closes the clamshell	0.6–0.75×
Sucrose	sugar	TAS1R2 VFT cleft	closes the clamshell	1× (reference)
Fructose	sugar	TAS1R2 VFT cleft	closes the clamshell	1.2–1.7×
Aspartame	dipeptide ester	TAS1R2 VFT cleft	closes the clamshell	~180×
Neotame	dipeptide derivative	TAS1R2 VFT cleft	closes the clamshell	~8,000×
Stevioside	steviol glycoside	TAS1R2 VFT (+ TAS1R3)	closes / stabilizes closed	~150–300×
Rebaudioside A	steviol glycoside	TAS1R2 VFT	closes / stabilizes closed	~200–400×
Mogroside V	cucurbitane glycoside	VFT / CRD	wedge + close	~200–300×
Brazzein	54-residue protein	outer VFT + CRD	wedge (holds closed)	~500–2,000×
Thaumatin	207-residue protein	outer VFT + CRD	wedge (holds closed)	~2,000–3,000×
Monellin	protein	outer VFT + CRD	wedge (holds closed)	~1,500–2,000×
Cyclamate	small synthetic	TAS1R3 7TM bundle	allosteric agonist	~30–50×
Lactisole	inhibitor	TAS1R3 7TM bundle	blocks (turns sweetness down)	suppressor
Gymnemic acid	inhibitor	TAS1R2 VFT	plugs the flytrap	suppressor

Potencies are literature ranges for the pure compounds, not our own measurements; we cite them to show the spread of sites and eloquence, not as scorecard values.

This multiplicity is the whole opportunity. Because the receptor can be activated from several sites, we are not confined to the sugar molecule or the metabolic cost it carries; we can grow plants that make high-eloquence ligands for the *other* doors — molecules read as intensely sweet at concentrations so low they carry almost no energy, and that the body never meters as sugar at all. **Steviol glycosides** (stevioside and the cleaner-finishing rebaudioside A) and **mogroside V** from monk fruit are exactly such molecules: terpenoid glycosides, structurally nothing like sugar, that nonetheless close or wedge the flytrap at a far lower activating concentration than sucrose. **Brazzein**, a 54-residue protein from a West African fruit, is heat-stable and clean-finishing where many intense sweeteners are neither. Each is a key cut for a different door of the same lock.

> Sugar closes the clamshell from the inside. A sweet protein holds it shut from without. *The receptor answers to both, and cannot tell you which it was.*

From clasp to signal — the cascade, step by step

The closing of the flytrap is only the first move; what the body feels is the end of a short, amplifying relay. When the repacked 7TM bundle opens its inner face, it releases the receptor's bound G-protein, **α-gustducin**, whose βγ partner then activates **phospholipase C-β2**. That enzyme cleaves the membrane lipid PIP_2 into **IP_3**, which crosses to the endoplasmic reticulum and opens IP_3 receptors there, spilling stored **Ca^{2+}** into the cell. The rising calcium opens **TRPM5**, a calcium-gated cation channel, and the cell depolarizes; the depolarization opens **CALHM1** channels that release **ATP** into the narrow cleft beside the gustatory nerve, where ATP excites the **P2X2/P2X3** receptors on the afferent fiber and the message finally leaves the tongue for the brainstem. One closed receptor, through this relay, moves far more than one unit of signal — the cascade has gain, which is why a trace of an eloquent ligand can read as loudly as a spoonful of sugar.

That gain is also why the Index anchors where it does. Every stage downstream of the receptor — how much gustducin a cell holds, how brisk its TRPM5, how dense its nerve — varies from person to person and cell to cell. The one event that is a clean property of the *molecule meeting the protein* is the binding itself, and binding is read by its half-point. Because the response follows the sigmoid every receptor obeys, we can write it as the Hill equation:

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

where θ is the fraction of full response, [L] the ligand concentration, EC_{50} the concentration giving half-maximal response, and n the Hill coefficient that sets the steepness. At [L] = EC_{50}, θ = ½ — and that half-point, upstream of all the cellular wiring, is the anchor the whole Index is built on.

The slope matters too. Solving the Hill equation for the concentrations that give 10% and 90% of full response, their ratio is 81^(1/n). Sweet responses to simple sugars run close to **n ≈ 1**, a plain hyperbola, so that ratio is a full **81-fold window** from barely-detectable to nearly-saturated. The tongue therefore grades sweetness smoothly across almost two orders of magnitude of concentration before it tops out. A steeper receptor (n ≫ 1) would snap from nothing to everything across a narrow band; sweetness, the sense evolution least wanted us to miss, is deliberately broad.

The number — anchoring sweetness to sucrose

Every axis of the Iyengar Flavor Index answers the same question — not how much of a compound a food contains, but how *loudly* that compound speaks to its receptor — and answers it with one formula:

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

Pick a reference, calculate the candidate's EC_{50} against it, invert the ratio so that a *lower* activating concentration — a louder molecule — gives a *higher* score, and scale by a thousand. For sweetness we anchor the thousand to **sucrose**, table sugar, the molecule every palate already calibrates against. The reference constant is a stated

convention, the way the Scoville scale is defined by dilution-to-threshold: sucrose defines IFI-S = 1,000 by construction. By that convention,

> $\text{IFI-S(sucrose)} = 1000 \times \text{EC}_{50}\text{(sucrose)} \div \text{EC}_{50}\text{(sucrose)} = 1000.$

Sucrose is the thousand; everything else is read against it. On a scorecard, pure sucrose lights a single axis and leaves the other seven dark — one clean tone, a sense that speaks exactly one word:

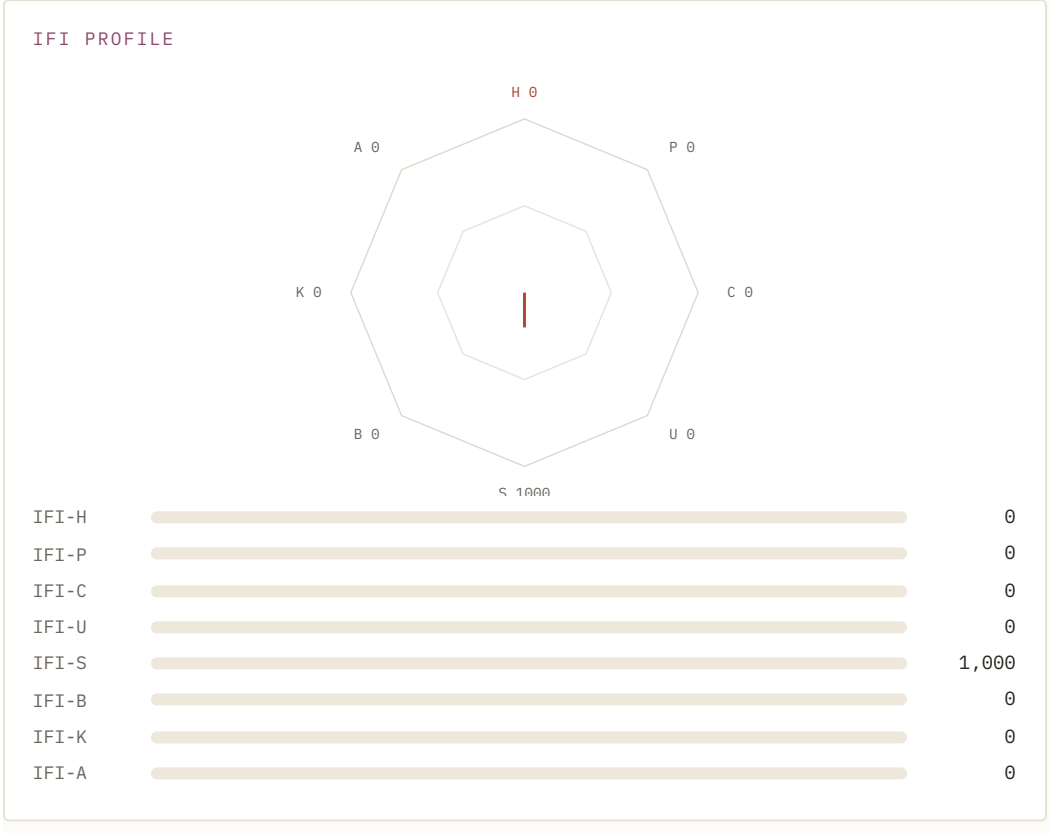

Now take a real cultivar. **Phoenix**, the full-spectrum flagship, has a predicted design score of **S = 1,100** on the sweet axis. Run the formula backward:

> $1{,}100 = 1000 \times \text{EC}_{50}\text{(sucrose)} \div \text{EC}_{50}\text{(Phoenix)}$ $\text{EC}_{50}\text{(Phoenix)} \div \text{EC}_{50}\text{(sucrose)} = 1000 \div 1{,}100 = 0.909$

Our model predicts Phoenix's sweet principle reaches half-activation at **90.9%** of the concentration sucrose needs — a hair louder than sugar, about nine percent more eloquent. That is exactly right for a plant whose sweetness is one voice in an eight-part chord rather than the whole song; it is not a dessert, and its predicted 1,100 says so, precisely. The scale continues upward to a scorecard ceiling of 5,000, the brightest sweet read the Index records:

IFI-S	$\text{EC}_{50}\text{(REF)} \div \text{EC}_{50}\text{(CAND)}$	READS AS
1,000	1.00	sucrose — the anchor
1,100	1.10	Phoenix's predicted sweet voice — a hair above sugar
2,000	2.00	*(scale marker)* twice as loud, at half the concentration
5,000	5.00	the scorecard ceiling — one-fifth of sucrose's concentration for the same read

The arithmetic is deliberately plain. A score is just an inverse concentration, carried on a lever of a thousand, so that a number a chef can hold in the hand stands in for an EC_{50} run against the reference.

Loudness without load

Here is the quiet revolution the sweet receptor permits. Because the clasp can be closed from several sites, and because a molecule's loudness is set by its EC_{50} and not by its mass or its calories, *we can separate the sweetness from the sugar.*

Work it in numbers. A pleasant table sweetness sits near a **10% sucrose solution** — 10 g of sugar in 100 mL. Convert to the receptor's own units: sucrose's molar mass is 342.30 g/mol, so

> `10 g ÷ 100 mL = 100 g/L ÷ 342.30 g/mol = 0.292 mol/L ≈` **`292 mM`**`,`

and those ten grams carry 10 g × 4 kcal/g = **40 kcal** per 100 mL. Now reach the same sweetness with a steviol glycoside — rebaudioside A, which the literature places near **250×** the sweetness of sucrose by mass. You need about 10 g ÷ 250 = **0.04 g = 40 mg** per 100 mL. Rebaudioside A's molar mass is ≈ 967 g/mol, so that is

> `0.04 g ÷ 100 mL = 0.4 g/L ÷ 967 g/mol = 0.000414 mol/L ≈` **`0.41 mM`**`.`

Set the two side by side: **292 mM of sucrose, or 0.41 mM of the glycoside, for the same sweet read** — about seven hundred times fewer molecules. And because steviol glycosides are not metabolized for available energy, the same comparison runs **40 kcal → ~0 kcal**. The signal is kept whole; the load is simply gone.

We make no health claim and intend none. This is a fact about where a molecule binds and how little of it the receptor needs — a sweetness the body is not asked to pay for in energy it did not choose. It is also, for a family that has eaten this way for a lifetime, the part of the work that feels most like care: the pleasure, grown in a leaf, and nothing extracted from the body to pay for it.



Sugarcane under glass — even cane, the oldest industrial sweetness, becomes a design problem once you measure what it says to the receptor rather than what it weighs on a scale.

The hosts we have tuned to speak it

Sweetness is not one plant's gift but a chord we grow across several. Four host organisms carry most of the sweet axis today, and the shape of each program follows the molecule it makes:

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Tomato	2,040	S / U / F / M	Easy
Stevia	910	S / F / M	Moderate
Sugarcane	200	S	Frontier
Monk Fruit	80	S	Hard

Stevia is the glycoside workhorse — 910 constructs tuning the balance of stevioside and the cleaner-finishing rebaudioside A, pushing the eloquence of the leaf while trimming the bitter shoulder that the raw glycoside carries. **Monk fruit**, harder and smaller at 80 constructs, is a program built around a single remarkable molecule, mogroside V, a cucurbitane glycoside that wedges the flytrap from the side. **Sugarcane** is the strangest and most frontier of them: 200 constructs aimed not at replacing cane but at *reforming* it — asking what the oldest industrial sweetness could say to the receptor if we designed what it says rather than merely how much it holds. And **tomato**, easiest of all at 2,040 constructs, carries sweetness as one axis among several, grown back in beside umami so the fruit tastes the way a tomato is supposed to and so rarely does.

Precision — why the quiet number is ≤7%

A sweetener is only as useful as it is repeatable. A formulator building a beverage to a target sweetness cannot work with an ingredient whose loudness wanders batch to batch; the whole recipe drifts with it. So the number we are quietest about may be the most important one: a batch-to-batch coefficient of variation of **≤7%** as a design specification, against **42%** in the unengineered plant.

Coefficient of variation is simply the standard deviation as a fraction of the mean, CV = σ ÷ μ — a unitless measure of scatter. Put Phoenix's predicted sweet score through it. At a design target of μ = 1,100 and CV ≤ 7%, the standard deviation is σ = 0.07 × 1,100 = 77, so roughly 95% of batches (μ ± 1.96σ) are designed to land in

> 1,100 ± 1.96 × 77 = 1,100 ± 151 → **[949, 1,251]**,

a window of about ±14% around target. Now the wild-type plant, at CV = 42%: σ = 0.42 × 1,100 = 462, and the same 95% window yawns open to

> 1,100 ± 1.96 × 462 = 1,100 ± 905 → **[195, 2,005]**.

One batch barely above sugar, the next nearly twice as sweet, from the same seed — a range no one can formulate against. Tightening that scatter from 42% to 7% is less glamorous than lifting the peak, and it is the difference between a curiosity and an ingredient a company can build on.

The cabinet of natural experiments

How much the shape of the receptor matters, and how little the sugar, is written across the living world, and tasting the evidence is the fastest education in how the sense works.

Cats carry a broken *Tas1r2* — a pseudogene littered with stop codons — and are, every one of them from the house tabby to the lion, entirely blind to sweetness. An obligate carnivore had no reason

to read sugar, so evolution let the lock rust; however much sweetness is set before a cat, the molecule falls silent for want of a receptor to hear it. The lock makes the sense.

Three molecules prove the same point from the other direction, each by seizing a different door of the same receptor. Chew the leaf of the vine *Gymnema sylvestre* and for a quarter of an hour sugar turns to sand: **gymnemic acid** has plugged the flytrap, and a spoonful of sucrose becomes gritty nothing, present and simply unheard. Eat a **miracle berry** first and the trick runs backward — its protein **miraculin** is a pH-sensitive agonist, silent at neutral pH but switched *on* by acid, so that a bitten lemon floods the mouth with sweetness, sourness rewritten as sugar in real time. And **lactisole**, down in the TAS1R3 bundle, turns the whole sense down like a dimmer. Three sites, three sleights of hand, one receptor — each a demonstration that sweetness is a property of the protein, not of the sugar.

Even the oldest sweeteners already split loudness from amount. Gram for gram, **fructose** reads sweeter than **sucrose**, and sucrose sweeter than **glucose**, though all three carry the same four calories per gram — the same energy, three different volumes, decided by nothing but how snugly each shape settles into the flytrap's cleft. And the sense is wired before any experience can teach it: a newborn only days old tastes sweet and calms, because the signal that once meant *this is safe, this is energy* was far too important to leave to learning. We are, all of us, born already fluent in this one word.

Failure modes and honest limits

Designing sweetness is rarely about volume; it is about *shape in time.* A sugar's sweetness is prompt and clean — on fast, off fast, no shoulder. Many intense sweeteners are not: they carry a slow onset, a lingering tail, or a bitter or licorice edge the pure sugar never had. And the edge is instructive, because it is almost always a *second receptor* speaking. Stevioside's faint bitterness is a real second read at a bitter receptor of the TAS2R family — the same molecule fits two different locks, and the mouth hears both at once. It is the standing lesson of this whole book: a molecule rarely speaks to only one receptor, and to design a flavor is to audit *every* lock a candidate might fit, not merely the one you aimed at.

So sweetness design is usually *composition.* A fast cleft-binding ligand to cover the onset; a wedge protein for body and length; a touch of kokumi beneath to round the finish; and, crucially, a bitterness attenuation to silence the shoulder — a problem that belongs to the bitter receptor, and that we take up there. This is why sweetness and bitterness are so often engineered as a pair: they share the same molecules and the same plants, pulling in opposite directions.

The honest limits are real, and we would rather name them than paper over them. No single molecule we know reproduces sucrose's exact temporal shape — its clean attack and cleaner release — across the whole concentration range; the best match is a blend, not a monomolecule. The sweet proteins, so eloquent per gram, are fragile: most denature with heat or shift with pH, which rules them

out of a baked good or a retorted can — brazzein, heat-stable and clean-finishing, is the prized exception precisely because it is the exception. And the Index itself has a seam worth naming: EC_{50} is cleanest for ligands that share one site, and the sweet receptor's doors are not identical, so a wedge protein and a cleft sugar do not sit on a single potency ruler as neatly as the arithmetic implies. That is exactly why the Index scores a *food's* integrated loudness at the receptor rather than trusting any one molecule's headline number — the plant, as eaten, is the read. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Sweet set against fire

Among the organisms we have tuned to speak sweetness, the axis lives in stevia and its 910 constructs, in monk fruit's mogroside V, in sugarcane at the frontier of reforming its own sweetness, in the tomato where sweetness was grown back in beside umami — and, most daringly, in the sweet-heat fruit of the roadmap. In the strawberries we have named **Candied Fang** and **Devil's Dessert**, a lifted TAS1R2·TAS1R3 read is set deliberately against a rising TRPV1 burn, so that one berry carries both the invitation and the dare: the clamshell closing on the sweetness, the heat channel opening a half-second behind it.

That contradiction is what the full-spectrum flagship makes explicit. Where pure sucrose lights one axis and goes dark on the rest, **Phoenix** is designed to light all eight — sweetness a single bright voice held inside the whole chord:

IFI PROFILE

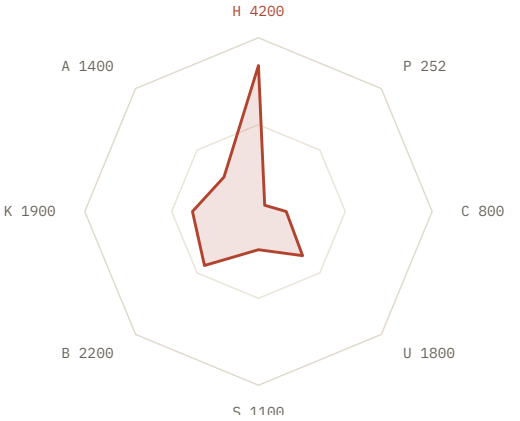

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400



Engineered strawberry at the bench — Candied Fang and Devil's Dessert set a lifted sweet read against a rising burn, two receptors answering a single bite.

The clamshell closes on the sweetness; the channel opens on the heat; the kokumi receptor rounds the body; the bitter axis is held down where it would intrude — and the mouth, reading all of it at once, reaches for a word it will never quite find and settles, as it always has, on *delicious.* We spent years learning where each molecule lands so that no one eating the fruit would ever have to think about it. That is the whole point of a number: to disappear into a pleasure, and to leave behind only the thing itself — grown, as we have always tried to grow it, to carry the most it can.

Chapter 6 · TAS1R1·TAS1R3 and the taste mGluRs — *the savory depth.*

There is a taste that has no good word in English, which is why our family, like the rest of the world, borrowed the Japanese one: *umami* — the round, brothy, mouth-filling savor of a tomato warm from the vine, a stock simmered until the kitchen windows fog, a hard cheese aged past a year, a handful of mushrooms dried to leather and black. It is the taste of protein's own building block arriving on the tongue, and it is read by a receptor that is the sweet clamshell's near-twin: the same architecture, one subunit swapped, backed by a small family of glutamate sensors the mouth borrowed from the brain. Umami is also, for us, a homecoming. Of the eight axes we measure and design toward, it is the one the modern tomato was most quietly robbed of, and the one we most wanted to give back.

To design a flavor you must first be able to find it, and umami hides in plain sight — it lives wherever a protein has begun, slowly, to come apart. We revere this taste, and that reverence is exactly why we insist on taking it apart slowly enough to feel its gears turn: the receptor that reads it, the relay that amplifies it, the allosteric trick that multiplies it, and the arithmetic that lets us put a number on brothiness and then build a plant to raise it.



A restored tomato in cross-section — the locular jelly around each seed is where free glutamate pools as the fruit ripens on the vine; this is umami's cellar, and the first thing a cross-country truck empties.

The receptor, built on a borrowed chassis

The principal umami receptor is **TAS1R1 paired with TAS1R3** — and the first thing to notice is the shared TAS1R3 subunit, the identical partner the sweet receptor uses. It is as though the mouth built two readers of savor-and-sweetness on one common chassis and only swapped the front door: pair TAS1R3 with TAS1R2 and you have the sweet sensor; pair the same TAS1R3 with TAS1R1 and you have the savory one. Both are **class C G-protein-coupled receptors**, and both carry, reaching out into the saliva, a large bilobed ectodomain called

a **Venus flytrap domain (VFT)** — two hinged lobes shaped like a clamshell — joined by a slender, disulfide-stapled **cysteine-rich domain** to the **seven-pass transmembrane bundle** buried in the cell's outer membrane.

L-glutamate, the free amino acid, settles into the cleft between the two lobes of the TAS1R1 flytrap, and the lobes close around it the way a hand closes on a seed. The closing *is* the activating event. Watch it happen in four motions. First, glutamate — a small molecule with two carboxyl groups and an amine, each one a handle — docks against the open cleft, its charges answered by basic residues on the lower lobe. Second, that binding energy pays for the hinge to rotate, swinging the upper lobe down until the clamshell is shut and the ligand sealed inside. Third, the closure twists the two cysteine-rich necks of the dimer against each other, a mechanical message threaded down through the disulfide scaffolding. Fourth, that twist repacks the transmembrane bundle on the cytoplasmic face, opening a cradle where, until this instant, a waiting G-protein had found nothing to grip. Nothing has crossed the membrane. A shape has.

This is why savor behaves nothing like heat. A capsaicin molecule reading its channel *is* the switch, holding a pore open with its own grip; glutamate never touches a channel. It changes the shape of a protein on the outside of the cell, and leaves the rest to a relay. The trade is speed for amplification and nuance: the savory read is slower to arrive and far more easily modulated, and that modulability is the whole reason umami is designable at all.

The read, station by station

Once the clamshell shuts and the inner cradle opens, the message begins a relay that is worth counting out, because every station in it is a place where a little signal becomes a larger one, and a place where our engineering can push.

1. Glutamate binds; the flytrap closes; the transmembrane bundle changes shape.
2. The receptor activates **gustducin**, the taste cell's own G-protein, which splits into an α-subunit and a βγ pair.
3. The **βγ** pair switches on the enzyme **phospholipase C-β2 (PLCβ2)**.
4. PLCβ2 cleaves the membrane lipid **PIP_2** into diacylglycerol and the diffusible messenger **IP_3**.
5. IP_3 crosses the cytoplasm and opens the IP_3 receptor on the cell's internal calcium stores.
6. **Calcium** pours out of those stores into the cytosol.
7. That calcium opens **TRPM5**, a calcium-gated, sodium-passing channel that is the shared final gate of sweet, umami, and bitter transduction, and the cell depolarizes.
8. The depolarization vents **ATP** through the **CALHM1/CALHM3** channel to the afferent nerve waiting below, where it lands on purinergic **P2X2/P2X3** receptors — and only now is the savor delivered.

Eight stations, and each catalytic one multiplies. One activated receptor switches on many gustducins before it resets; one PLCβ2 cleaves many PIP_2 before it lets go; the calcium step recruits its neighbors in a small regenerative burst. Treat each stage as a rough multiplier and the gain of the whole cascade is their product. If one receptor turns over on the order of twenty G-proteins, each PLCβ2 makes on the order of fifty messengers, and the calcium release adds a further tenfold, then

> $\text{gain} \approx 20 \times 50 \times 10 \approx 10^4$.

A few molecules of free glutamate — far too few to matter if each had to force a channel open by hand — raise a signal roughly ten thousand times their number. But the gift comes with a bill: eight stations of binding, catalysis, and diffusion do not come free. If each contributes between fifteen and forty milliseconds, the message takes on the order of 120 to 320 milliseconds to arrive — a tenth to a third of a second, held honestly as a range. It is exactly the beat-late swell you feel when the savor of a broth blooms *after* the spoon has left your mouth and stays. Heat is a door kicked open; umami is a message passed hand to hand down a long hall, slower to arrive and slow, gratefully, to leave.

A second opinion: the taste mGluRs

The TAS1R1·TAS1R3 dimer is the principal reader, but it is not the only one. Taste tissue also expresses truncated, taste-specific variants of the brain's **metabotropic glutamate receptors** — **taste-mGluR1** and **taste-mGluR4** — each a class C GPCR with a shortened Venus flytrap of its own. These read glutamate too, but at lower affinity and over a broader span of concentration, a second and third opinion on the same molecule. In the brain their full-length cousins tune synapses; on the tongue, shorn of part of their N-terminal domain, they widen the window over which savor can be registered, catching the signal where it is faint and refusing to saturate where it is strong.

Several readers, one word. The mouth hedges its bets on a signal it cares about this much, and the reason is metabolic. Free glutamate is the body's cue that it is eating protein — the most precious thing on any plate, the material a body is literally built from. A sense that important is not left to one sensor that might fail or be fooled; it is read in parallel, by a high-affinity dimer and a pair of lower-affinity sentinels, so that the tongue reports *protein is here* across the whole range from a trace to a flood. When we write umami into a plant, we are writing to all of these readers at once, which is part of why the axis is so robust to measure: the signal is heard in triplicate before the brain ever receives it.

Synergy — the allosteric multiplication

Umami's defining trick is not the glutamate read alone. It is **synergy** — one of the most striking positive-allosteric effects in all of sensory biology, and the clearest place in the mouth to watch our Index move.

Add a **5′-ribonucleotide — inosine monophosphate (IMP)**, abundant in meat and fish and the reason dried bonito, anchovy, and cured ham deepen a sauce, or **guanosine monophosphate (GMP)**, from dried mushrooms — and the umami signal does not merely add. It *multiplies*. The ribonucleotide binds its own pocket on the outer face of the TAS1R1 flytrap, just beside the hinge, and acts as a molecular wedge that stabilizes the closed clamshell. With the hinge braced shut from outside, glutamate lets go far more reluctantly, and the receptor reaches its half-point at a much lower glutamate concentration than before. In the language of the Index this is simple: the ribonucleotide drives the glutamate **EC_{50} down**.

Hold the Index formula in mind and the consequence is exact. Because a score is built on the inverse of EC_{50}, a fall in EC_{50} by some factor is a rise in IFI-U by that *same* factor. Write it out. Let a glutamate-only savory read have

```
IFI-U(glutamate alone) = 1,000 × G₅₀ ÷ E_g,
```

where G_{50} is the reference half-point and E_g the glutamate's effective EC_{50}. Now let the ribonucleotide lower that EC_{50} by a potentiation factor φ, from E_g to E_g ÷ φ. The new score is

```
IFI-U = 1,000 × G₅₀ ÷ (E_g ÷ φ) = φ · (1,000 × G₅₀ ÷ E_g) = φ
· IFI-U(glutamate alone).
```

The ribonucleotide does not add to the score; it multiplies it. A glutamate-only savor reading a hollow IFI-U of 300 becomes, under a potentiation of φ ≈ 8.7, a read of 2,600 — the *same* glutamate, heard once quietly and once loudly, depending only on what sits in the hinge beside it. This is why savory depth is a composition problem and not a dosage problem: you do not always add more glutamate; you add the molecule that makes the glutamate already present speak louder.

You can watch the same multiplication on the dose–response curve, which for a cooperative receptor follows a Hill function,

```
R = R_max · [c]ⁿ ÷ (EC₅₀ⁿ + [c]ⁿ).
```

Set glutamate sitting exactly at its half-point, [c] = EC_{50}, and by definition the response is R_max ÷ 2 — fifty percent. Now add a ribonucleotide that lowers the EC_{50} threefold, to EC_{50} ÷ 3, while the glutamate concentration does not change at all. The ratio [c] ÷ EC_{50} is now 3, and with a Hill coefficient near n = 1 the response climbs to 3 ÷ (1 + 3) = 0.75 — three-quarters of maximum. The receptor jumps from half-lit to three-quarters-lit with not one extra molecule of glutamate, purely because the curve slid left beneath a fixed dose. That leftward slide is the whole of what a pinch of dried mushroom does to a tomato sauce.

The discovery of this synergy is itself a piece of culinary history. The chemist Kikunae Ikeda isolated glutamate from kombu kelp in 1908 and gave the taste its name; decades later the ribonucleotides of

dried bonito and shiitake were shown to potentiate it. It is why the Japanese stock *dashi*, built from exactly kelp (glutamate) and bonito (inosinate), is one of the most efficient umami machines ever devised — two ingredients chosen, centuries before anyone could read an EC_{50}, to load the receptor from both sides at once.

> Glutamate closes the savory clamshell. A ribonucleotide leans on the hinge. *Together they say depth the mouth cannot un-hear.*

The Index, in the key of umami

A flavor you cannot measure is a flavor you cannot design, so before we build anything we anchor a number. The Iyengar Flavor Index scores every axis against the half-maximal activation of its receptor:

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

Read it as a sentence about loudness. A louder stimulus is one the receptor answers at a lower concentration — a smaller EC_{50} — so putting EC_{50} in the denominator makes loudness climb as the number falls. The reference is chosen so that the native stimulus lands at exactly 1,000: set candidate equal to reference and the ratio is one, the score is 1,000, the anchor holds. For umami we fix that reference at the native receptor's half-activation by L-glutamate, writing it G_{50}. By convention we set **G_{50} = 5 mM** — a representative value within the low-millimolar range that heterologous assay data suggest for this receptor — fixed once by definition, like the dilution-to-threshold anchor of the Scoville scale, so no reader mistakes the convention for a single measured value.

Five millimolar is a tangible thing, and converting it is the first worked example. L-glutamic acid has a molar mass of **147.13 g/mol**, so

> 5 mM × 147.13 mg/mmol = 735.65 mg/L = **73.6 mg per 100 mL**.

As the crystalline seasoning most kitchens know it — monosodium glutamate monohydrate, molar mass **187.13 g/mol** — the same 5 mM is 5 × 187.13 = 935.65 mg/L, about **93.6 mg per 100 mL**: a light pinch in a small glass of water. That is the concentration at which the reference is set, and the rung from which every umami score is measured up or down.

Now run real scores back through the formula. Rearranged, the score tells you exactly how concentrated a food's savory extract must be, relative to reference, to reach the receptor's half-point: $EC_{50}(candidate) \div G_{50} = 1{,}000 \div IFI\text{-}U$. A shipping tomato, picked green and reddened in a warehouse, reads a hollow IFI-U of about 300; a vine-ripened heirloom allowed to finish reads about 2,600; and the full-spectrum pepper we pushed furthest carries a predicted IFI-U of 1,800 on its savory axis. Each resolves into a single ratio you can check with a calculator.

STATE	IFI-U	EC_{50} ÷ G_{50} = 1000 ÷ IFI	EC_{50} (MM, AT G_{50} = 5)	MG GLUTAMATE / 100 ML
Reference (native glutamate)	1,000	1000 ÷ 1000 = 1.000	5.00	73.6
Shipping tomato	300	1000 ÷ 300 = 3.333	16.67	245.3
Phoenix · SS-0096 (pepper)	1,800	1000 ÷ 1800 = 0.556	2.78	40.9
Restored heirloom tomato	2,600	1000 ÷ 2600 = 0.385	1.92	28.3
Scorecard ceiling	5,000	1000 ÷ 5000 = 0.200	1.00	14.7

Nothing in that table is invented; every ratio is a canonical score run once through the formula, and the final column is only that ratio multiplied by G_{50} and by glutamate's molar mass. Read down it and the relationship is plain: as IFI-U climbs, the EC_{50} falls — the richer-reading food reaches the receptor's half-point at an ever-lower concentration. The restored heirloom reaches half-max at 1.92 mM-equivalent where the shipping tomato needs 16.67, a ratio of

> `1.92 ÷ 16.67 = 0.115, or equivalently 2,600 ÷ 300 =` **`8.67×`**.

The restored fruit's savory extract is 8.67 times louder at the receptor — you can dilute it almost ninefold further and still half-light the cell. And it reaches that loudness by two levers pulling the same way: more free glutamate per gram, which lowers the extract's effective EC_{50} directly, and the fruit's own co-produced ribonucleotides, which lower it again by synergy. The two multiply.

Lay the pure synergy lever out on its own and the ladder reads cleanly — a glutamate-only base of IFI-U 300, climbed by nothing but the wedge in the hinge:

RIBONUCLEOTIDE LOAD	POTENTIATION φ	EFFECTIVE IFI-U = 300 · φ	READS LIKE
none	1.0	300	an interrupted, hollow tomato
light	≈ 2	600	a fresh tomato warmed through
moderate	≈ 6	1,800	the savor a superhot carries on its U axis
full	≈ 8.7	2,600	an heirloom allowed to finish on the vine
dashi-class	> 10	> 3,000	kelp married to dried bonito

The potentiation factors are illustrative — the receptor literature puts φ anywhere from severalfold to past an order of magnitude — but the endpoints they land on are the canonical scores, and the arithmetic in between is only multiplication. A number that behaves this predictably is a number you can design against.

A score is only a measurement if it comes back the same tomorrow, so the last piece of the arithmetic is error. Our engineered lines are designed to hold a batch-to-batch **coefficient of variation of 7% or less** — the standard deviation as a fraction of the mean, σ ÷ μ. Because IFI-U is proportional to the inverse of a measured EC_{50}, that relative error passes almost one-for-one into the score: an IFI-U of

2,600 carries a one-sigma band of 2,600 × 0.07 ≈ **±182**. And because the score is a *ratio* of two measured EC_{50}s — the candidate's and the reference's — their errors add in quadrature:

```
CV(IFI) ≈ √(0.07² + 0.07²) = √0.0098 ≈ 0.099, about 10%.
```

On a restored-tomato IFI-U of 2,600 that is a band of roughly ±257 — tight enough that a tomato reading 2,600 and a shipping fruit reading 300 are genuinely different plants, not two draws from one noisy urn. A savor you cannot reproduce is not a soft number; it is not a number at all. In our family's work reproducibility came before loudness, because a figure you cannot stand behind is a small dishonesty dressed as precision, and we would rather give the reader a quiet true number than a loud unreliable one.

Only free glutamate is read

One distinction untangles nearly all the confusion that has ever surrounded this taste: **only free glutamate is read.** The glutamate bound up inside intact protein — the overwhelming majority of it in any steak, any bean, any wheel of cheese — is silent to the receptor, folded away in peptide chains the tongue cannot parse. Umami is the taste of glutamate *liberated*, and almost every traditional process prized for deepening savor is, at bottom, a process that frees it. Ripening a tomato on the vine, aging a cheese for a year, fermenting a soybean into miso, drying a mushroom, browning a crust — each breaks protein down into free amino acids, and frees the ribonucleotides alongside them, manufacturing umami out of time and enzymes.

A fresh steak holds enormous glutamate and tastes only faintly of it; the same steak dry-aged tastes profoundly of it, because aging is simply the slow release of a signal that was present all along. This is why a short table of *free* glutamate tells you more about a food's savor than any protein label ever could. The representative figures run from kombu kelp at well over two thousand milligrams per hundred grams, through a hard aged cheese well past a thousand and a dried shiitake near it, down to a vine-ripened tomato at a couple hundred and a warehouse tomato at almost nothing — not because the warehouse fruit lacks protein, but because it was never allowed to free any. Most tellingly of all, free glutamate is the single most abundant free amino acid in human milk, by a wide margin and rising through the weeks of lactation — so that umami, not sweetness, may be among the very first tastes a person ever knows. We did not invent this depth. We learned where the plant stores it, and how to let it finish.

The tomato, interrupted

The tomato is the parable of this whole axis. A vine-ripened heirloom accumulates free glutamate in the locular jelly around its seeds as it reddens, climbing into the range of a couple hundred milligrams per hundred grams — a concentration a well-aged cheese carries several times over, which is exactly why both taste so deeply of themselves. The accumulation runs by the plant's own route:

glutamate decarboxylase (GAD) and **glutamate dehydrogenase (GDH1)**, the enzymes of the fruit's native GABA shunt, which load the ripening fruit with the amino acid as it colors on the plant.

Harvest that fruit green for a cross-country truck and you arrest the very metabolism that makes it taste like anything. By the time it reddens in a warehouse under ethylene, the glutamate it never made is simply absent, and its savory read collapses from the heirloom's 2,600 to a hollow 300 — a sevenfold-and-then-some loss of loudness, the 8.67× of the earlier arithmetic run in reverse. The modern commodity tomato is not a different fruit. It is a tomato *interrupted*.

Our tomato work, **2,040 constructs** deep, is in large part the restoration of that native GAD/GDH1 flux — our model predicts lifting the fruit's IFI-U from the interrupted 300 back toward the 2,600 of a tomato allowed to finish, and holding it there at a batch CV of 7% or less, where a field heirloom's savor swings wildly with season, soil, and the hour it was picked. This is not an invention. It is a *return*: the switching-back-on of a pathway the plant already owns, so that the fruit does on a loading dock what it would have done on the vine.

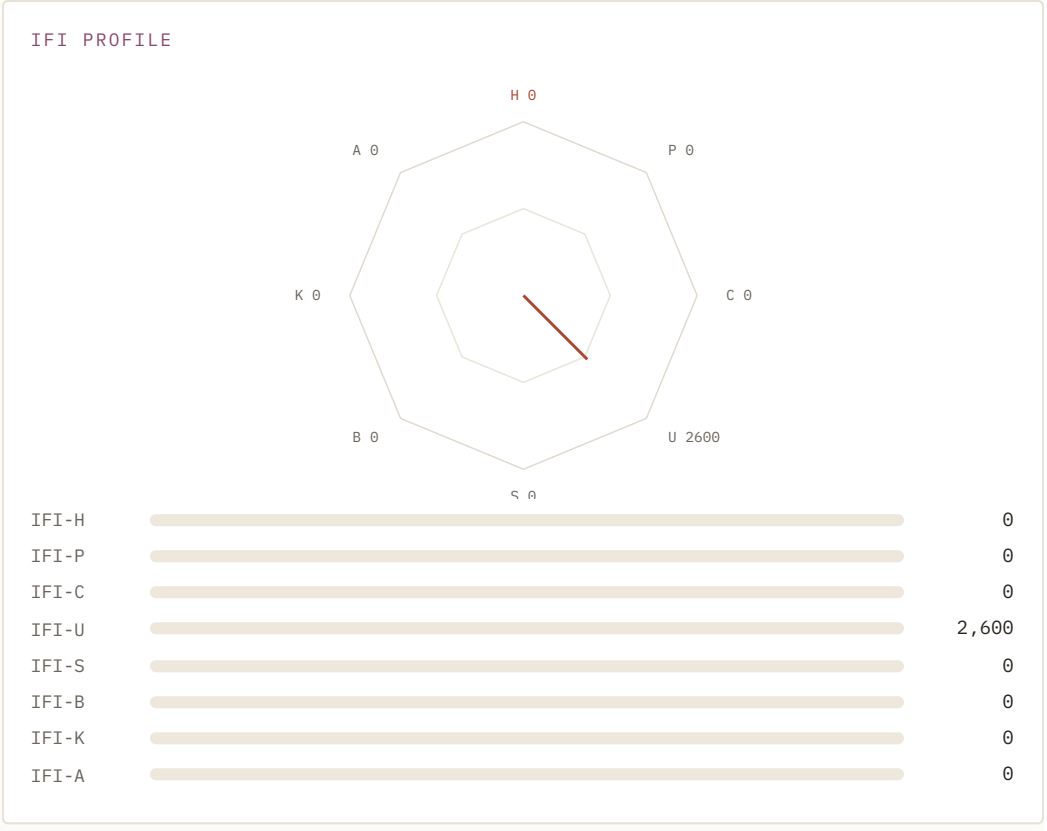

A tomato restored — IFI-U 2,600, up from the interrupted 300, by the plant's own enzymes, in our model's prediction. We say nothing here about its other seven axes, because this is the savory chapter and savor is the number that was stolen; the restoration of the tomato's sweetness is a second lever, pulled with the same hand, that this axis does not need to claim.

Building savor into the bean and the fungus

Where the tomato stores glutamate alone, a handful of other organisms make glutamate *and* its ribonucleotide partner together, so that the synergy is native to the thing itself — no second ingredient required. These are the hosts we reach for when we want depth that stands on its own, and our family comes to them with a vegetarian's particular gladness, because the brothy savor people reach to meat for turns out to be, at the receptor, a glutamate-and-ribonucleotide read a bean and a mushroom can raise as loudly as any stock.



Soybean, before the ferment — the protein reservoir a koji mold unlocks into free glutamate and ribonucleotide at once, the synergy already built into the living organism.

ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Tomato	2,040	S · U · F · M	Easy
Soybean	360	U · K	Frontier
Tea	280	U · B · A	Hard
Yeast	240	U · K	Hard
Koji mold	80	U	Hard
Shiitake	80	U	Hard
Mushroom	40	U	Hard

Read the roster and you can see a whole vegetarian pantry of depth. **Soybean**, 360 constructs deep and carrying both umami and kokumi, is the raw reservoir a **koji mold** — 80 constructs of its own — unlocks in miso and soy, the mold's proteases freeing glutamate while its nucleotidases free the ribonucleotides beside it, synergy assembled inside one ferment. **Shiitake** and **mushroom**, 80 and 40 constructs, are the plant kingdom's guanylate factories, the GMP source that makes a mushroom stock deepen a tomato one past the sum of its parts. **Yeast**, 240 constructs across umami and kokumi, is the savory backbone of a hundred spreads and broths. And **tea**, 280 constructs, carries its umami as L-theanine beside its bitterness and astringency — proof that savor is not only the province of the broth pot.

When we write depth into one of these, we are not imitating meat. Meat and plant are both, in the end, only carriers of the same two molecules the same receptor reads, and a soybean, a mushroom, and a fermented paste can present them as fluently as any animal —

more cleanly, in fact, and grown to hold more of what nourishes. We are simply speaking the receptor's own first language in the organism that can hold it most cleanly, which is the whole of what a plant-first kitchen has always quietly done.

Umami inside a pepper — the full-spectrum proof

The most persuasive demonstration that umami is portable — that it is a receptor read and not a property of broth — is a plant nobody would ever go looking for savor in. **Phoenix** (SS-0096) is the superhot we designed to register on all eight axes at once, with a predicted heat of approximately 2.94 million SHU, and set into that fire, deliberately, is a predicted savory read of **IFI-U 1,800**.



Phoenix (SS-0096) in the controlled chamber — the first cultivar designed to register on all eight axes, carrying a predicted IFI-U 1,800 on the savory channel in a fruit most people meet only as fire.

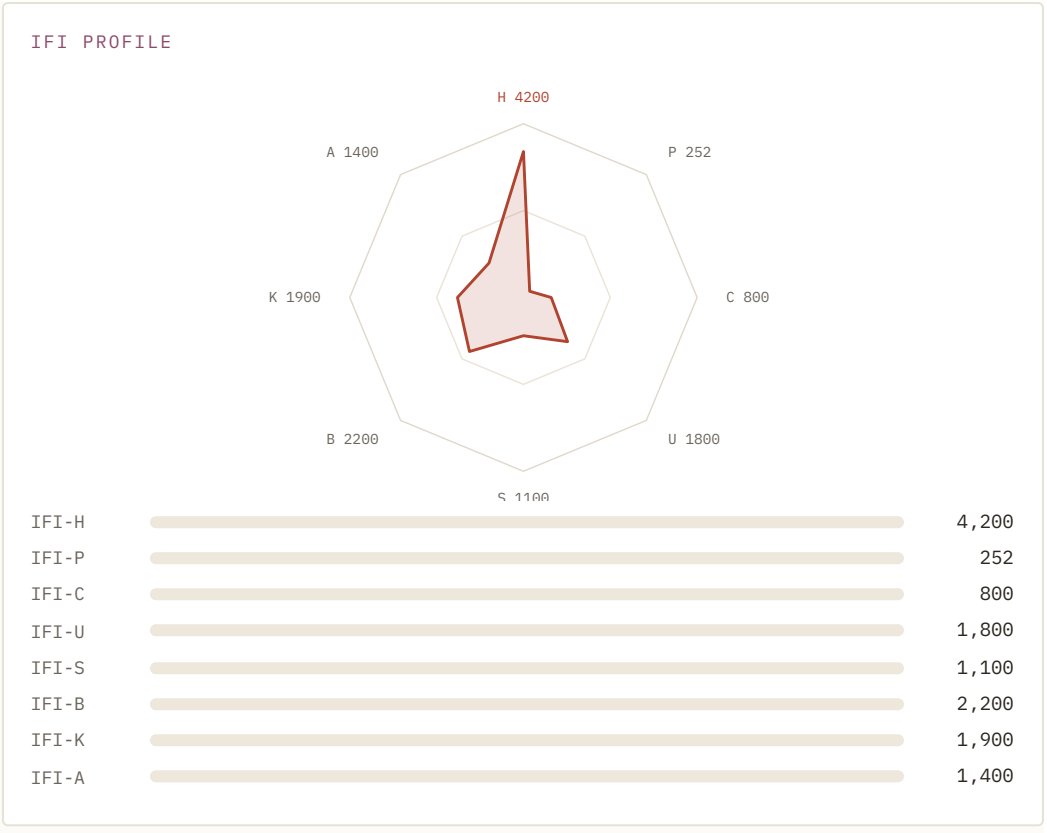

Read the savory line against the ladder already built. At IFI-U 1,800, Phoenix is designed so the receptor's half-point falls at 1000 ÷ 1800 = 0.556 of the reference concentration — it is 1.8 times louder on the savory axis than native glutamate, and fully 1,800 ÷ 300 = **six times** louder than a shipping tomato, in a fruit whose first gift is a wall of capsaicin. A pepper has no broth-making tradition, no dashi, no age; its umami is there because we wrote the glutamate-and-ribonucleotide chemistry into it directly, the same way we wrote the heat. That a single fruit is designed to carry 4,200 of heat and 1,800

of savor at once — each anchored to a different receptor, each designed to hold within a few percent batch-to-batch — is the clearest proof we can offer that these eight tastes are eight separable, writable channels, and that savor, far from being locked to meat and time, is simply one of them.

What the number cannot say — the honest limits

A single number is a form of respect for the reader, but only if it is honest about what it leaves out, and umami leaves out more than most. We hold these limits in plain sight.

The first is **sodium**. Glutamate's taste is strongly sodium-dependent — it is *monosodium* glutamate that the world seasons with, not the bare acid, and the perceived savor of a given free-glutamate concentration rises and falls with the salt beside it. The Index anchors IFI-U at a fixed reference ionic condition, but a dish served under-salted will read, to a mouth, quieter than its score; a dish served salt-forward, louder. The number measures the receptor's potency, not the whole table's chemistry.

The second is the **gap between the tube and the tongue**. An EC_{50} measured on a receptor in a dish is a clean, reproducible quantity, which is why we build on it — but perceived intensity is shaped by the food's matrix, by competing and masking tastes, by temperature, and by adaptation across a meal. A score faithfully measures how loudly a molecule speaks to its receptor, and only *estimates* how loud the eater will find it.

The third lives at the **shared final gate**. Sweet, umami, and bitter all vent their signal through the same TRPM5 channel, so stacking those metabotropic axes is not simply additive: load two of them hard at once and they compete for one channel population, and a cultivar designed to be both intensely savory and intensely sweet may deliver a little less of each than its two scores, read alone, would promise. We design the composition, not just the peaks.

The fourth is **measurement itself**. Free glutamate is the quantity that matters and the hard one to pin — it shifts with every hour of ripening, every degree of storage, every day of a ferment — so an organism's savory read is a read of a moving target, which is exactly why the batch CV and the error propagation are not housekeeping but the heart of the claim. And the reference G_{50} is a stated convention, fixed by definition and not derived from any single measurement; change it and every umami score shifts together, which is why it is locked and never quietly moved. None of this weakens the number; it keeps it honest about the world it is measured in.

The receptor's first language

Umami is a homecoming, and the arithmetic only deepens the reason. Of the eight axes, this is the one most purely about nourishment — the body's own signal that protein, the thing it is built from, has arrived — read in triplicate by a high-affinity dimer and two patient sentinels, amplified ten-thousandfold through its long second-messenger relay, and multiplied again by a ribonucleotide leaning on a hinge. It is the taste a newborn likely meets before any

other, written into mother's milk ahead of sweetness itself. To raise it in a tomato by switching back on the plant's own GAD and GDH1, or in a bean and a mushroom where glutamate and guanylate are already made side by side, is not to counterfeit depth but to let an organism finish a sentence it was always trying to say.

That is the quiet ethic under the whole of this axis, and under our family's work on it: food grown to hold the most it can of what feeds a body, and designed carefully enough that *the most* stops being a wish and becomes a number we can stand behind. Glutamate closes the savory clamshell; a ribonucleotide braces the hinge; the cascade swells its beat-late swell; and a plant, given back its own enzymes, speaks the receptor's first language more fluently than it has in a hundred years of being hurried to market. We learned where the plant stores its savor. The rest of the work is only letting it finish. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 7 · TAS2R38 — *the bitter sentinel.*

Every other receptor in this book, we try to make speak more loudly. This one, we try to make speak less. Bitterness is the mouth's alarm — the broadly tuned warning that evolved to catch the alkaloids and glycosides that signal poison — and it is magnificent at its job, which is exactly the problem, because a great many things the body would benefit from are bitter, and a great many foods carry a bitterness that masks everything else they have to offer. The design question for bitterness is inverted: not how to activate the sentinel, but how to quiet it without blinding it. It is the only axis in the whole Index we read as a subtraction, and the only one where the number climbs as the signal falls — and holding that inversion honestly, in the molecule and in the mathematics, is the work of this chapter.



Coffee at the bean — chlorogenic-acid lactones and the roast's own breakdown products are what the bitter receptors read; attenuate them and the cup's sweetness and body step forward.

Humans carry about **twenty-five bitter receptors**, the TAS2R family, class A GPCRs each tuned — some narrowly, some very broadly — to different bitter chemotypes, so that between them they cover an enormous chemical space of potential toxins. Some, like TAS2R16, read a specific class (β-glucopyranosides); others, like the broadly tuned TAS2R10 and TAS2R46, answer to dozens of unrelated bitter compounds, a deliberate generalism that trades precision for coverage — when the stakes are *poison*, a smoke detector that goes off for too many things is safer than one that misses. That is a different design philosophy from every sweet or umami receptor,

which the body built to be exquisitely specific; bitterness was built to be suspicious. The best-studied member is **TAS2R38**, the receptor for **phenylthiocarbamide (PTC)** and the related compound **PROP** (6-n-propylthiouracil), both of which carry the N–C=S thiourea group the receptor is tuned to find — and it is famous for a genetic coin-flip written into its sequence, which we will come to once we have watched the receptor do its actual job.

How a mouth reads "bitter," step by step

A TAS2R is seven membrane-spanning helices threaded through the apical membrane of a Type II taste cell — a bundle of coils with a shallow pocket open to the saliva. Reading a bitter molecule is not one event but a relay, and it is worth walking the whole length of it, because every step is a place where the loudness can, in principle, be set, and because the receptor only earns its reverence once you see how much machinery it spends on a single word of warning.

1. **Binding.** A bitter ligand — goitrin from a cut cabbage, a chlorogenic lactone from a dark roast — diffuses into an orthosteric pocket formed where transmembrane helices 3, 5, 6, and 7 lean together in the outer third of the receptor. For a broadly tuned receptor the pocket is forgiving, lined so that many unrelated shapes can find a seat; for a narrow one it is particular, a near-lock-and-key.
2. **The nudge.** Occupancy tilts helix 6 outward by a few ångströms — the conserved class A activation motion, the same hinge rhodopsin uses to report light — opening a cytoplasmic cradle on the inside of the cell.
3. **The G protein.** Into that cradle docks **gustducin**, the taste cell's own heterotrimeric G protein (α-gustducin with β_3 and γ_{13}). The receptor acts as a nucleotide-exchange factor, swapping GDP for GTP on the α subunit, and the trimer splits into α and βγ.
4. **The amplifier.** The freed βγ pair activates **phospholipase C-β2**, which cleaves the membrane lipid PIP_2 into diacylglycerol and **inositol trisphosphate (IP_3)**. A single occupied receptor can drive many enzyme turnovers before it lets go; this is the step where a faint chemical whisper becomes a cellular shout, and it is why a few molecules of a bitter alkaloid can taste like a warning shouted across a room.
5. **The calcium.** IP_3 binds the IP_3 receptor on the endoplasmic reticulum and opens it; stored Ca^{2+} floods out into the cytosol.
6. **The gate.** That calcium opens **TRPM5**, a Ca^{2+}-activated monovalent cation channel. Sodium pours in and the cell depolarizes.
7. **The release.** Depolarization opens **CALHM1/CALHM3**, a large-pore channel, and the cell exhales **ATP** directly into the synaptic cleft — an unusual, channel-based neurotransmission with no vesicle.
8. **The nerve.** ATP lands on **P2X2/P2X3** receptors on the gustatory afferent fibre, and the message — *bitter, here, this much* — climbs to the nucleus of the solitary tract in the brainstem and on to the gustatory cortex, where it arrives as the thing you actually feel.

We traced the βγ limb because it carries the depolarizing current, but the α subunit works in parallel: freed α-gustducin switches on a phosphodiesterase that drives cytosolic cyclic AMP *down*, and that falling cAMP releases a brake on the very calcium machinery the other arm is loading — a second, slower hand pushing the same door. Both arms converge on the calcium of steps 5 and 6, which is part of why a bitter read is so quick and so hard to fool: a molecule that slips one limb is usually caught by the other.

Eight steps, and the sensation of "too bitter" is the compounded gain of all of them: a steep amplifier in the middle (steps 3–4) sitting on a binding event at the top (steps 1–2). Both halves matter to a designer. You can quiet the top — present the receptor with less ligand, or a ligand it reads less avidly — and the whole cascade below reports less, because step 4's amplification multiplies a smaller input. What you cannot safely touch is the amplifier itself: the gustducin–PLC-β2–TRPM5 relay is shared across *all* the TAS2Rs and the sweet and umami cells besides, so dampening it would blind far more than one bitter note. Every intervention we make lives at step 1 — at the ligand, before the fire — which is both the cleanest place to work and the only place a plant's own chemistry actually reaches.

The coin-flip in the gene

Bitterness, more than any other axis, is personal — the receptor itself differs from mouth to mouth, which is why "too bitter" is a real disagreement between people and not merely a difference of opinion. TAS2R38 carries three variable amino-acid positions — **49, 262, and 296** — and they travel together as two common haplotypes. **PAV**, the "taster" form, reads PTC and PROP intensely; **AVI**, the "non-taster" form, barely reads them at all. The split was discovered by accident in 1931, when the chemist Arthur Fox let some PTC dust drift into the air of his lab and a colleague tasted bitterness while Fox tasted nothing.

HAPLOTYPE	POSITION 49	POSITION 262	POSITION 296	WHAT THE MOUTH DOES
PAV · taster	Proline	Alanine	Valine	reads PTC / PROP intensely
AVI · non-taster	Alanine	Valine	Isoleucine	barely reads them at all

Those three residues sit in the transmembrane helices that shape the binding pocket and the activation hinge; swapping proline for alanine at position 49 alone stiffens the way helix 2 packs, and the receptor that results reaches for the thiourea group far less willingly. Inheritance is close to a textbook two-allele story — roughly three in four people carry at least one PAV copy and taste the bitterness; the rest are spared it — with the PAV/PAV homozygotes, especially those with dense fungiform papillae, forming the "supertasters" for whom a raw Brassica is genuinely vicious. Whether a particular cabbage tastes savage or mild to you is decided, in part, by which versions of this single gene you were handed. We dwell on it because it sets a hard honest limit on the work: there is no single attenuation that is *correct* for every mouth, only one that is right for the population you are cooking for — and the axis has to say so plainly rather than pretend the receptor is the same organ in all of us.

The mathematics of turning it down

The Iyengar Flavor Index puts every axis on one number by anchoring it to the half-maximal activation of a receptor:

```
IFI(c) = 1000 × EC₅₀(ref) ÷ EC₅₀(c)
```

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. The reference compound scores exactly **1,000** on its own axis by construction — $EC_{50}(ref) \div EC_{50}(ref) = 1$ — and everything else is read as louder or softer against it. For the seven activating axes, *louder is the goal*, and a lower candidate EC_{50} (more potent, half-response reached at less compound) lifts the score. Work one example to fix the idea: a cultivar that registers **4,200 on heat** has, by the formula, $EC_{50}(c) = EC_{50}(ref) \times 1000 \div 4200 =$ **$0.238 \times EC_{50}(ref)$** — its capsaicin reaches half-activation of TRPV1 at under a quarter of the reference concentration, a read better than four times as loud.

Bitter is the one inversion. Here loudness is the enemy, so the axis is scored not on how fiercely the bitter compound speaks but on how much quieter we made it — the same reciprocal structure, read on the suppression side: a half-effect concentration of attenuation stands in for EC_{50}, and a more effective quieting earns a higher **IFI-B**. A designed **B of 2,200** therefore implies a suppression reaching half-effect at 1000 ÷ 2200 = **0.45 ×** the reference's concentration — an attenuation about 2.2 times as effective as the anchor. The number goes *up* as the mouth's experience goes *down*. That is the whole trick of the axis, and it is why "Zero-Bitter" is a score you climb.

How much ligand must actually move to earn it? Here the simple pharmacology repays close reading, because the answer is not linear and the nonlinearity is the entire reason attenuation-at-the-source beats brute force. Take the one-site occupancy model, the honest sketch: the fraction of maximal receptor response is

```
R = [L] ÷ ([L] + EC₅₀)
```

where [L] is the ligand concentration at the receptor. (Real bitter dose–response often climbs more steeply than this — a Hill coefficient above 1 — which only sharpens everything below.) Picture a dark roast sitting at **$[L] = 3 \times EC_{50}$**, three times its receptor's threshold — a strong, saturating read. Its starting response is $R_0 = 3 \div 4 =$ **0.75**, three-quarters of the way to maximal. Now ask for the canonical attenuation band — a **20–65% reduction** in that perceived bitterness — by two routes. **Route A** lowers the ligand in the plant: rearranged, the residual needed is $[L'] = EC_{50} \times R \div (1 - R)$. **Route B** deploys a competitive **bitter-blocker** that raises the receptor's apparent EC_{50} by a factor $(1 + [I]/K_i)$ without ever triggering it.

PERCEIVED-BITTER REDUCTION	RESIDUAL RESPONSE R	ROUTE A — LIGAND CUT NEEDED	ROUTE B — EC_{50} RAISED
20% (band floor)	0.60	50%	2.0×
50% (midpoint)	0.375	80%	5.0×
65% (band ceiling)	0.26	88%	8.4×

Read the first column against the third and the lesson jumps out: to drop the *felt* bitterness by a fifth you must strip *half* the compound, and to drop it by the canonical 65% you must remove **88%** of the ligand — because on the saturating shoulder the receptor is already near-full, and the first molecules you take away barely register until you get down near the threshold. This is exactly why piling sugar on a bitter cup is a losing game: masking fights a saturated receptor with distraction and wins only a little per spoonful, while attenuation walks the ligand concentration back down the steep part of the curve where every reduction finally counts. It is also why a true blocker is so efficient — doubling the apparent EC_{50} (a blocker present at roughly its own K_i) already buys the 20% floor, and the 65% ceiling needs the blocker at about 7.4 × K_i, a steep but finite ask.

To put a bean's chemistry onto the same axis as a receptor, we convert mass to molarity, because EC_{50} lives in micromolar and a label lives in milligrams. Chlorogenic acids run about **354 g/mol**; a strong cup carrying roughly 120 mg of them in 200 mL works out to 0.6 g/L ÷ 354 g/mol ≈ **1.7 mmol/L — about 1,700 µM**, far up the saturating shoulder of a receptor whose threshold sits in the single-to-tens-of-micromolar range. That one conversion is why coffee bitterness feels like a wall rather than a dial: the mouth is operating where the curve is nearly flat.

And the quiet has to be *repeatable*. An unengineered plant is a lottery — the bitter load of one broccoli crown against the next drifts by tens of percent, so "less bitter" in a field is a hope, not a specification. Across our engineered material the batch-to-batch coefficient of variation holds to **≤7%**, against about **42%** in wild-type. If the target attenuation is the 50% midpoint, a 7% CV puts roughly two-thirds of harvests inside 50% ± 3.5 points and nearly all of them inside a band about seven points wide — tight enough that a roaster or a formulator can build around the number instead of tasting every lot and guessing. Reproducibility is not a footnote on the mathematics; on an axis defined by subtraction, it *is* the product.

What the alarm is actually reading

It helps to know what the sentinel reads, because it is rarely the obvious culprit. In coffee most people blame caffeine — but caffeine accounts for only a small share of the cup's bitterness; the dominant bitter compounds are **chlorogenic-acid lactones** and, at darker roasts, **phenylindanes**, both of them *created by the roast itself* from the bean's native acids. Coffee bitterness, in other words, is largely a thing the roaster makes, which is exactly why attenuating it in the bean's chemistry — before the fire — is so much cleaner than fighting it with sugar in the cup. In the Brassicas it is the **glucosinolate** breakdown products, including the sulfurous **goitrin** that gives raw broccoli, kale, and cabbage their bite; in cacao, the

theobromine and **polyphenols**; in hops and beer, the **iso-α-acids**. Each is a legitimate read at a real receptor — the bitterness is not an illusion to be masked but a signal to be turned down at its source.



Broccoli at the cut — the sulfurous edge of a raw Brassica is myrosinase turning glucosinolates into goitrin and its kin the moment the tissue is wounded; this is what the bitter sentinel reads, and what we tune down without silencing.

There is a small enzymatic drama worth naming in the Brassicas, because it changes where you intervene. The glucosinolates themselves are nearly odourless and stored apart from the enzyme that activates them; the instant a leaf is bitten or cut, **myrosinase** meets its substrate and hydrolyses the glucosinolate into the pungent, bitter isothiocyanates and goitrin that the mouth then reads. The bitterness, like coffee's, is partly *made at the moment of eating* — which means a designer has two honest handles: lower the stored glucosinolate pool, or modulate the myrosinase step, each of them a tap on the ligand before step 1 of the cascade ever fires. The axis is not telling the plant to abandon its chemistry; it is asking the plant to carry a little less of the specific compound that walls the flavour off.

These are the hosts where the bitter axis is the point, and the shape of the programme is right there in the construct counts — the breadth of the work on grape and the Brassicas, the careful early reach into coffee:

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Grape	1,610	B / A / M	Hard
Broccoli	380	B / M	Moderate
Tea	280	U / B / A	Hard
Broccoli-Kale	280	B	Hard
Orange	120	B / C	Hard
Potato	80	B	Easy
Hops	80	B	Moderate
Coffee	40	B	Hard
Cabbage	40	B	Moderate
Kale	40	B	Moderate
Radish	40	B	Hard
Rapeseed	40	B	Hard

Mask, or attenuate

There are two ways to turn the alarm down, and the Index keeps them firmly apart. You can **mask** — pile on sweetness or fat so the brain is distracted from a receptor that is still firing at full volume, the brute-force path of every sugared coffee drink — or you can **attenuate** at the source: lower the concentration of the bitter ligand in the plant, or deploy a true **bitter-blocker**, a molecule that occupies or allosterically dampens the TAS2R without triggering it, so the sentinel simply reads less. Nature already makes such blockers — **homoeriodictyol**, a flavonoid from an herb long used in traditional medicine, is a known bitter-suppressing compound, and the principle is well established. Blockers come in two mechanistic flavours, and the difference decides how they behave in a real mouth. A **competitive** blocker sits in the same orthosteric pocket the bitter ligand wants, so it raises the receptor's apparent EC_{50} but can be out-competed by piling on more ligand — the surmountable kind, and the kind the "EC_{50} raised" column of the table above assumes. An **allosteric**, non-competitive blocker binds a separate site and damps the response outright, lowering the ceiling the ligand can ever reach; it is not out-competed, but it risks reaching further across the broadly tuned family than you intended. We favour the competitive, surmountable path and the source-reduction path precisely because both stay local to the compound and the receptor in front of them, and leave the sentinel's coverage of everything else intact. Masking adds calories and hides the signal; attenuation removes a designed fraction of the ligand and leaves everything else the plant carries untouched. Attenuation is cleaner, quieter, and ours: across the bitter-forward hosts we target that **20–65% reduction** in the bitter read, enough to let a coffee's native sweetness and a cacao's fruit and a broccoli's green vegetal brightness come forward — without the sugar, and without silencing the receptor's protective role for anything genuinely toxic. We turn the alarm down; we do not cut its wires.



Cacao at the pod — theobromine and a dense raft of polyphenols are the bitter and astringent reads here; tune the first down a measured amount and the bean's own fruit and florals come up out of the dark.

The platform that writes these edits is the same one that engineers heat the other direction — a protein language model to score candidate substitutions, a genome-scale flux-balance model to confirm which enzyme is the true rate limiter of a pathway, and CRISPR to write the change in cleanly. On the heat axis that method found the double variant S39L+L345G — two Pun1 substitutions no breeder had ever combined; on the bitter axis we point the identical

toolchain the opposite way, at the enzymes that *build* the bitter ligand or the pool it is stored in, lowering the load the receptor meets rather than lifting it. Nowhere is the balance of the Index clearer than on a full-spectrum design, where the bitter axis sits quietly down inside a radar that is loud almost everywhere else:

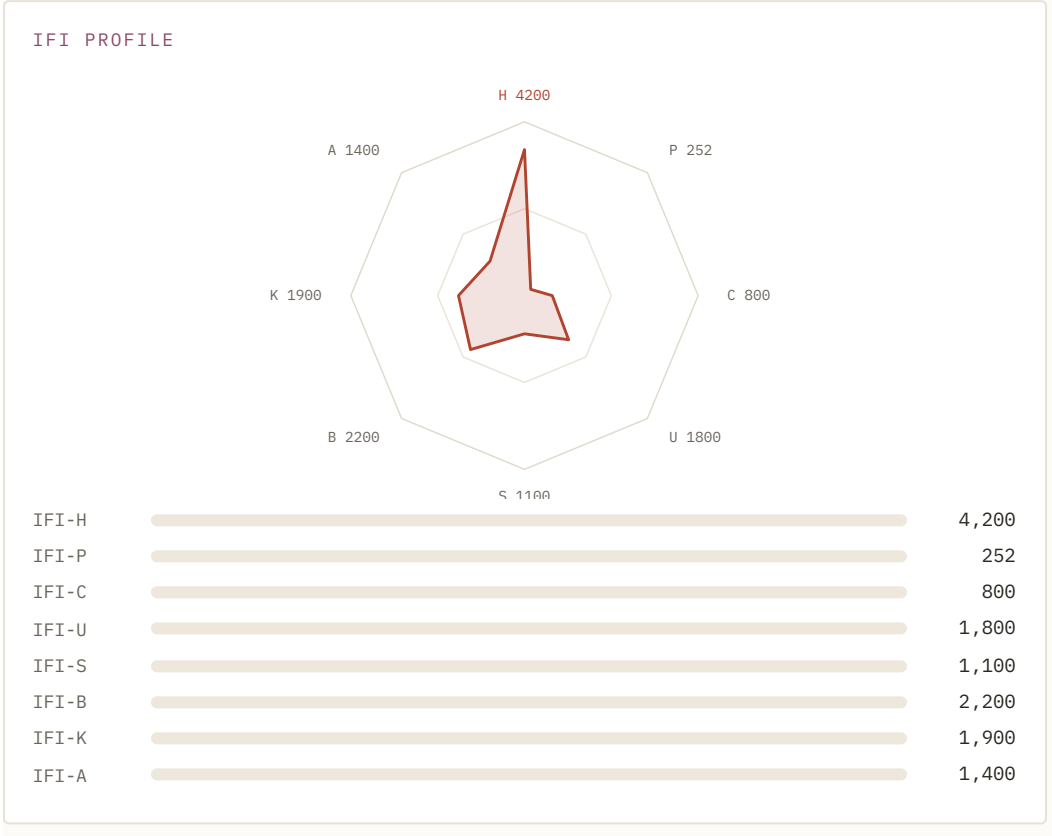

Read that profile as a sentence: four thousand two hundred of heat, two thousand two hundred of engineered *quiet* on the bitter axis, and the other six axes each set to their own intended loudness — one plant speaking all eight at once, with bitterness designed down the way heat is designed up. The B of 2,200 is not a bitterness the plant *has*; it is a bitterness it has been built to **withhold**, scored on the same thousand-fold reciprocal as every other axis, pointed the other way.

> Bitterness is a smoke alarm. You do not want it gone — *you want it to stop screaming at the toast.*

The sentinel beyond the mouth

The bitter sentinel is not confined to the tongue, and a designer is obliged to know where else it speaks before quieting it anywhere. The same TAS2R proteins are expressed in the lining of the **airway**, where bitter compounds relax bronchial smooth muscle and quicken the beat of the cilia that sweep the lungs clean, and in the **gut**, where they sit on the hormone-releasing enteroendocrine cells that help govern the release of appetite-and-fullness signals. The body liked this threat-detector so well that it installed copies throughout — a reminder that a taste receptor is a chemical sensor first and a flavour organ second. We say this carefully and claim nothing from it: these are facts of where the protein lives, not promises about what a quieter plant will do in a body. But they are the reason our edits stay at the ligand in the edible tissue and never reach for the receptor's amplifier — because the receptor you quiet in a leaf is, molecule for molecule, a receptor doing other work elsewhere, and reverence for the plant has to include reverence for the machinery that reads it.

The honest limits

Four limits bound this work, and naming them is part of keeping the lushness honest.

We cap the quiet, on purpose. The band ends at a 65% reduction and never at zero, because a bitter read is a protective read; a plant tuned to *no* bitterness would also be a plant unable to warn you about a genuinely bitter toxin that found its way in. The sentinel is turned down, deliberately short of off.

No single setting fits every mouth. The PAV/AVI coin-flip means a non-taster barely needed the attenuation while a PAV/PAV supertaster might still find the plant sharp; we tune to the median of the population a given host is grown for, and the axis is honest that some of "too bitter" lives in the taster's own receptor, not the food. A number on a page cannot resolve a disagreement that is written into two different genes.

The curve punishes greed. Because the dose–response saturates, chasing the top of the band means stripping most of the ligand — and in coffee and cacao those same chlorogenic acids and polyphenols are also flavour precursors and antioxidants, carrying far more than their bitterness. Pull too hard and you flatten the cup you were trying to free. We prefer the gentler end of the band for exactly this reason: the goal is to let the plant's other notes through, not to hollow it out.

It only works if it is reproducible. An attenuation that lands at 50% in one harvest and 25% in the next is not a design, it is a rumour; the whole value rests on the **≤7% CV** design target that keeps the delivered quiet inside a narrow, build-around-able band. We build reproducibility into the design specification rather than assume it, because on a subtraction axis a design without a predicted batch CV is indistinguishable from nothing happening.

This is the axis where the nourishment thread runs closest to the surface, and where we are most careful with our words. Many of the compounds a bitter read guards — the glucosinolates of Brassica, the polyphenols of cacao and coffee — are precisely the ones people are told are worth eating, and bitterness is a real reason they go uneaten. To attenuate that bitterness is, quietly, to make a bitter-but-worthwhile plant more likely to actually be eaten, in its whole form, with everything it carries intact. We will not phrase that as a health claim, because it is not one; it is an intention, and a reason our family finds the bitter axis among the most worth the work. A child who finishes the broccoli because it no longer fights back has been served by a receptor turned gently down — nothing more is claimed, and nothing less is meant.

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

Chapter 8 · CaSR — *the invisible richness.*

There is a quality in the best cooking that has no flavor of its own and yet changes everything — a fullness, a continuity, a sense that the dish is *complete*, that it coats and lingers and holds together rather than arriving in separate notes. The Japanese word is *kokumi*, "rich taste," and for a long time it was treated as a chef's affectation, a mysticism of the hand. It is not. It is a receptor — a specific protein, reading specific molecules — and the receptor turns out to be one the body already keeps for an entirely different purpose: **CaSR, the calcium-sensing receptor.** When we first set out to give a plant this quality on purpose, we understood we were not chasing a taste at all. We were learning to write, gene by gene, to a protein that most of biology uses to manage a mineral.



Garlic at the clove in cross-section — its γ-glutamyl peptides are among the most potent natural kokumi compounds, the reason a dish built on alliums feels whole.

The receptor that reads the mineral and the meal

CaSR is a class C G-protein-coupled receptor — the same architectural family as the sweet and umami dimers. Picture a very large protein that spends almost all of itself outside the cell: a bilobed clamshell, the **Venus flytrap domain**, roughly a thousand amino acids of sensing surface, tethered through a stiff cysteine-rich neck to the seven-pass transmembrane bundle that every GPCR shares. Two of these chains clasp together into a disulfide-linked homodimer, so the working receptor is a pair of flytraps held mouth to mouth, and the two halves talk to each other as they open and

close. This is not a protein built to report a single molecule arriving once. It is built to integrate — to read the slow, crowded chemistry of a fluid and return one continuous verdict about how rich that fluid is.

In the body at large, that verdict is about calcium. CaSR sits on the parathyroid gland and the kidney, where it monitors the concentration of free calcium ions in the blood and tunes the release of parathyroid hormone to hold that concentration almost perfectly steady. It is one of the quietest and most important governors we carry. The drugs that act on it — the **calcimimetics**, cinacalcet chief among them — were developed for disorders of calcium balance and bone, never for dinner. That a flavor sense and a mineral-regulating endocrine sense should share one receptor is one of biology's small, startling economies, and it is the first thing to hold in mind about kokumi: the molecule that makes a broth feel like being cared for is read by the same lock that keeps your bones supplied.

On the tongue, the same receptor is expressed in the taste epithelium, and the cascade it runs when something binds is worth walking through slowly, because the shape of that cascade is the shape of the sensation. A ligand settles into the flytrap and the clamshell closes; the closure twists the cysteine-rich neck; the twist rearranges the transmembrane bundle; the rearranged bundle couples to **Gq/11** on the inner face of the membrane. Gq switches on **phospholipase C-β**, which cleaves the membrane lipid PIP_2 into two messengers — **IP_3** and diacylglycerol. IP_3 opens calcium stores inside the cell, cytosolic calcium rises, the cell depolarizes, and transmitter goes out to the gustatory nerve. Eight steps from mouth to message, and notice that the very last thing the cell does, before it speaks, is raise its *own* internal calcium — the receptor that reads calcium outside answers by releasing calcium inside. The biology rhymes with itself.

The crucial subtlety is *where* the kokumi molecules bind. Calcium itself is an orthosteric agonist — it sits in the floor of the flytrap and drives the receptor directly. The γ-glutamyl peptides that give food its richness do not, by and large, do this. They bind at the lip of the flytrap as **positive allosteric modulators**: they do not open the trap themselves so much as make it easier to close, lowering the amount of calcium the cell needs to reach half of its full response. This is the same trick cinacalcet plays from a pocket in the transmembrane bundle, from the other side of the membrane. A kokumi peptide is, mechanistically, a calcimimetic you can eat — a molecule that tells the receptor the broth is richer in calcium than it truly is, and the receptor, believing it, leans into every other taste arriving at the same moment.

There is one more lever on the flytrap floor, and it explains a kinship cooks have always known in their hands. CaSR reads free **L-amino acids** as co-agonists: L-tryptophan and its kin, docking in the orthosteric cleft, nudge the trap toward closing and raise its sensitivity to everything else. So a broth already loaded with the free glutamate of umami arrives at the receptor half-persuaded, and the γ-glutamyl peptides need only tip it the rest of the way. Umami and kokumi are not neighbours by accident. They share a flytrap —

which is why savory depth and richness so often rise together in the same slow stock, and why we so often build them into the same plant.

What a γ-glutamyl peptide actually is

The molecules that do this share one odd piece of chemistry, and the whole axis turns on it. An ordinary peptide bond forms between the α-carboxyl of one amino acid and the amino group of the next; it is the bond every protease in your gut is built to cut. Glutamate, though, carries a second carboxyl out on the end of its side chain — the **γ-carboxyl** — and when the peptide bond forms *there* instead, you get a **γ-glutamyl** linkage: the same atoms, bonded through the side door. That single structural choice does two things at once. It hides the bond from the standard peptidases, so these molecules survive cooking, fermentation, and the first moments of chewing intact; and it presents the receptor with a shape the orthosteric amino-acid machinery of the flytrap does not fully recognize, which is exactly why they read as modulators rather than tastes.

The archetype is **glutathione**, γ-Glu-Cys-Gly — the small tripeptide that every one of your cells manufactures by the gram as its chief antioxidant, here moonlighting as a flavor molecule. Trim the glycine and you have **γ-glutamylcysteine**; swap the cysteine for valine and extend it and you reach **γ-Glu-Val-Gly**, a peptide so potent it is used as a reference standard for the sensation, active where glutathione is merely present. Calcium, the receptor's namesake, belongs on the same list as the orthosteric baseline. Arranged by the concentration each needs to half-activate the receptor, they form a ladder of potency that climbs from the mineral through the humble cellular tripeptide to the engineered apex.

Where do the peptides come from in a plant? Not from one enzyme we can rewrite with a single stroke, which is the honest structural fact about this axis. Glutathione is built in two ATP-driven steps — **glutamate-cysteine ligase** (GCL) joins the γ-bond, **glutathione synthetase** (GSS) adds the glycine — and then a third enzyme, **γ-glutamyltransferase** (GGT), works as the kitchen's transpeptidase, lifting the γ-glutamyl group off glutathione and setting it onto other amino acids to mint the wider family of peptides. The alliums take this further: garlic and onion store their pungency precursors as γ-glutamyl-S-alk(en)yl-cysteines, a reservoir of γ-glutamyl peptides held in the clove, which is why aging and slow cooking — black garlic, a long-sweated onion — deepen richness as GGT works the pool and the sharpness cooks away. Time is a kokumi enzyme. So is fermentation. The receptor reads a *pool* of these peptides, not a single product, and that is both the beauty and the difficulty of designing for it.



Onion in cross-section — concentric scales of stored γ-glutamyl peptides; the savory foundation that outlasts its own sharpness.

The mathematics of a taste you cannot point to

The Iyengar Flavor Index scores every axis the same way, by how *loudly* a stimulus speaks to its receptor rather than how much of it a food holds:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the half-activation point read straight off a dose-response curve. The reference compound is anchored, by construction, at 1,000: put the reference into the formula against itself and the ratio is one, times a thousand. Everything more potent than the reference needs *less* of itself to reach half-activation, so its EC_{50} is smaller, so its score climbs above 1,000; everything weaker falls below. The number is an inverse: low EC_{50}, high Index.

For kokumi we anchor the axis on **glutathione**, the one γ-glutamyl peptide present in every living cell and therefore the fairest zero point a tongue could have. In our standard assay — human CaSR expressed in cultured cells, read by calcium-sensitive fluorescence — glutathione half-activates the receptor at **2.00 mM**. That is EC_{50}(reference), and it fixes IFI-K(glutathione) = 1,000 by convention. Every other kokumi ligand is then read against it.

LIGAND	ROLE AT CASR	EC_{50}	IFI-K
Ca^{2+}	orthosteric agonist	4.00 mM	500
Glutathione (γ-Glu-Cys-Gly)	reference modulator	2.00 mM	1,000
γ-Glu-Cys	positive allosteric modulator	1.33 mM	~1,500
γ-Glu-Val	positive allosteric modulator	0.50 mM	4,000
γ-Glu-Val-Gly	positive allosteric modulator	0.16 mM	12,500*

Work two of these by hand, because the arithmetic is the whole idea. γ-Glu-Val half-activates the receptor at 0.50 mM — four times less peptide than glutathione needs — so IFI-K = 1,000 × 2.00 ÷ 0.50 = **4,000**. And γ-Glu-Val-Gly, half-active at 0.16 mM, gives IFI-K = 1,000 × 2.00 ÷ 0.16 = **12,500**. That asterisk matters: the scorecards normalize every axis to a ceiling of 5,000, and γ-Glu-Val-Gly is the rare natural ligand whose raw score runs clean off the top of the chart. On a radar it saturates; in the formula it keeps climbing. We

report the true 12,500 in the text and let the scorecard show the capped bar, because pretending the ceiling is the truth would be the one dishonesty the Index cannot afford.

The dose-response curve underneath each EC_{50} is not a straight line but a sigmoid, and its steepness tells you something you can taste. We fit it with the Hill relation,

```
response = R_max × [L]ⁿ ÷ (EC₅₀ⁿ + [L]ⁿ),
```

and for CaSR the Hill coefficient *n* runs high — near 2.5 — because the receptor is a cooperative dimer, two flytraps persuading each other to close. A high Hill number makes the curve steep, almost switch-like: below threshold the broth is thin, and then across a narrow band of concentration it *snaps* into coherence. That steepness is the molecular reason kokumi feels less like a dial and more like a dish suddenly deciding to be whole.

Walk the curve at a Hill coefficient of 2.5 and the snap becomes quantitative. Let *x* be the ligand concentration in multiples of the EC_{50}; the response fraction is x^2.5 ÷ (1 + x^2.5). At a quarter of the EC_{50} the receptor is barely 3% engaged; at half, about 15%; at the EC_{50}, 50% by construction; at twice, roughly 85%; at four times, 97%. A single fourfold step in concentration — from half the EC_{50} to twice it — carries the receptor from 15% to 85% of its full response. That narrow band is the whole working range of kokumi: the difference between a stock that is almost rich and one that has closed, measured in a doubling and a doubling again.

Because chefs and formulators think in weight, not molarity, the Index carries its conversions. A millimolar of a roughly 300-gram-per-mole peptide is about 300 milligrams per liter — near enough 300 parts per million in dilute water. Multiply each EC_{50} by its molecular weight and the ladder reappears in the units a label uses:

LIGAND	MW (G/MOL)	EC_{50} (MM)	EC_{50} (PPM)
γ-Glu-Val-Gly	303.3	0.16	~48.5
γ-Glu-Val	246.3	0.50	~123
γ-Glu-Cys	250.3	1.33	~333
Glutathione	307.3	2.00	~615
Ca^{2+}	40.1	4.00	~160

So γ-Glu-Val-Gly reaches half of the receptor's response at under fifty parts per million, while glutathione needs better than six hundred — a twelve-fold span of potency compressed into a few atoms of difference. Its bare *detection* threshold, the faintest hint a trained panel can catch, sits lower still, a few parts per million; the EC_{50} the Index actually scores is the half-activation point, and the gap between the two is itself a lesson in why "active at a trace" and "scores high on the axis" are related but not the same measurement.

To make the scale concrete at the stove: a 300-millilitre bowl of broth brought exactly to the γ-Glu-Val-Gly half-activation point carries 0.16 mmol/L × 0.300 L × 303 mg/mmol ≈ **14.6 milligrams** of the peptide — a dose you could balance on the tip of a knife, dissolved past seeing into the whole bowl, and enough to move the

receptor halfway to its ceiling. The potency of this axis is not a figure of speech. It is a few milligrams doing the work of a long afternoon's simmer.

Continuity, not loudness — and why the number is softer here

The researchers who first pinned kokumi down name it with three words, and together they map the axis exactly: *mouthfulness*, the sense that a flavor fills the whole mouth rather than sitting on the tongue; *continuity*, that it carries unbroken from the first instant to the finish; and *thickness*, a feeling of body and layered depth. None of the three is a taste you could isolate. All three are qualities of how the *other* tastes are delivered. Kokumi is a gain control — the invisible hand that makes a good broth feel like being cared for — and that single fact reshapes how the number is earned.

You cannot ask a panel how kokumi a plain solution is; there is nothing there to rate. So the measurement is always a *difference*. We give tasters a base of umami, salt, and a little sweetness, then the same base with a γ-glutamyl peptide added, and we ask not "how strong" but "how full, how continuous, how thick," and above all how long the finish lasts. The honest readout is not the height of a peak but the **area under the curve of perceived intensity over time** — the slow fade of the finish rather than the loudness of the arrival. Where heat gives you a clean, fast, single-channel burn to calibrate against, kokumi gives you a modulation you can only read against a matrix, by subtraction.

We summarize that modulation as a single figure, the gain coefficient μ — the ratio of a base taste's perceived intensity with the peptide present to its intensity without. In one plain vegetable stock, on one trained panel, adding γ-Glu-Val-Gly to its EC_{50} lifted rated umami from 5.0 to 6.8 on a ten-point scale: $\mu \approx 6.8 \div 5.0 = 1.36$, a 36% deepening of a taste the peptide does not itself carry. The salt in the same bowl rose too, and the finish lengthened by seconds. Change the stock, though, and μ changes with it — a base already dense with glutamate has less headroom to gain — so we are careful never to quote μ as a fixed Index number. The receptor EC_{50} is what we anchor and publish; μ we report only against a named base, as the honest, matrix-bound thing it is.

That difference costs precision, and the Index pays for it in the error bars. The heat axis holds a batch-to-batch coefficient of variation at or below **7%** — tight enough to report heat nearly to the unit. Kokumi cannot match it, and the reason is arithmetic. IFI-K is a *ratio* of two half-activation values, each measured with its own scatter, and when you divide two noisy numbers their relative errors add in quadrature. If the reference EC_{50} carries a 6% coefficient of variation and the candidate carries 9%, the score inherits

> $\mathrm{CV}(\text{IFI-K}) = \sqrt{(6^2 + 9^2)} = \sqrt{117} \approx 11\%.$

Eleven percent, not seven — which is precisely why we report IFI-K to two significant figures and resist the temptation to pretend a kokumi score is as sharp as a Scoville count. Honesty about the

width of a number is a form of respect for the person who has to design with it. The richness is real; the measurement of it is simply harder-won, and the Index says so out loud.

The plants that keep the secret

The organisms that carry this axis are the savory foundations of nearly every cuisine on earth, and in our work they are the quiet engineering beneath the loud axes — the layer that makes a designed sweetness feel natural, a restored umami feel finished, a whole plate feel like it was made by someone who meant it.

ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Garlic	560	K / P	Moderate
Soybean	360	U / K	Frontier
Yeast	240	U / K	Hard
Onion	40	K	Moderate
Vanilla	40	K	Hard

Across those five carriers sit 1,240 constructs — more than a thousand distinct attempts to raise a plant's γ-glutamyl pool without touching the rest of what it is. The alliums lead because they already keep the raw material: garlic's 560 constructs and onion's 40 explore the GGT-and-glutathione axis in the one family that stores these peptides by design, which is the real reason a dish built on sweated onions and garlic feels complete long after their sharpness has cooked away. Yeast, at 240 constructs and reading on both umami and kokumi, is glutathione-rich by nature — the reason yeast extracts have carried savory cooking for a century — and soybean, at 360, pairs the same kokumi with umami because fermentation into miso and shoyu liberates free glutamate and γ-glutamyl peptides together, the two halves of savory depth made by the same slow culture. This is the axis that lets our vegetarian cooking reach the completeness people wrongly assume belongs only to bone and marrow: richness grown into a plant, and drawn from it whole.



Vanilla pods — the world's most complex and most faked flavor, carrying a genuine kokumi mouthfeel that single-molecule vanillin cannot.

And then vanilla, at 40 constructs — the most surprising name on the list, because its complexity includes a true kokumi mouthfeel that is part of why real vanilla and its single-molecule vanillin substitute are not the same experience, whatever a chemical analysis of their headspace says. The missing thing is not an aroma. It is continuity, read at CaSR, that the one-molecule copy was never built to carry.

Of our flagship superhots, only one lights this reader at all, and its profile is worth laying out whole, because kokumi living inside a plant our model predicts to reach approximately 2.95 million SHU is the clearest proof that the axes are genuinely separable.

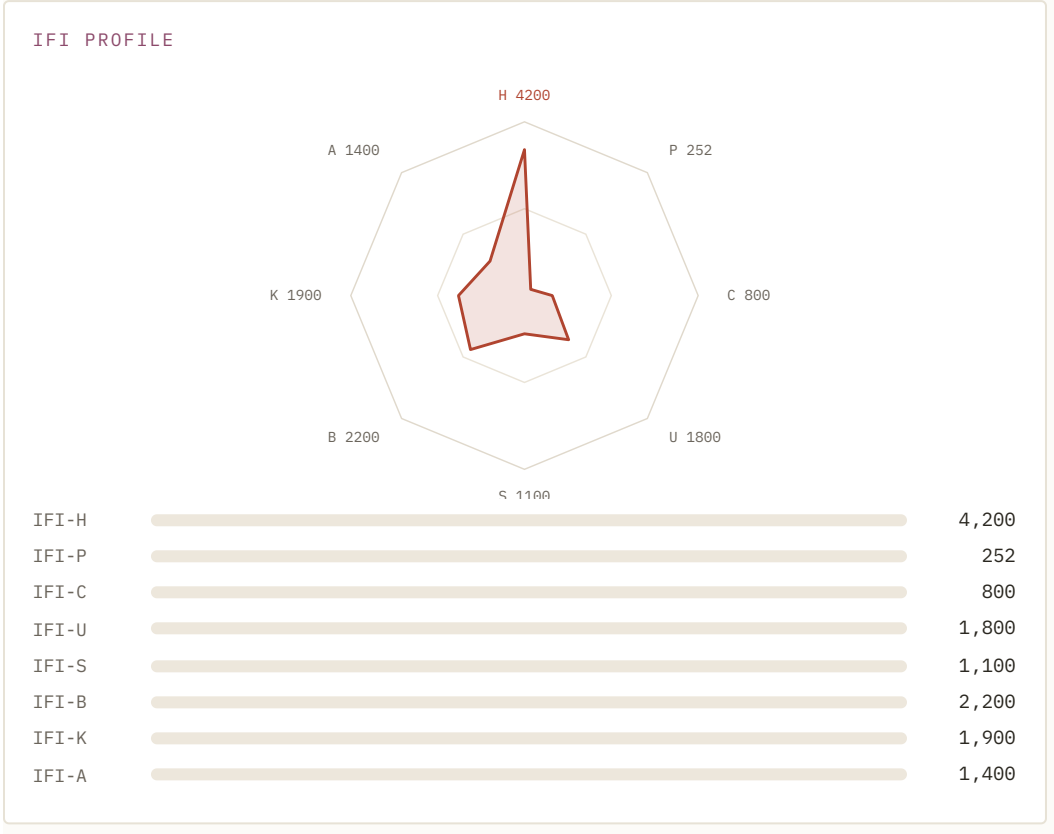

Our model assigns Phoenix **IFI-K 1,900**, and the formula runs just as cleanly backward as forward. Invert it — EC_{50} = 1,000 × EC_{50}(reference) ÷ IFI-K = 1,000 × 2.00 ÷ 1,900 — and the equivalent kokumi load in that tissue computes to an effective half-activation near **1.05 mM** of γ-glutamyl-equivalents: a richness denser in effect than a pure glutathione solution and a shade past γ-glutamylcysteine on the ladder. A pepper that could sear the paint off a wall also quietly carries the continuity of a long-sweated allium, and the Index can state both in the same breath, to the same receptor-anchored standard, without either one bleeding into the other.



Phoenix under the microscope stage — the only one of our flagship cultivars to score on the kokumi axis in our model, at a predicted IFI-K 1,900.

Edge cases, failure modes, and the honest limits

An axis you cannot taste alone is an axis that fights back when you try to measure it, and it is only fair to name where the method strains.

The first and sharpest trap is **calcium itself**. Because Ca^{2+} is a genuine orthosteric agonist of the receptor, a mineral-rich sample will read as kokumi even with no peptide present at all — hard water,

a calcium-set tofu, a bone stock, each nudges the receptor from the floor rather than the lip. Every clean IFI-K we report is measured against a controlled free-calcium background for exactly this reason; skip that control and you will credit a peptide for a mineral's work. The confound is not a flaw in the receptor. It is the receptor doing its original job while we ask it to do its second one.

The second is **saturation**. The steep Hill curve that makes kokumi snap into coherence also means there is a top to it: push the γ-glutamyl pool far past the receptor's half-activation and the pleasant fullness tips into something metallic and faintly soapy, the sensation of a modulator with nothing left to modulate. More is emphatically not better, and a design that chases the highest possible IFI-K rather than the right one will overshoot into a taste nobody wants to finish.

The third is that **IFI-K is a derived, matrix-dependent number in a way a Scoville count is not**. Heat is nearly absolute — capsaicin on TRPV1 reads much the same in water or in oil. Kokumi is a gain applied to other tastes, so its felt magnitude depends on what it is amplifying; the same peptide load that transforms a thin broth barely moves a dish already dense with glutamate. We anchor the receptor-level EC_{50} to make the score reproducible, but we are honest that the *eating* of kokumi lives in a matrix, and that the modulation coefficient — how much a given CaSR engagement lengthens and fuses the other axes — is reported for a stated base, not as a universal constant. Add to that the real **inter-individual variation** in CaSR, the inherited differences that make one taster read richness where another reads only salt, and you have an axis that is true, repeatable at the bench, and still irreducibly relational at the table.

We make our peace with all of it, because the alternative is to pretend kokumi is simpler than it is. It is the most democratic of the eight readers — made by time, by fermentation, by the alliums any kitchen already keeps, and by the calcium in the water — and felt by everyone, named by almost no one. Of the eight axes this is the one most people have tasted their whole lives and never once identified, and the one whose absence they feel most keenly without knowing what to call the lack: the thin, disjointed quality of a dish that has umami and salt and sweetness in correct amounts and still, somehow, does not cohere. Add the γ-glutamyl peptides and it closes up, the way a chord resolves into its tonic.

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

> Kokumi has no sound of its own. *It is the resonance that makes every other note ring longer.*

Chapter 9 · TRPM8 — *cold without cold.*

Breathe in after a mint, and the cool air feels colder than the warm air before it, though nothing about the room has changed. This is the exact mirror of the capsaicin trick, run in the opposite direction: a molecule that convinces your nervous system of a temperature that is not there — not a false fire, but a false frost. The receptor that reports it, **TRPM8**, is the body's cold sensor, and menthol picks its lock as neatly as capsaicin picks the heat channel's. We find ourselves returning to this receptor more than any other in our cooling work, because it is the one that taught us the discipline the whole family now keeps: that a flavor is not a quantity of a compound but a shape arriving at a lock, and that the cold on your tongue after a leaf of mint is a geometry, not a temperature.



Peppermint in leaf — the glandular trichomes on its surface hold the menthol that opens TRPM8 and makes the mouth report a cold that was never in the room.

The channel that reports winter

TRPM8 is a **transient receptor potential** channel of the melastatin family — a large tetramer, four identical subunits of roughly eleven hundred amino acids each, assembled around a single central pore. Each subunit folds into six membrane-spanning helices, S1 through S6. The last two, S5 and S6, from all four subunits together line the ion gate at the center; the first four, S1 through S4, bundle off to the side as a *voltage-sensor-like domain*, a structure borrowed from the voltage-gated channels that fire nerves and hearts. The conduction path has two constrictions in series: an outer selectivity filter at the

pore mouth, only weakly selective — it passes sodium and calcium alike, with a modest preference for calcium — and a lower gate where the four S6 helices cross and seal the pore shut with a ring of oily residues. Closed, that hydrophobic seal is impassable to ions not because it is physically too narrow but because water will not wet it; opening dilates the bundle crossing by a fraction of a nanometre, the seal breaks, and the pore floods. When the gate opens, sodium and calcium rush inward, the nerve ending depolarizes, and a cold message leaves for the brainstem. Nothing in that message carries the word *menthol*; the brain only ever learns that a cold fiber fired.

The receptor is **polymodal**, which is the quiet marvel of it. One channel integrates three different inputs into one read. It opens to **cold**, below roughly 26 °C. It opens to **membrane voltage**. And it opens to chemical agonists — chief among them **menthol** from the mints and the synthetic super-cooler **icilin**. The chemical pocket is not at the gate but tucked inside the voltage-sensor bundle, where menthol slots into a small cavity formed by S1–S4; a single residue there, a tyrosine partway up the second helix, is so central that changing it one letter leaves the channel cold-sensitive but almost blind to menthol. The molecule does not force the gate. It settles into that side cavity and shifts the energetics of the whole protein so that the gate, left to itself, now prefers to be open.

Two more dependencies make the cold read honest to the state of the cell. TRPM8 needs the membrane lipid **PIP_2** beneath it to open at all — the same lipid whose depletion helps shut the heat channel — so the cooling signal is tied to the cell's own lipid housekeeping and cannot run away unsupervised. And icilin, unlike menthol, demands **intracellular calcium** to do its super-cooling work, which is why icilin's cold arrives with a flicker and a lag that menthol's does not. These are not footnotes. They are the reasons a designed cool can be made clean and self-limiting rather than relentless, and we lean on every one of them.

The thermodynamics of a feeling

Here is the physics, because it is beautiful and because it is the whole mechanism of the illusion. Temperature, voltage, and ligand do not each open the channel by a separate door. They converge on a single quantity — the channel's **open probability**, the fraction of the time the gate stands open — and they trade against one another in one shared account of free energy.

Write the open probability as a function of voltage in the two-state form every gated channel obeys:

```
P₀(V) = 1 / (1 + e^(−zF(V − V₁′₂)/RT))
```

Here z is the effective gating charge the membrane field acts on, F is Faraday's constant, R the gas constant, T the absolute temperature, and $\mathbf{V_{1'2}}$ the half-activation voltage — the voltage at which the gate is open exactly half the time. At body temperature, 310 K, the term RT/F works out to (8.314 × 310) / 96,485 = **26.7 mV**, the natural voltage scale of the whole relationship.

TRPM8's gating charge is small, near $z \approx 0.9$, so voltage alone is a weak lever; what matters is where cold and menthol push $\mathbf{V_{1/2}}$. Warm and unliganded, the channel sits with a wildly depolarized half-activation, out past +120 mV — a voltage no resting nerve ever reaches — so at a resting −60 mV the gate is all but shut. Put the numbers in:

> exponent = −0.9 × (−60 − 120) / 26.7 = +6.07, so P_0 = 1 / (1 + e^6.07) ≈ **0.0023.**

Two parts in a thousand. The channel is silent; the mouth reports no cold, which at body temperature is exactly correct.

Now flood the cavity with menthol. Menthol's entire action is to slide $\mathbf{V_{1/2}}$ leftward — to make the gate willing to open at voltages and temperatures it would otherwise ignore. At a saturating dose the half-activation can march from +120 mV down past −40 mV. Recompute at the same resting −60 mV and the same body temperature:

> exponent = −0.9 × (−60 − (−40)) / 26.7 = +0.67, so P_0 = 1 / (1 + e^0.67) ≈ **0.34.**

From two parts in a thousand to better than one in three. A channel that was mute now fires a third of the time, at normal body temperature, with no cold anywhere in the room — and the brain, which only ever hears *a cold fiber fired*, believes it completely. Real cold does the identical thing by the identical arithmetic. Lower the temperature toward 10 °C and $\mathbf{V_{1/2}}$ slides left on its own, no menthol required; the voltage scale itself tightens, because RT/F at 283 K is only 24.4 mV. Carry the half-activation to −30 mV by cold alone and recompute at the same resting −60 mV: exponent = −0.9 × (−60 − (−30)) / 24.4 = +1.11, so P_0 = 1 / (1 + e^1.11) ≈ **0.25.** A genuine chill near 10 °C and a mouthful of menthol drive the one channel to open probabilities of the same order — a quarter to a third of the time — and the fiber fires alike for both. The channel cannot tell the two apart. Neither can we. That is not a flaw in the sense; it is the sense working perfectly on a molecule built to exploit it.

> Capsaicin tells the body it is burning. Menthol tells it it is freezing. *Both are lying, and the body believes them completely.*

Where capsaicin lowers a heat threshold to make warm feel scalding, menthol lifts a cold threshold to make warm feel chilled. The two most dramatic flavor illusions the mouth can host are the same trick in mirror image, performed on the two channels that bracket comfortable temperature from above and below.

Reading the cold as a number

A feeling this precise deserves a number, and the Index gives it one. Every axis is anchored the same way: a single score tied not to how *much* of a compound a food holds but to how *loudly* that compound

speaks to its receptor. The IFI-C reference is a stated convention — a fixed anchor defined by our platform, like the dilution-to-threshold definition of the Scoville scale — not a literature-derived value.

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

EC_{50} is the concentration that drives the receptor to half of its full response. It is read off the dose–response curve, which for TRPM8 and menthol is cooperative — a tetramer with more than one place to bind behaves with a Hill coefficient near 2:

$$\text{response} = R_{max} \cdot [L]^n \,/\, (\mathrm{EC}_{50}{}^n + [L]^n),\ n \approx 2$$

Take a characteristic menthol EC_{50} of **80 µM** on our calcium-imaging assay as the cool-burn reference, and the curve is steep and legible. At 20 µM the channel runs at 400/(6,400 + 400) ≈ 6% of full response; at 40 µM, 20%; at 80 µM, exactly half, by definition; at 160 µM, 80%; by 320 µM it is saturated near 94%. Doubling the dose near the foot of the curve nearly quadruples the cold — which is why a strong mint overwhelms so suddenly, and why a designed one must be tuned, not merely amplified.

It helps to leave molarity for the kitchen's units at least once. Menthol's molar mass is 156.27 g/mol, so an 80 µM reference is 80 × 10^{-6} mol/L × 156.27 g/mol = 0.0125 g/L = **12.5 mg/L**, about 12.5 ppm, or 0.00125% by weight. Our engineered peppermint is designed to reach half-activation near 20 µM — 3.1 mg/L — so a quarter of the menthol speaks as loudly, because what little it makes is the exactly correct shape.

The score is built on a consistent assay foundation. We express human TRPM8 in a clean cell line, load the cells with the ratiometric calcium dye Fura-2, and step a dose series of the candidate across them while reading the 340-over-380-nanometre fluorescence ratio — a direct, dye-based proxy for the calcium the open channel lets in. Each plate is normalized to its own saturating response, so a day's variation in how much dye the cells took up cancels out of the ratio. We fit the resulting points by nonlinear least squares to the Hill equation, which hands back both the EC_{50} and the slope *n* in one pass, and we run every candidate in independent triplicate so that the spread across plates becomes the standard deviation — and, divided by the mean, the coefficient of variation the whole platform is held to. A score without that spread attached is half a measurement.

Fix the reference at EC_{50} = 80 µM, hold the assay at 25 °C and −60 mV so every candidate is read under identical conditions, and the formula turns potency into a score anyone can read on one page.

STIMULUS AT TRPM8	EC_{50} (MM, ASSAY)	IFI-C = 1,000 × 80 / EC_{50}	NOTE
(−)-menthol — reference	80	**1,000**	the anchor, by definition
Engineered peppermint (menthol-equiv.)	20	**4,000**	our deepest cool program
Field peppermint (menthol-equiv.)	160	500	a gentle garden mint
(+)-menthol — the wrong hand	480	167	same atoms, inverted geometry
WS-23 — odorless synthetic coolant	40	2,000	the cool, with the smell removed
Eucalyptol (1,8-cineole)	3,800	21	weak at TRPM8; its "cool" is mostly aroma
Icilin — synthetic super-agonist	0.36	~222,000 → **capped 5,000**	off the scale; needs intracellular Ca^{2+}
Phoenix cool fraction · SS-0096	100	**800**	the composite read beneath the heat

Two lines in that table carry the chapter's argument. **Icilin** is roughly two hundred times more potent than menthol, and its raw IFI-C runs past 222,000 — far beyond the per-axis normalization ceiling of 5,000 that every scorecard honors. We cap it, and we do not put it in food. A super-agonist that saturates the Index also saturates the tongue, flickers with the cell's own calcium, and overstays; it is a laboratory tool, not a kindness to the body, and the clean cold we want is the one that arrives, speaks, and leaves. The second line is **(+)-menthol**, the mirror image of the real thing — same atoms, same order, opposite hand — reading 167 against menthol's 1,000. Nothing was lost but geometry, and geometry was almost everything.

Eight shapes, one of them cold

Menthol carries three stereocenters. Three handed carbons give 2^3 = **eight** stereoisomers — the same atoms in the same connectivity, differing only in their arrangement in space — grouped as four diastereomers, each a pair of mirror images. Of the eight, only one delivers the clean, sweet, unmistakable cold. The receptor's pocket is itself handed, and it accepts only the hand cut to fit.

DIASTEREOMER	ENANTIOMER	POTENCY VS (−)-MENTHOL	CHARACTER OF THE COOL
menthol	**(−)** · 1R,2S,5R	**1.00**	clean, sweet, deep — the one
menthol	(+) · 1S,2R,5S	0.17	weak, faintly dusty
neomenthol	(±)	≤ 0.12	thin, herbaceous, a bitter edge
isomenthol	(±)	≤ 0.08	musty, medicinal
neoisomenthol	(±)	≤ 0.05	near-mute, green

This is the whole thesis of the book compressed into one molecule. The atoms are identical across all eight; the *geometry* is everything; flavor is a lock reading a shape. When we tune a mint we are not asking for more terpenoid — we are steering flux toward the single correct stereoisomer and away from the seven off-forms and the

neighboring molecules (**menthone**, **isomenthol**, **pulegone**) that muddy it. A better mint is not a mint with more oil. It is a mint whose oil is more exactly the one shape the lock was cut for.

That shape is the end of a short, elegant terpenoid pathway the platform tunes as a single unit. It begins with geranyl diphosphate, passes through **limonene** — the first committed ring, laid down by limonene synthase — and then runs a disciplined sequence of oxidations and reductions that set each of the three stereocenters in turn, arriving at **(–)-menthol**. Miss a single reductive step and the flux pools at menthone or pulegone instead; let the wrong reductase dominate and it drifts to neomenthol. The art is not in making the molecule. It is in making the *right* one, eight times out of eight.

Many cools, not one

"Cooling" undersells the sensation, because there is no single cool. A pure TRPM8 read is clean and dry — the cold of peppermint, of a breath of winter air. The camphoraceous, medicinal cool of eucalyptus is a different thing entirely, and the sweeter, rounder cool of spearmint is different again, because each plant presents its own blend of agonists and accompanying aromatics, and each "cool" is composed, not given. The cooling plants are not interchangeable, and the Index shows exactly why.

COOLING PLANT	CONSTRUCTS	DIFFICULTY	DOMINANT AGENT	THE COOL IT GIVES
Peppermint	2,360	Hard	(–)-menthol	clean, dry, winter-air
Lavender	160	Hard	camphor + linalool	soft, floral, a low cool
Spearmint	80	Hard	(R)-carvone	sweet, round, gentle — weak at TRPM8
Eucalyptus	80	Hard	eucalyptol	medicinal, penetrating; aroma-led
Camphor	40	Hard	camphor	sharp, counterirritant, cold-hot

Spearmint is the quiet lesson in that table. Its signature terpene is not menthol at all but **(R)-carvone**, a weak TRPM8 agonist, which is precisely why spearmint cools so much less than peppermint and leans instead on sweetness and aroma to read "cool" at all. Carvone makes the same point menthol does, in miniature: (R)-carvone is the smell of spearmint, while its mirror image, (S)-carvone, is the smell of caraway — the same atoms, the opposite hand, two unmistakably different plants on the tongue. The nose, like the cold channel, reads geometry and nothing else. Eucalyptol tells the same story from the other side: a modest direct agonist, scoring a mere 21 on our assay, so most of eucalyptus's famous chill is actually its smell, read by the nose and *called* cold by the brain out of long association. The flavor industry proved the point decades ago with the **"WS" compounds**, first made by a razor company chasing a cooling aftershave: synthetic coolants engineered to hit TRPM8 with no minty smell whatsoever. WS-23 scores 2,000 on the same reference and arrives

as pure, odorless cold — proof that the cooling and the aroma are separable signals the receptor and the nose read independently, and therefore two dials we can set apart.

Peppermint is our deepest cool work — 2,360 constructs, every one of them an attempt to make the single correct molecule more purely and more abundantly. A representative radar for our engineered peppermint reads almost entirely on one axis, the shape of a plant built to do one thing exactly:

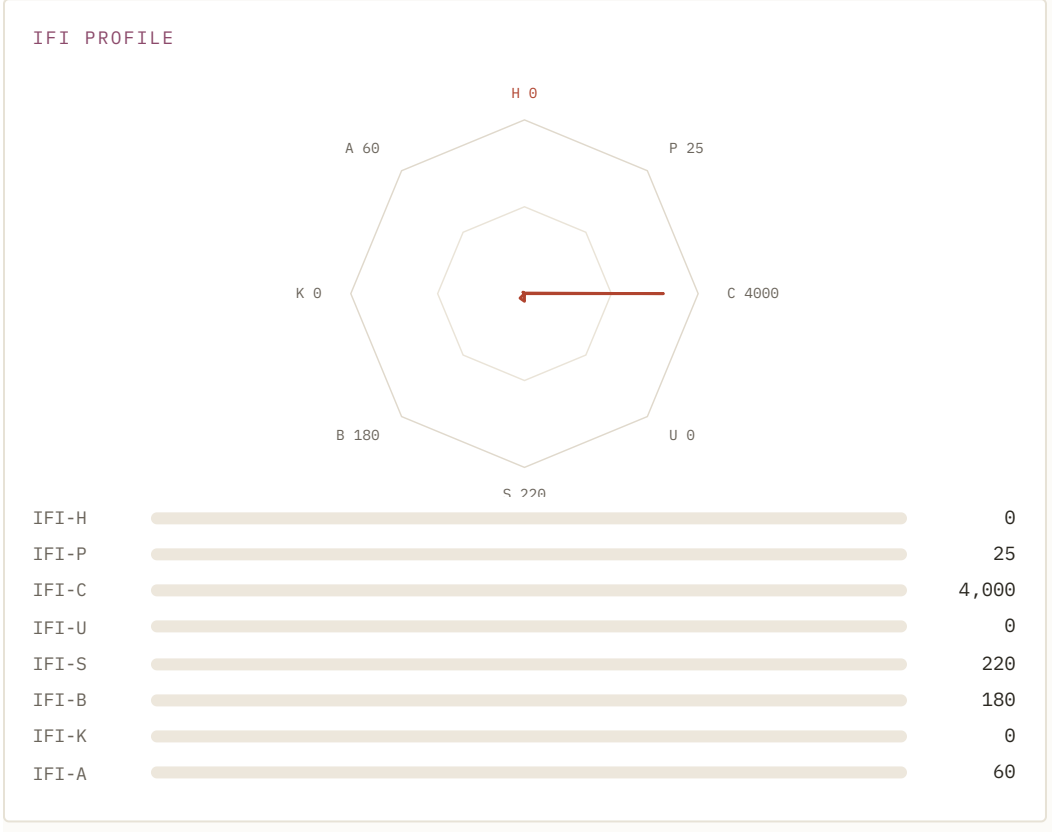

Designing the cooling axis, then, is largely composition: how clean, how sweet, how medicinal, how lasting the frost should be, and whether any warmth should sit beneath it. The hard part is never the strength. It is the purity.



Eucalyptus — a different cool entirely: eucalyptol is a weak direct agonist at the cold channel, so most of this plant's famous chill is aroma the brain has long since learned to call cold.

Error, consistency, and the honest limits

A score is only worth the error bars around it, and this is where consistency stops being housekeeping and becomes the whole credibility of the Index. Because IFI-C is a ratio of two concentrations, its relative error is the two relative errors added in quadrature:

$$\sigma(\text{IFI-C})/\text{IFI-C} = \sqrt{\left[(\sigma_{ref}/EC_{50_ref})^2 + (\sigma_{cand}/EC_{50_cand})^2 \right]}$$

Hold both measurements to our platform's batch coefficient of variation of **≤7%**, and the score inherits $\sqrt{(0.07^2 + 0.07^2)} \approx$ **9.9%**. An engineered peppermint reading IFI-C 4,000 therefore carries roughly ±400 — tight enough that 4,000 and 3,600 are genuinely different cools, and a formulator can design to the number. Run the same arithmetic on a wild-type batch, whose capsaicinoid-pathway cousins drift at a CV near **42%**, and the score inherits $\sqrt{(0.42^2 + 0.07^2)} \approx 43\%$ — a peppermint "4,000" that might honestly be anything from 2,300 to 5,700. A number with that much slack is not a measurement; it is a mood. The ≤7% discipline is what lets the cold be a figure you can stand on.

The honest limits are worth naming plainly, because a reverent book is still a careful one. The deep mechanism of *cold* transduction — whether the chill reaches the gate through the pore, the voltage-sensor bundle, or the surrounding lipid — is still genuinely debated in the literature; what is settled is where **menthol** binds and how it shifts the gate, which is the part we engineer. EC_{50} itself is conditional: it slides with holding temperature and voltage, so a score is only comparable against scores read under the same fixed assay, which is why we pin ours to 25 °C and −60 mV and publish the conditions. Scoring a cool-burn food raises a subtler problem still, because two channels share one mouth: a chili crisp drives TRPM8 and the heat channel together, and a naïve reading would let the fire bleed into the frost. We deconvolve it on the bench — running the candidate twice, once against a TRPM8-selective blocker and once against a heat-channel blocker, and taking the difference — so that the published IFI-C reflects genuine cold drive alone. The steep slope of the curve is its own quiet tax, too: with *n* near 2, a ten-percent error in menthol near the foot of the dose–response moves the perceived cold by nearly twenty, so the consistency the Index demands is not a luxury but the price of the sensation sitting where you designed it. And three failure modes shadow every cooling build. The first is **stereochemical leakage** — flux spilling to neomenthol, isomenthol, or the wrong enantiomer — which does not merely weaken the cool but dirties it, and is the single hardest thing to hold. The second is **pulegone**, a near-neighbor on the pathway that is genuinely hepatotoxic; feeding the body right means tuning flux decisively away from it, every time, with margin to spare — a nourishment constraint before it is a flavor one. The third is **desensitization**: lean too hard on menthol and the cold fibers dull with repetition, the cool fading even as the dose holds, so a mint designed for a second mouthful must be gentler than one designed for a first.

The cool held against the fire

Cool-burn is, by design, a *composite* axis: IFI-C reads both the TRPM8 cooling and the warmth it can be set against, because the most interesting cold experiences are the ones held in tension with heat. A numbing-cool chili, a mint with a low ember beneath it, the tingling *mála* of a Sichuan oil — these are two channels, TRPM8 and the heat channel, played at once, the mouth reporting frost and fire

from the same bite and finding the contradiction delicious. That duet, frost and fire resolving on one tongue, is a study worth its own long sitting.

Our fullest proof of it is a pepper, not a mint. **Phoenix** — SS-0096, our model predicts approximately 2.95 million SHU, a plant designed gene by gene to activate all eight axes at once — carries a heat of 4,200 and, beneath it, a deliberate cool-burn of **800**, the menthol-equivalent frost we tuned in to sit under the fire like cold stone under a flame:

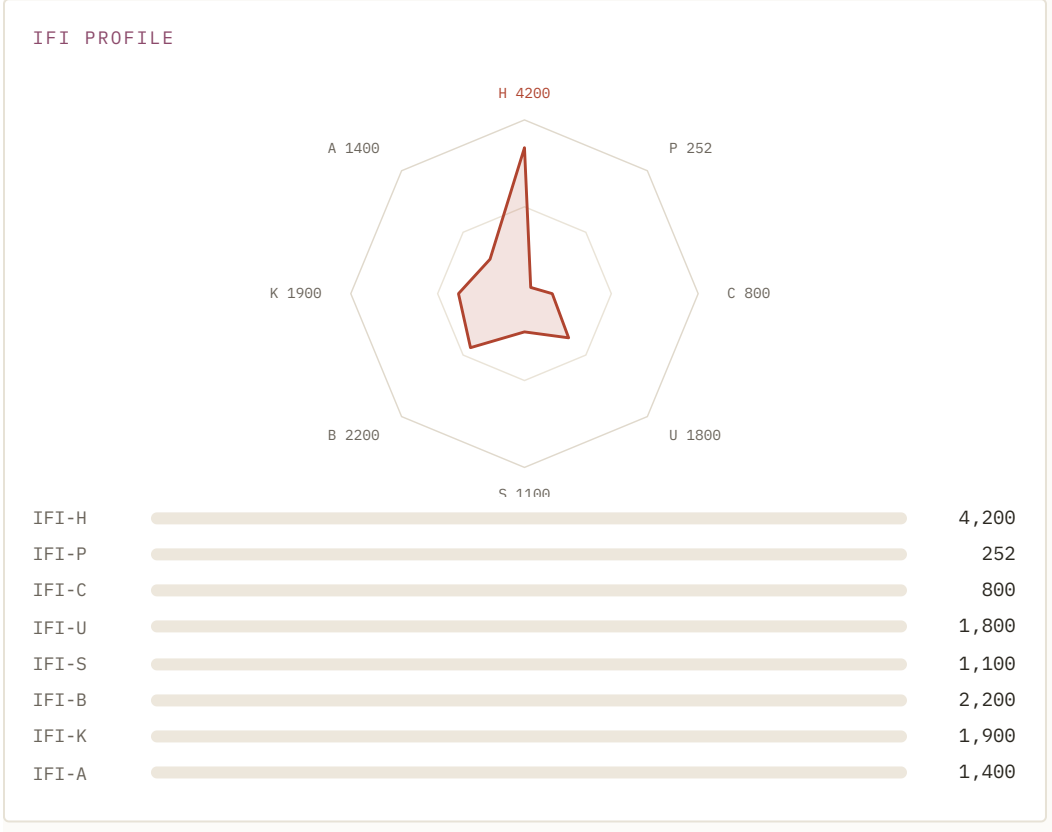

That cool read of 800 is the same EC_{50} 100 µM line from the worked table — not an accident of the plant but a value we chose and placed, the cold designed to arrive a half-beat behind the heat and keep the burn from becoming only burn.



Phoenix · SS-0096 — a superhot carrying a designed cool of 800 beneath a heat of 4,200: frost placed deliberately under fire, so the burn resolves instead of merely escalating.

The cold is also, quietly, a medicine of its own, and this is where the family's care for the body meets the chemistry. (–)-menthol is the counterirritant in every muscle balm and cough lozenge, working partly by the same desensitization that tames a chili and partly because the cold fibers it opens can gate pain signals in the spinal cord. That gating is a real circuit, not a metaphor: the cool-sensitive fibers menthol recruits feed inhibitory interneurons in the dorsal horn, which damp the slower pain traffic arriving alongside them, so the cold is quite literally crowding out the ache at the first synapse

before it ever reaches the brain. And its illusion lingers, kindly: for a minute after a mint, plain cool water tastes colder than it is, because the TRPM8 channels are still sensitized, still reporting a winter that the mint, not the water, put there.

We began with a breath of cool air that felt colder than it was, and we end with the reason. Nothing in the room changed; what changed was the shape sitting in a pocket on a nerve, nudging a single number — the open probability — until the body reported a cold it could not find. The whole of this work, the two thousand three hundred and sixty constructs behind one clean mint, comes down to making that shape exactly, and giving the body a cold it can trust: clean, correct, self-limiting, and kind. That is the only cold worth designing. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 10 · TRPA1 — *the sharp channel.*

There is a kind of heat that goes up rather than down — the wasabi that climbs behind the nose, the raw garlic that catches at the back of the throat, the mustard that flares in the sinuses and vanishes as fast as it came. It is not the slow deep burn of a chili, and it is not read by the chili's receptor. It is **pungency**, read by **TRPA1**, and the way TRPA1 is activated is unlike anything else the tongue does: the molecule does not merely settle into the channel and wait. It *bonds* to it. Where the heat of a pepper is a lock turned by a key that leaves again, pungency is a key that fuses to the lock — a genuine chemical reaction, measured in the same currency as every other axis of the Index, but arriving by a mechanism all its own.

We love these plants the way a cook loves them: as the sharp spine of a vegetarian table, the garlic and ginger and mustard that give a meatless pot its backbone and its bite. For most of our lives we could taste exactly what they did and say only that they were "sharp," "clean," "biting." The purpose of this chapter is to make that word a number — to follow the sharpness all the way down to the atom that reacts, and all the way back up to a score you can set on purpose.



Black mustard in seed and leaf — crush it with water and the enzyme myrosinase liberates allyl isothiocyanate, the electrophile that bonds covalently to TRPA1.

The channel that is also a spring

TRPA1 is a TRP channel — tetrameric, cation-permeable, a cousin of the heat sensor TRPV1 and the cold sensor TRPM8 — four identical protein subunits folded together around a central pore that, when it opens, lets sodium and calcium pour into the nerve ending and fire it. That much it shares with its relatives. What it does not share is the enormous structure it carries on its intracellular side: an **ankyrin repeat domain** of roughly sixteen stacked structural hairpins, the largest such domain of any TRP channel, and the feature that gives the receptor its name — the *A* is for ankyrin. Picture a long molecular

spring coiled beneath each subunit, four of them bundled under the pore. That spring is not decoration. It is the sensor, and it is wound to be sprung by chemistry.

On the channel's intracellular N-terminus, where the spring anchors, sits a small cluster of reactive **cysteine** residues — principally **C621, C641, and C665** in the human channel. Cysteine is the one amino acid whose side chain ends in a sulfur-hydrogen bond, a thiol, and a thiol is nucleophilic: its sulfur is electron-rich and hungry to attack anything electron-poor. In most proteins these cysteines would be buried or paired off into stabilizing disulfide bridges. In TRPA1 they are deliberately exposed, their sulfurs presented to the cytoplasm like tripwires. The channel is, in effect, a chemical alarm whose trigger is a set of sacrificial sulfur atoms waiting to be modified. The modification is only the trigger; the force has to travel. Each adduct stiffens and reshapes the ankyrin spring that carries it, and that change is relayed through the linker coupling the spring to the pore's gate — so a reaction at an intracellular tripwire is felt, finally, as a mechanical pull on the gate that lets the ions through.

There is a second amplifier built into the same machine. TRPA1 is permeable to calcium, and the calcium that floods in through an open channel binds back onto the channel's own regulatory sites and *potentiates* further opening — a feed-forward loop in which the first trickle of current recruits more current. It is part of why pungency, once it starts, arrives with a rush rather than a fade-in: the channel helps to open itself. The same loop also sets a limit, because sustained calcium eventually drives the channel into a desensitized, refractory state — so the rush is self-amplifying at its start and self-terminating at its end, a pulse shaped by the very ion it admits.

A bond, not a handshake

The ligands that spring this channel are **electrophiles** — "electron-loving" molecules, electron-poor at some atom and therefore drawn to the electron-rich sulfur of a cysteine. When an electrophilic agonist meets the exposed cysteine cluster, it does not dock and release. It reacts. A new covalent bond forms, the spring is pulled, and the pore opens. Three distinct reactions do this work, depending on the molecule:

- **Isothiocyanates** — the sharp principle of mustard, wasabi, and horseradish — carry a highly electrophilic central carbon in their –N=C=S group. A cysteine thiol attacks that carbon and forms a **dithiocarbamate** adduct, a stable thioether-like linkage.
- **α,β-unsaturated carbonyls** — **cinnamaldehyde** from cinnamon, **acrolein** from smoke and scorched fat — present a carbon-carbon double bond conjugated to a carbonyl, a classic **Michael acceptor**. The thiol adds across the double bond in a Michael addition.
- **Thiosulfinates** — **allicin**, made the instant garlic is cut — carry a reactive S–S(=O) group that undergoes **thiol–disulfide exchange** with the cysteine, swapping partners and leaving its mark.

In every case the key chemically alters the lock. This is categorically different from the reversible, lock-and-key binding that reads heat or sweetness, where the molecule can leave as readily as it came and the receptor resets the moment it diffuses away. Here the receptor must *repair* itself. The cell resolves the adducts over seconds to minutes — glutathione and other cellular thiols reduce or displace them, and the neuron renews its cysteines — which is exactly why pungency is a flare and not a plateau: fast, bright, and then gone, on the timescale of the chemistry that undoes it. It is also why these molecules are so volatile and so good at reaching the nose. The same small, reactive compounds that bond to a cysteine on the tongue evaporate readily into the air of a mouthful and find the trigeminal nerve endings lining the upper airway, where a chili's heavy, oily capsaicin never travels. That is the whole reason wasabi "goes up your nose" and a pepper never does.

Because the activation is a reaction and not a binding, potency has two ingredients a reversible agonist lacks: the intrinsic reactivity of the electrophile, and the time it dwells against the cysteines. The behavior mirrors an irreversible enzyme inhibitor, whose effect is governed by a second-order constant — reactivity multiplied by residence — rather than a single equilibrium affinity. A mildly reactive electrophile given long enough can reach the same adduct burden as a ferociously reactive one in an instant, which is why a slow-building mustard and a flash of fresh wasabi can arrive at the same place on the ruler by different roads, and why exposure time belongs inside the measurement rather than beside it.

> A chili binds its receptor and lets go. A mustard *bonds* to its receptor and burns the bridge. *One is a visitor; the other leaves a mark.*

The weapon assembled at the bite

The most beautiful thing about these molecules is that, in the living plant, most of them do not yet exist. They are made at the moment of injury — the plant's own defense, and a masterclass in compartmentalization that we have had to learn to respect and to re-engineer.

In mustard and its relatives, the pungent isothiocyanate is absent from the intact leaf. The plant keeps a stable, odorless **glucosinolate** — sinigrin, in black mustard — locked in one cellular compartment, and the enzyme **myrosinase** that will dismantle it locked in another. Nothing happens while the tissue is whole. Only when a cell is crushed — by a chewing insect, by a grater, by a tooth — do the two compartments break open and mix, myrosinase hydrolyzes the glucosinolate, and the reactive electrophile is manufactured on demand, in the mouth of whatever bit the plant. Garlic runs the same play with different reagents: intact cloves hold odorless **alliin** apart from the enzyme **alliinase**, and the knife that cuts them produces **allicin** within seconds. The sharpness of these foods is, quite literally, a wound response — a chemical weapon that does not exist until the instant of damage. It is freshest at the cut and decays as the

reactive product breaks down, which is why a crushed garlic sitting ten minutes mellows, and why a jar of prepared horseradish is a shadow of the fresh root.

The two mustards split the difference by chemistry alone. Black mustard's sinigrin yields the small, volatile allyl isothiocyanate that leaps into the nose; white mustard's sinalbin yields the heavier, non-volatile p-hydroxybenzyl isothiocyanate that never climbs to the sinuses and instead lingers as a slow heat on the tongue. Same enzyme, same reaction, two molecular weights — 99.15 against 165.21 — and therefore two entirely different sensations, the clearest proof that pungency is set by the product and not by the plant.



Garlic in cross-section — intact cells hold alliin and alliinase apart; the cut mixes them, and allicin, the electrophile TRPA1 reads, is made within seconds.

The chemistry is not even fixed once the plant is picked. Fresh ginger owes its warmth to **gingerol**; dry that ginger, or heat it, and gingerol loses a molecule of water and becomes **shogaol**, a sharper and more pungent compound. The arithmetic is exact: 6-gingerol has a molecular weight of 294.39 g/mol, water removes 18.02, and 294.39 − 18.02 = 276.37 g/mol, the mass of 6-shogaol. The dehydration rewrites the key, which is why dried ginger *bites* where fresh ginger merely warms. The lesson for us is that to tune pungency honestly we cannot tune only the plant. We have to tune the stored precursor, the enzyme that liberates it, the compartments that keep them apart until the bite, *and* the knife, the heat, and the drying rack that will finish the food after it leaves us.

Reading loudness — from dose–response to one number

All of this chemistry has to resolve into a single scalar, because the Index scores every axis the same way and pungency is no exception. The bridge from biology to number is the dose–response curve.

Expose the channel to an agonist at concentration $[A]$ and measure the fraction of maximal response it produces. The relationship follows a Hill function:

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

where EC_{50} is the concentration that drives the receptor to exactly half of its full response — set $[A] = EC_{50}$ and the equation returns one-half — and n, the Hill coefficient, measures how steeply the

curve rises. For a sensor like TRPA1, with several reactive cysteines on each of four subunits all contributing to one opening, the effective n is often steeper than one: the sites act together, so response climbs faster than a simple single-site binding would predict. The table that follows shows the same curve at $n = 1$ and $n = 2$, written as the fraction activated at multiples of EC_{50}.

[A] ÷ EC_{50}	FRACTION ACTIVE (N = 1)	FRACTION ACTIVE (N = 2)
0.25	0.20	0.06
0.50	0.33	0.20
1.00	0.50	0.50
2.00	0.67	0.80
4.00	0.80	0.94

The steeper $n = 2$ column is the designer's friend and the designer's hazard at once: cooperativity means a small change in delivered electrophile moves the response a long way near threshold, so the setpoint is sensitive — a gift for precision, a trap for sloppiness.

EC_{50} is the number we actually hang the axis on, because it captures *loudness* rather than mere quantity. A molecule that reaches half-activation at a low concentration is a loud agonist; one that needs a high concentration is a quiet one. To turn loudness into a score, we invert EC_{50} against a fixed reference and scale by a thousand:

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

The reference is anchored so that a candidate matching it exactly scores **1,000**. Halve the candidate's EC_{50} — make it twice as potent — and the score doubles to 2,000. Double its EC_{50} and the score falls to 500. For the pungency axis, the reference constant is **allyl isothiocyanate (AITC) at 2.1 µM** — a stated convention defined by this index, in the same spirit as Scoville's own dilution-to-threshold: IFI-P 1,000 is what a compound with that exact half-activation concentration scores by definition, not a value derived from a single literature measurement.

A worked case grounds it. Our full-spectrum flagship **Phoenix** (SS-0096) is designed to achieve a predicted pungency of **IFI-P 252**. Running the formula backward:

$$252 = 1{,}000 \times 2.1\ \mu M \div EC_{50}(\text{Phoenix})$$
$$EC_{50}(\text{Phoenix}) = 1{,}000 \times 2.1\ \mu M \div 252 = 8.33\ \mu M$$
$$EC_{50}(\text{Phoenix}) \div EC_{50}(\text{ref}) = 8.33\ \mu M \div 2.1\ \mu M \approx 3.97$$

so Phoenix's designed pungent principle reaches half-activation at approximately 3.97× the reference concentration — a calibrated sharpness that complements rather than competes with the cultivar's seven other active axes. That one number sits inside its full eight-axis profile, the only cultivar we have designed to register on every channel at once:

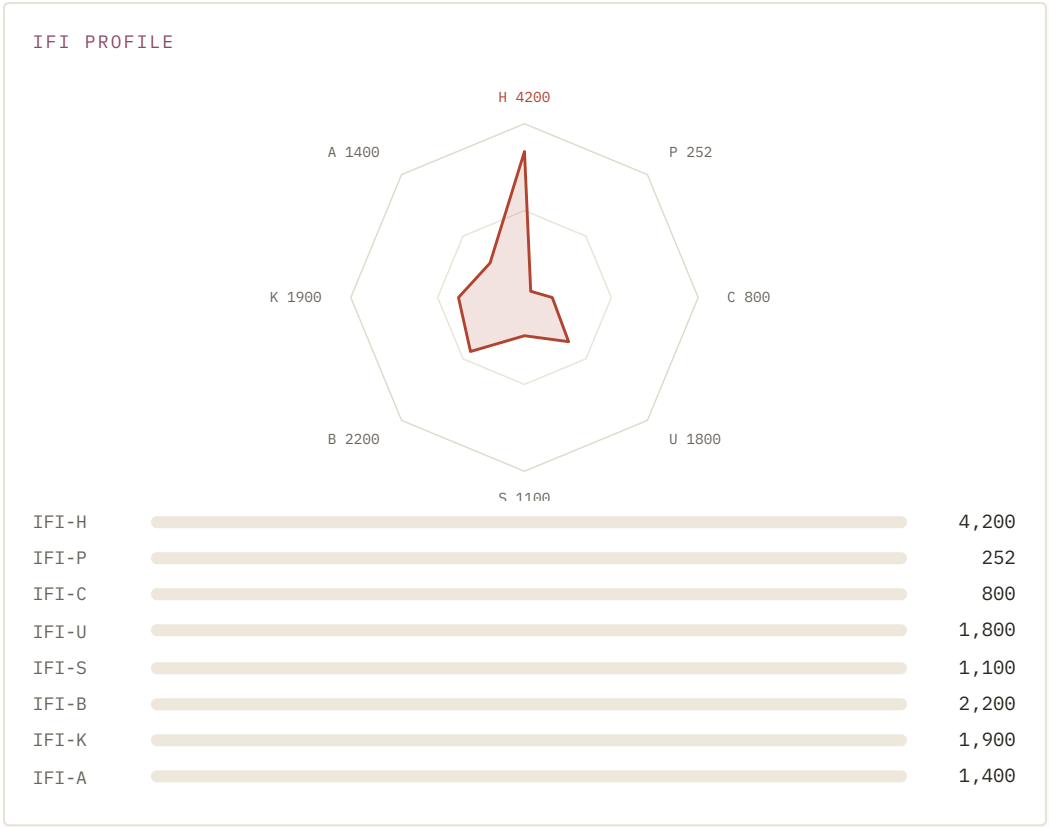

A pepper with a genuine pungency reading is itself a small marvel — heat and sharpness are different channels, and Phoenix was built to speak both — but the point here is the method, not the plant. Any food that activates TRPA1 can be placed on the same ruler.

The formula runs forward just as cleanly. A brief that asks for a pungency of IFI-P 1,500 is asking, in receptor terms, for an electrophile whose half-activation concentration is EC_{50}(ref) ÷ 1.5 = 0.667 of the reference — a target the three design levers then have to hit and hold within tolerance. The number on the page and the molecule in the cell are the same statement, written twice.

A ladder of sharp plants

The pungent plants are the condiment backbone of the world's kitchens, and nearly all of them we can now address with constructs on the platform. Each one delivers a different electrophile, made from a different precursor by a different enzyme, with a different molecular weight that will matter the moment we convert a recipe's concentration into a receptor's dose.

HOST	CONSTRUCTS	DIFFICULTY	PUNGENT PRINCIPLE	PRECURSOR → LIBERATING ENZYME	MW (G/MOL)
Black Mustard	2,790	Hard	allyl isothiocyanate	sinigrin → myrosinase	99.15
White Mustard	200	Hard	p-hydroxybenzyl isothiocyanate	sinalbin → myrosinase	165.21
Garlic	560	Moderate	allicin	alliin → alliinase	162.27
Ginger	120	Moderate	gingerol → shogaol	(dehydration on drying / heat)	294.39 → 276.37
Watercress	120	Hard	phenethyl isothiocyanate	gluconasturtiin → myrosinase	163.24
Cinnamon	40	Hard	cinnamaldehyde	(free electrophilic aldehyde)	132.16

Black mustard, at **2,790 constructs**, is the deepest pungency library we hold, with white mustard's 200 beside it; garlic's 560 straddle two axes at once, since garlic reads **pungency through TRPA1 and kokumi through CaSR** — the double identity that made it the foundation of almost every savory cuisine. That dual identity is why garlic's 560 constructs must be read on two rulers at once — a single edit can move allicin's sharpness and the calcium-sensing richness together or push them apart — and why we treat it as one of the most delicate hosts in the set despite its merely Moderate difficulty. Difficulty runs Moderate to Hard across the set, because tuning a two-enzyme, injury-triggered system is harder than tuning a single stored compound.



Fresh ginger — gingerol's warmth; dry or heat it and the molecule sheds water to become shogaol, a sharper key.

With the chemistry on the table, the Index entries follow directly from the EC_{50} ratio. The ladder here reads the formula across a range of candidate potencies, with Phoenix's canonical setpoint marked where it falls.

EC_{50}(CANDIDATE) ÷ EC_{50}(REF)	IFI-P = 1,000 ÷ RATIO	READING
0.50	2,000	twice as loud as reference
0.70	1,429	well over the reference
1.00	1,000	the reference anchor
1.25	800	softer, rounder sharpness
2.00	500	half as loud as reference
3.97	252	Phoenix pungency setpoint

The unit conversion is where a kitchen measurement becomes a receptor dose, and it is pure arithmetic once the molecular weight is in hand. Suppose a mustard delivers allyl isothiocyanate to the mucosa at 100 parts per million. For a dilute aqueous film, 100 ppm is 100 mg/L, which is 0.100 g/L; dividing by the molecular weight,

> $0.100\ \text{g/L} \div 99.15\ \text{g/mol} = 1.009 \times 10^{-3}\ \text{mol/L} \approx 1{,}009\ \mu\text{M}$

Against a half-activation threshold that sits in the low-micromolar range, roughly 1,000 µM is some two orders of magnitude over EC_{50} — which is why even a smear of fresh mustard saturates the channel and floods the response to its ceiling. The same 100 ppm of the heavier allicin (162.27 g/mol) is only $0.100 \div 162.27 = 6.16 \times 10^{-4}$ mol/L ≈ 616 µM, fewer moles for the same mass, a reminder that "parts per million" is a mass unit and the receptor counts molecules.

Reaching a precise, pleasurable sharpness — a defined point on the ladder rather than a wall — means delivering the electrophile a controlled, modest multiple over threshold, not the hundredfold overshoot that raw plant tissue throws at the first bite.

Where the bond can fail

A standard this exact earns its keep by being honest about its edges, and pungency has more of them than any other axis.

Not every key leaves a mark. Alongside the covalent electrophiles sit a few reversible agonists, **menthol** the most striking: it binds TRPA1 without reacting and, depending on concentration and species, can either open the channel or block it — a genuinely bimodal key. It is proof that the channel has more than one way in, even though the reactive way is its signature, and a caution that an EC_{50} measured for menthol on one species' TRPA1 may invert on another's.

The channel is an alarm, and alarms have reflexes. Because TRPA1 is the body's chemical-irritant sentinel, its agonists recruit the whole protective repertoire of the airway — the cough, the streaming eyes, the reflex to expel. It is why cutting an onion makes you weep: the onion's lachrymatory factor, assembled by its own injury enzymes, is itself a TRPA1 agonist that finds the trigeminal endings around the eye. A pungency design that ignores volatility and reach can produce a food that is "sharp" on paper and an irritant in the room.

Covalent agonism has a ceiling. Strong electrophiles do not discriminate perfectly; pushed high enough they adduct cysteines on other proteins, deplete the cell's glutathione, and cross from pleasurable sharpness into irritation and tissue stress. There is a band — loud enough to delight, gentle enough to be food — and the whole craft is to stay inside it. Cranking electrophile load is not a strategy; it is how a condiment becomes a hazard.

The second bite is softer. Covalent activation spends the receptor's reactive cysteines, and they are renewed only as fast as the cell's thiol machinery can reduce the adducts and refold the sites. Deliver a second dose before renewal catches up and fewer tripwires remain to trip — the response fades with repetition, a tachyphylaxis a reversible agonist never shows. It is why the third forkful of a sharp mustard reads milder than the first, and why a pungency metered across a whole dish must be designed differently from one built for a single clean hit.

The sharpest molecules are the most fragile. Volatility is what lets an isothiocyanate climb into the nose, and it is also what lets it evaporate out of a jar and cook off in a pan. Tuning for a bright, nose-reaching hit and tuning for shelf stability pull in opposite directions; a mustard built for the sharpest possible top note is, almost by physics, the one with the shortest useful life. And the compartmentalization that makes the chemistry possible can fail in storage: freeze-thaw, bruising, or slow enzymatic leak lets precursor and enzyme meet before the bite, spending the electrophile in the package so the food arrives flat.

The EC_{50} itself is an operational number. Classical EC_{50} assumes a reversible equilibrium, and covalent activation is not strictly reversible on the timescale of a measurement — the adduct forms, the current rises, and resolution lags. We therefore read IFI-P from a half-maximal *peak* response under a controlled exposure time, which makes it a reproducible, physiologically faithful number but one that is defined with respect to timing, not a textbook equilibrium constant. We state that rather than hide it; a number is only respect for the reader if its definition travels with it.

And the mouth integrates what the Index separates. TRPA1 is frequently co-expressed in the very same sensory neurons as the heat channel TRPV1, so pungency and heat can be layered within a single nerve. A chili-garlic crisp speaks fire and sharpness through two channels in one cell, and the mouth reads a single complex burn that is in truth a duet. The Index scores the channels one at a time because that is what you can design against; eating reassembles them, and no profile fully predicts the lived sum.

What we tune, and why

Set against those edges, the design levers are refreshingly concrete. Pungency has three, and we tune all three together: the **stored precursor** (how much glucosinolate, alliin, or gingerol the tissue holds), the **liberating enzyme** (how fast and how completely myrosinase or alliinase converts it at the bite), and the **compartment integrity** that keeps the two apart until the moment of damage and then releases them cleanly. A sharpness is a setpoint across all three, not a single dial, which is exactly why the hosts carry a Moderate-to-Hard difficulty and why the mustard library runs to 2,790 constructs. We model the three as a small coupled system — precursor supply, enzyme turnover, and the fraction of compartments that rupture cleanly at the bite — and solve for the combination that lands the target score at the lowest electrophile load, because the gentlest route to a given loudness is always the safest and the most stable.

The payoff of that discipline is reproducibility, and reproducibility is itself a number. Model IFI-P as roughly proportional to the delivered electrophile in the sensitive band, and the relative error carries straight through: the coefficient of variation on the score tracks the coefficient of variation on the load. Our reproducibility benchmark puts wild-type variation near **42%**; the platform holds batch CV at or below **≤7%**. At Phoenix's predicted setpoint of 252, those two worlds look nothing alike. A 7% CV gives a one-sigma spread of ±18 and a 95% interval of roughly 217–287 — a sharpness you can promise. A 42% CV gives a one-sigma spread of ±106 and a 95% interval near 45–459 — a range so wide the setpoint stops meaning anything, a mustard that is faint one jar and savage the next. Precision is not a luxury layered on top of flavor; for an axis built on a volatile, injury-made, covalently reactive molecule, it is the only thing that makes a *designed* sharpness possible at all.

We hold these plants in particular reverence, and not only for their chemistry. Garlic and ginger and mustard are the aromatics a vegetarian kitchen is built on — the first things into the hot oil, the backbone a meatless pot leans on for depth and bite. There is

something fitting in the fact that their sharpness is a wound chemistry, a defense the plant assembles at the instant it is cut: the pungency we love is the plant at its most alive and most itself, answering injury with intelligence. Our work is to read that answer exactly, to set it where a cook wants it, and to deliver it clean and consistent — so that the sharpness arrives as a pleasure given to the body with intention, and the plant's own brilliance reaches the table undiminished. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 11 · The Salivary PRPs — *the sense that isn't a taste.*

Drink a young, tannic red or oversteep a black tea and your mouth goes suddenly, strangely dry — the surfaces drag, the cheeks pull toward the teeth, the tongue feels stripped of its own slickness. This is astringency, the eighth of our axes and this book's accent, and it is the strangest of the eight because it is not, in the proper sense, a taste at all. No tannin binds a taste receptor and tells the brain *astringent*. Instead the tannin attacks your saliva — and what you feel is the sudden absence of something you never noticed you had. For years our family circled this axis the way you circle a word on the tip of the tongue, because everything we knew how to do assumed a receptor to bind and a dose to raise. Astringency has neither in the ordinary sense, and the first thing it taught us is a thing we have come to love about it: to design astringency, you must first learn to design a loss.



Tea in cross-section — the catechins packed in the leaf are condensed tannins in waiting, molecules that will bind and precipitate the proteins that keep your mouth smooth.

> Every other axis is a molecule the mouth hears. Astringency is a molecule that *takes the mouth's own voice away* — and the silence is the sensation.

The proteins that make the mouth smooth

Saliva is not just water. It is a lubricant, and most of its slickness comes from a family of molecules called **proline-rich proteins (PRPs)** — floppy, open-structured proteins, unusually rich in the amino acid proline, that coat the surfaces of the mouth and let wet tissue glide over wet tissue without catching. To understand

astringency you have to understand first why these proteins are shaped the way they are, because their shape is the whole mechanism.

Most proteins fold. They collapse into a compact globule with a hydrophobic core, and their function lives in a precisely sculpted pocket. The salivary PRPs do the opposite. They are **intrinsically disordered** — they never fold at all, staying extended and loose, a chain that samples a thousand shapes a second and settles into none. The reason is proline. Every other amino acid lets the protein backbone swivel freely at two hinge points; proline's side chain loops back and bonds to its own backbone nitrogen, forming a rigid five-membered ring that locks one of those hinges. String enough prolines together and the backbone can no longer coil into a normal helix or sheet. Instead it is forced into a **polyproline-II helix** — a left-handed, stretched-open spiral with about three residues per turn, flat and splayed, every proline ring presented outward like a run of little paddles. This is not a defect. It is a surface. The polyproline-II helix is a molecular stage built to be bound, and the thing it was built to be bound by is tannin.

Saliva carries several kinds of these proteins, and the division of labor is worth knowing, because it tells you the body made this trade-off on purpose. The **acidic PRPs** also manage calcium and tend the mineral pellicle on tooth enamel. The **basic and glycosylated PRPs** carry the highest tannin affinity of all and appear to have evolved for one job: to meet dietary tannins and be ruined by them first. Alongside them ride the related lubricating proteins **statherin**, which keeps calcium phosphate from crystallizing where it shouldn't and helps the film glide, and the **histatins**, small and antimicrobial. That the basic PRPs appear to be a *sacrificial tannin-sink* — proteins the body manufactures expressly to be destroyed by tannins so that the tannins never reach something more important — is one of the loveliest facts in all of taste. It is a defense you experience as a sensation. The mouth spends a protein to protect itself and charges you, in the same instant, the small toll of dryness.

What a tannin is, exactly

A tannin is a polyphenol — a molecule studded with phenolic hydroxyl groups, the —OH hung off an aromatic ring — and large enough, and multivalent enough, to bridge a protein in more than one place at once. That bulk and that multivalency are what separate a tannin from an ordinary small polyphenol. Two great families do the work, and they are built completely differently.

The **condensed tannins**, or proanthocyanidins, are the ones in grape skin, in tea, in cacao, in the seed coats of sorghum and the skins of berries. They are chains of **flavan-3-ol** units — principally **catechin** and its mirror-image stereoisomer **epicatechin** — linked one to the next through carbon–carbon interflavan bonds. A single catechin or epicatechin unit has the formula $C_{15}H_{14}O_6$ and a mass of 290.27 grams per mole; a dimer forms by losing two hydrogens across the new bond, so a procyanidin dimer runs 578.52 grams per mole, a trimer half again as much, and so on up a ladder we describe by **mean degree of polymerization (mDP)** — the average chain length in a given extract. Two chemical details on these units set their bite.

One is **galloylation**: an ester of gallic acid hung on the ring, as in tea's epigallocatechin gallate, which adds a whole extra phenolic plate and sharply increases protein affinity. The other is chain length itself: a lone monomer barely precipitates protein, while a well-grown oligomer wraps and bridges with ease. Potency climbs with mDP and with galloylation — a fact we lean on constantly when we decide how much grip a construct should carry.

The **hydrolyzable tannins** are built on a sugar instead of a flavan chain. A core of glucose is esterified with gallic acid at several of its hydroxyls; pile on more galloyl groups and you have a **gallotannin**, or let two neighboring galloyls couple oxidatively and you have an **ellagitannin**, carrying the biaryl hexahydroxydiphenoyl group that releases ellagic acid when it breaks. Oak and chestnut give the winemaker these; pomegranate gives us the **punicalagins**, among the largest of them, at 1,084.72 grams per mole — nearly four times the mass of a single epicatechin unit. That size matters in a way the arithmetic makes plain. A gram of punicalagin is 3.7 times *fewer* molecules than a gram of epicatechin, yet each punicalagin molecule carries so many galloyl and HHDP plates that per molecule it bridges far more protein. Mass concentration and molar concentration tell two different stories about the same grip, and we have learned to keep both in view.

A worked unit conversion keeps the two honest. Take a model-saliva assay reading expressed, as they often are, in milligrams per milliliter of catechin-equivalent. One milligram per milliliter is one gram per liter; divided by 290.27 grams per mole, that is **3.45 millimolar** epicatechin-equivalent. The same 0.50 mg/mL of punicalagin, divided by 1,084.72, is only **0.461 millimolar — 461 micromolar** — barely more than a seventh the molar concentration for the same mass on the scale. When we report an astringency load we always say which currency we are counting in.

The binding event, step by step

Put tannin and PRP in the same drop of liquid and three things happen in order, and the order is the mechanism.

First, the tannin coats the protein. The flat aromatic faces of the polyphenol stack against the proline rings of the polyproline-II helix — a hydrophobic association, ring against ring, water squeezed out from between them — while the phenolic hydroxyls reach over and hydrogen-bond to the backbone carbonyls and to the glutamine side chains that pepper these proteins. Neither force alone would hold; together they clamp. This is why proline specifically, and the open helix it builds, is the target: no folded protein presents such a long, bare, bindable face.

Second, the coated proteins cross-link. A tannin is multivalent — many binding plates on one molecule — and a PRP is multivalent too, a long chain with many sites. So one tannin bridges two proteins, another bridges two more, and the solution knits itself into a growing mesh. This is the step that turns binding into something you can feel, because a single tannin on a single protein changes nothing you would notice; a network does.

Third, the network falls out of solution — it precipitates. Past a threshold ratio the aggregates grow heavier than the liquid can hold, nucleate, and drop out as an insoluble complex, taking the lubricating film down with them. In that instant the slick coat on your oral surfaces is chemically stripped away, and your touch receptors — the trigeminal somatosensory endings that report pressure and friction — register the sudden rise in drag between tongue and palate, between cheek and tooth. *That* is the dryness. Astringency is the feeling of your own lubrication being removed, reported not by a taste receptor but by the raw mechanics of unlubricated tissue. There is good evidence that the exposed epithelium and the precipitating particles also tug at mechanosensitive and TRP channels secondarily, which is where astringency gets its faint edge of irritation riding on top of the pure drag — but the drag is the heart of it.



Grape in cross-section — skin and seed are where the condensed tannins live, and the proanthocyanidin spine they carry is the structure a serious red wine is built on.

Where the tannin comes from — the pathway, drawn in enzymes

We cannot tune a sensation we cannot source, so the design work begins in the plant's own chemistry, enzyme by enzyme. The condensed and hydrolyzable tannins take two different roads out of the same metabolism, and we have had to learn to drive both.

The condensed tannins begin on the great trunk every vascular plant shares, the phenylpropanoid pathway: phenylalanine is stripped to cinnamic acid by **PAL**, hydroxylated to *p*-coumaric acid by **C4H**, and activated to *p*-coumaroyl-CoA by **4CL**. Here the flavonoid branch leaves the trunk. **Chalcone synthase (CHS)** condenses that coumaroyl-CoA with three malonyl-CoA units into a chalcone;

chalcone isomerase (CHI) closes it into the flavanone naringenin; **F3H** hydroxylates it to a dihydroflavonol; and **dihydroflavonol-4-reductase (DFR)** reduces that to a **leucoanthocyanidin**, the hinge of the whole story. From the hinge, two enzymes make the two monomers of every condensed tannin. **Leucoanthocyanidin reductase (LAR)** converts the leucoanthocyanidin straight to **catechin**. The other route runs through **anthocyanidin synthase (ANS)**, which oxidizes the leucoanthocyanidin to a coloured anthocyanidin, and then **anthocyanidin reductase (ANR)** reduces that to **epicatechin**. Catechin and epicatechin are then condensed, unit onto unit, into the proanthocyanidin chains whose length — that mDP again — sets how hard they grip. The LAR and ANR branches are our two taps on tannin structure, and the ratio we set between them, together with how far we let the chains polymerize, is most of the astringency dial on a berry or a vine.

The hydrolyzable tannins take the other road, out of the shikimate pathway. Gallic acid is drawn from 3-dehydroshikimate, loaded onto glucose by a glucosyltransferase to make **β-glucogallin**, then galloylated again and again up to **pentagalloylglucose** — the branch point of this family. Add still more galloyl groups and you climb toward the gallotannins; let adjacent galloyls couple oxidatively and you build the ellagitannins, and in pomegranate, the punicalagins. Two completely different syntheses, condensed and hydrolyzable, converging on one sensation at the salivary film — which is exactly why a chapter on astringency has to hold two chemistries at once.



Pomegranate — the jeweled arils are sweet-tart cover for the punicalagins beneath, the hydrolyzable tannins that give the juice its clean, drawing grip.

The hosts we tune for astringency

Six organisms in our curated library carry astringency as a named axis, and the shape of the work lives in how differently they wear it. Grape and tea are multi-axis plants — astringency braided with bitterness, and in tea with umami — while the berries and the sorghum speak astringency almost alone.

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY	TANNIN FAMILY
Grape	1,610	B · A · M	Hard	condensed (proanthocyanidin)
Tea	280	U · B · A	Hard	condensed, galloylated
Blueberry	120	A	Moderate	condensed
Pomegranate	120	A	Hard	hydrolyzable (punicalagin)
Sorghum	120	A	Frontier	condensed
Raspberry	40	A	Moderate	condensed + ellagitannin

Across these six hosts sit **2,290 constructs** — 1,610 of them in grape and wine alone, where the tannin backbone is the whole architecture of the thing, and a bare 40 in raspberry, where astringency is a grace note on a berry built mostly of aroma and sweetness. The counts are a fair map of where the difficulty and the value concentrate: the condensed-tannin plants dominate, the one hydrolyzable-tannin host (pomegranate) is rated Hard for the reach of its punicalagin chemistry, and sorghum sits at the Frontier because coaxing its seed-coat proanthocyanidins into a precise, food-grade register is still the edge of what we can do.

Measuring a sensation that has no receptor

For the seven chemical axes we anchor the Index to a receptor's half-activation — a clean EC_{50}, the concentration that drives the channel to half its maximal response. Astringency denies us that, because there is no receptor binding the tannin and no half-activation to read. So we measure the *event* instead, and we triangulate it three ways.

The first and most direct is the **turbidimetric half-precipitation assay**. Mix a graded series of the polyphenol with a fixed amount of model-saliva protein and watch the solution cloud as the complex drops out, reading the haze by nephelometry or weighing the pelleted protein. The concentration that precipitates **half** of the precipitable protein is our standing EC_{50}-analog — we call it the half-precipitation point, and it plays the exact arithmetic role an EC_{50} plays on the other axes. The second is **tribology**: we drag two soft, wet surfaces across each other, lubricated by saliva with and without tannin, and measure the coefficient of friction directly — an instrument built to feel what the tongue feels, reporting the mechanical truth of the axis as a rising drag. The third is the **time-intensity curve**, where a trained panel rates the sensation second by second, because astringency's signature is as much in its shape over time as in its height. We fuse the three: the half-precipitation point gives us the number, the friction curve confirms the mechanism, and the time-intensity trace keeps us honest about the part a single number can never show.

The number, worked

The Index is one formula, and astringency obeys it like every other axis once we agree what the EC_{50}-analog is:

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

Read it as loudness, not quantity. It does not ask how much tannin a food holds; it asks how *potently* that tannin precipitates protein, relative to a fixed reference, scaled so the reference itself scores exactly 1,000 — substitute the reference's own value top and bottom and the ratio is one, times a thousand. A candidate that reaches half-precipitation at a *lower* concentration than the reference is the more potent, scores *above* 1,000, and the whole axis normalizes for the scorecard radar at a ceiling of 5,000.

Let us turn the handle. Define the relative potency as the IFI divided by 1,000, and the candidate's half-precipitation point as a fraction of the reference's:

IFI-A	POTENCY VS REFERENCE	EC_{50}(CANDIDATE) AS FRACTION OF EC_{50}(REFERENCE)
500	0.5×	2.00
1,000 *(the reference itself)*	1.0×	1.00
1,400 *(Phoenix, A-axis)*	1.4×	0.714
2,000	2.0×	0.500
5,000 *(scorecard ceiling)*	5.0×	0.200

Our model predicts an A-axis score of 1,400 for Phoenix — SS-0096, the single cultivar in our whole library designed to register on all eight axes at once, astringency among them. Run the formula backward: an IFI-A of 1,400 means EC_{50}(candidate) = 1,000 ÷ 1,400 × EC_{50}(reference) = **0.714 × EC_{50}(reference)**. Phoenix's designed tannin load is predicted to reach half-precipitation at 71.4% of the reference concentration — 1.4 times as potent, gram for gram of effect. If we anchor the arithmetic with an illustrative reference half-precipitation point of 1.00 mg/mL catechin-equivalent (3.45 mM), Phoenix's sits at 0.714 mg/mL, which is **2.46 mM** — the same calculation that converts any reading between the chemist's millimoles and the formulator's milligrams. Here is how that single astringency reading sits inside the full chord of the plant we built it into:

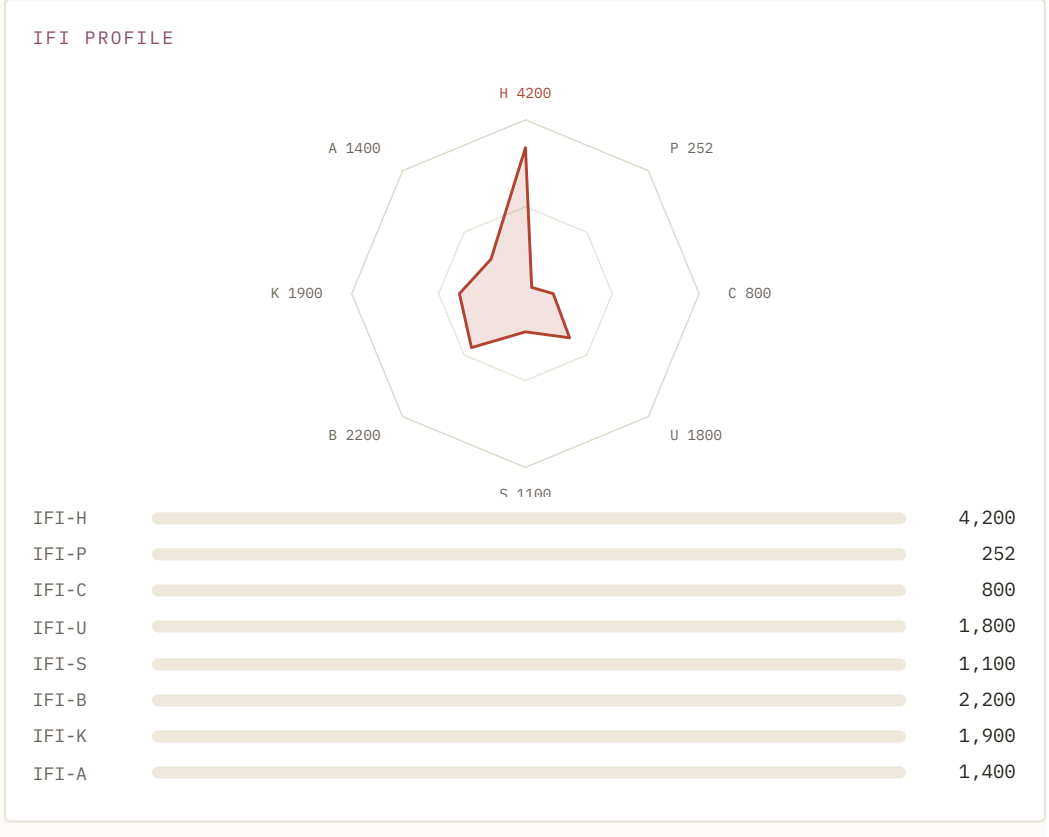

Eight numbers, one plant. The astringency at 1,400 is the quiet voice in a loud room — a third the height of the heat at 4,200 — and that is exactly as it should be. On a cultivar this intense the tannin is not there to grip; it is there to give the sweetness at 1,100 something to lean against, so the whole does not read as flat or cloying. Designing the 1,400 was harder than designing the 4,200, because the target was *restraint*, and restraint has a narrower window than force.

A number without its error bar is a wish, so we carry the error explicitly. Our ≤ 7% batch CV design specification — the same discipline that governs our capsaicin batches and that we built expressly to beat the roughly **42%** swing of wild, unselected material — sets tight bounds on what any predicted score means in practice. At 7%, a predicted A-axis score of 1,400 carries a standard deviation of 0.07 × 1,400 = 98, so a 95% batch window of 1,400 ± 1.96 × 98 ≈ **1,400 ± 192**, from 1,208 to 1,592 — tight enough to design against. At 42% that same reading would scatter across 1,400 ± 1,152, from 248 to 2,552: nearly an order of magnitude of spread, a number you could not formulate to and could not reproduce. The ≤ 7% is not housekeeping. It is the difference between astringency as a specification and astringency as a rumor.

Why it builds instead of fading

Astringency behaves unlike the chemical tastes in one unmistakable way: it **builds** with repeated sips rather than adapting away. A tannic wine's grip tightens across a glass; an oversteeped tea's pucker deepens cup to cup. This is the opposite of heat's merciful desensitization, where TRPV1 quiets down under a steady stimulus — and the reason is written in both the dose-response curve and the kinetics.

The dose-response first. Model the fraction of salivary protein precipitated as a Hill relation, $f(C) = C^n \div (K^n + C^n)$, where K is the half-precipitation point and n the Hill coefficient that measures cooperativity. At C = K, f = 0.5 by definition — half the film gone. Because tannin–PRP binding is multivalent and cooperative, one binding event making the next one easier, n runs greater than one, and the curve steepens. Take an illustrative n = 2: doubling the load to C = 2K drives f from 0.50 to 4 ÷ 5 = **0.80**, while halving it to 0.5K drops f to 0.25 ÷ 1.25 = **0.20**. Against a non-cooperative n = 1 — where 2K gives only 0.67 and 0.5K gives 0.33 — the cooperative curve is far sharper, a near-threshold that tips from "a little drag" to "a lot" across a small change in load. That steepness is why a small overshoot of tannin feels like a cliff, not a slope.

The kinetics do the rest. Each sip removes a bolus of PRP from the film; saliva resupplies PRP only at a steady trickle between sips. Write it as a simple balance — the free-PRP pool after sip $n+1$ equals the pool after sip n, minus what this sip precipitated, plus what the glands replaced in the interval. Whenever removal outruns resupply, the sequence only decreases: less lubricant before each successive sip, more friction reported, the grip climbing monotonically across the glass. Heat fades because the receptor adapts; astringency builds because the lubricant is spent faster than the body can rebuild it. Two axes, two arrows, pointing opposite ways in time.

This is also why astringency, alone among the eight, has a remedy you can add to the cup. Give the tannin a protein to ruin *before* it reaches your saliva and it never strips your film at all. Milk's caseins do this greedily — the reason a brisk black tea takes milk and a tannic chocolate softens against cream — and a bowl of yogurt, a cube of paneer, a soft fresh cheese, or even soy milk will sacrifice their proteins to the same end, which is how the vegetarian table has always tamed its most tannic pleasures. Fat helps too, coating and carrying the tannin past. The remedy is the mechanism run in reverse: spend an outside protein so the mouth need not spend its own.

A few edge cases decide whether a reading holds in the real mouth, and they are failure modes if ignored. Tannin–protein precipitation is **pH-sensitive**, peaking near the protein's isoelectric point; the low pH of wine (around 3.3 to 3.6) and the mouth's buffering shift the curve, so the same tannin load reads more astringent in a more acidic matrix — sourness and astringency are not independent. **Ethanol** solubilizes tannin and suppresses precipitation, so a higher-alcohol wine wears the same tannin more softly, a confound we correct for rather than pretend away. And **polysaccharides** — grape mannoproteins, arabinogalactans — sequester tannin and blunt the grip, one of the real levers behind a wine that tastes "rounder" without less tannin on paper.

The honest limits

We would rather state what this axis cannot do than oversell what it can. Astringency is the one score in the Index anchored to a mechanical, kinetic event rather than a clean receptor, and that buys three honest limitations we carry openly.

It is **protocol-bound**. A half-precipitation point depends on which model protein you precipitate — real basic PRP, a gelatin surrogate, serum albumin — and on the ionic strength, pH, and temperature of the assay. Change the protein and the absolute number moves. So IFI-A is only reproducible across labs when the reference protein and conditions are fixed and published with the value; we treat the method as part of the number, not a footnote to it. It **saturates**. Once the tannin load precipitates essentially all the available PRP, the film is as stripped as it can get — the Hill curve plateaus, and extra tannin past that point no longer adds drag. It instead spills into harshness and into bitterness, which is a separate reader entirely: the smaller flavan-3-ols and their first oligomers activate the bitter TAS2Rs even as the larger, polymerized tannins precipitate the PRPs, so a young red or a square of dark chocolate delivers a real bitterness *and* a real astringency at once, from one family of molecules through two unrelated mechanisms. Teasing them apart — bitter at the receptor, astringent at the protein — is half of reading a tannin correctly, because the two want opposite remedies and respond to opposite moves on the mDP dial: shorter chains lean bitter, longer chains lean astringent. Mistake one for the other and you tune the wrong enzyme.

Balance, not volume

Astringency is where flavor becomes most clearly a matter of balance rather than volume, because a little is structure and a lot is punishment. In wine it is the tannic backbone that makes a red feel serious and ageworthy — and aging is largely this very chemistry run slow, the tannins polymerizing and dropping quietly out of the bottle over years until the grip softens of its own accord; past a threshold, though, that same backbone is merely harsh. In tea it is the briskness that makes a cup feel alive, and oversteeped it is the pucker that sends people reaching for milk. In fine chocolate and in pomegranate it is the grip that keeps sweetness from cloying. So across our astringency hosts the work is almost never about *more*. It is about placing the tannin exactly: enough condensed-tannin structure, routed through the LAR and ANR branches that make the catechin and epicatechin of every proanthocyanidin, to give a wine its spine or a tea its briskness, without tipping into the dry harshness that the salivary proteins, overwhelmed, report as too much. Of the eight readers this is the one that asks us most often to hold back — a fitting accent for a book whose whole argument is that the art of flavor is not loudness at all, but knowing, axis by axis, exactly how loud each reader should be allowed to speak.

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



PART III OF III

Reading the Whole World

convergence · synergy and suppression · the design inversion · one tongue, every plant

Chapter 12 · The Convergence — *where eight streams become one flavor.*

You have never tasted an axis. You have never, in your whole life, experienced a pure TRPV1 read or an isolated CaSR modulation. What you experience is *flavor* — a single, seamless, instantaneous whole — and the eight readers are a fiction of the laboratory, a way of decomposing something the brain hands back as indivisible. This final part of the book is about that reassembly: how eight private molecular languages become one experience, why that experience carries things no taste receptor reads at all, and how understanding the convergence lets us design not eight numbers but one coherent moment at a table. We spend most of our days pulling flavor apart — one receptor, one EC_{50}, one score at a time. This is the chapter where we put it back together, because the putting-back-together is the part a person actually eats.



The 8/8 cultivar in hand — one object that fires on all eight readers at once is also the clearest proof that the eight are, in the end, read as one.

The common currency — one message behind eight doors

The convergence begins long before the brainstem. It begins inside the cell, in the first thousandth of a second of transduction, and the first surprise is how little the eight axes share at the door and how completely they agree once inside.

Three of the eight are *ionotropic*. Heat (TRPV1), cool-burn (TRPM8), and pungency (TRPA1) are cation channels set directly into the bare nerve endings of the trigeminal nerve, and when the right molecule

arrives — a capsaicinoid for TRPV1, a menthol for TRPM8, the allyl isothiocyanate of mustard for TRPA1 — the channel simply opens. Calcium and sodium pour across the membrane, the ending depolarizes on the spot, and a fiber fires. There is no relay and no second messenger; the molecule *is* the switch. This is why heat arrives so fast and leaves so slowly: it is wired to the pain-and-temperature system, not routed through the taste buds at all.

Four others are *metabotropic*. Sweetness (TAS1R2·R3), umami (mGluR1 together with TAS1R1·R3), bitterness (TAS2R38), and the quiet continuity of kokumi (CaSR) are G-protein-coupled receptors sitting on specialized taste cells, and when their ligand binds, nothing crosses the membrane yet. Instead a cascade runs inward: phospholipase C-β2 cleaves its lipid substrate, the fragment releases calcium from the cell's own internal stores, that calcium throws open TRPM5, and the depolarization that follows flings wide the CALHM1 channels through which the taste cell finally pushes its message out — not as a classical vesicle but as a pulse of ATP onto the afferent fiber pressed against it. Three different receptors — sweet, umami, bitter — pour into one shared downstream river. Knock out TRPM5 and all three go dark at once, because the receptors are many but the cascade is one. Kokumi's CaSR mostly rides that same river as a volume knob, lifting the gain on everything around it without a loud voice of its own.

The eighth is stranger than any of them. Astringency has no lock-and-key receptor at all. Its "receptor," in our index, is a *reaction*: salivary proline-rich proteins meet the tannins and polyphenols of tea or grape or fine chocolate and precipitate, stripping the mouth of its lubrication, and the brain reads the loss of slip as a mechanical drag through ordinary touch fibers. It is the one axis that is felt rather than tasted — a change in friction, not a chemical word. We keep it in the eight because it is real, reproducible, and designable; we only stay honest that its machinery is a different kind of thing.

AXIS	RECEPTOR	SIGNAL CLASS	WHAT LEAVES THE CELL
IFI-H · Heat	TRPV1	Ionotropic channel (trigeminal)	direct depolarization → glutamate, CGRP
IFI-C · Cool-Burn	TRPM8	Ionotropic channel (trigeminal)	direct depolarization
IFI-P · Pungency	TRPA1	Ionotropic channel (trigeminal)	direct depolarization
IFI-S · Sweetness	TAS1R2·R3	Class-C GPCR (taste cell)	PLCβ2 → TRPM5 → ATP via CALHM1
IFI-U · Umami	mGluR1 / TAS1R1·R3	Class-C GPCR (taste cell)	PLCβ2 → TRPM5 → ATP via CALHM1
IFI-B · Zero-Bitter	TAS2R38	GPCR (taste cell)	PLCβ2 → TRPM5 → ATP via CALHM1
IFI-K · Kokumi	CaSR	Class-C GPCR (modulatory)	gain on the shared ATP cascade
IFI-A · Astringency	Salivary PRPs	Protein–tannin precipitation	lost lubrication → mechanical drag

Look down the last column. Eight doors, three kinds of hardware, and in the end the same handful of signals leave: a rise in calcium, a swing in voltage, a transmitter spilled onto a nerve. This is the first and deepest meaning of convergence. However a receptor is

engaged, the cell's output is written in one alphabet, and from that point inward the body no longer knows or cares which door was opened. The decomposition into eight is ours. The reassembly into one has, in a sense, already happened by the time the signal leaves the tongue.

Eight numbers, one vector — the arithmetic of a read

If the body collapses eight doors into one message, our job is the reverse: to hold the eight apart precisely enough to design them, and the index is how we do it. Each axis is a single number anchored to half-maximal activation of its receptor,

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

so that a score of 1,000 is, by construction, the loudness of the native reference on that axis, and everything is read against that line. Each reference constant is a stated convention — defined, the way Scoville's dilution-to-threshold was defined, to fix a consistent anchor — not a single cross-axis measurement scale. A score of 1,000 on the heat axis and a score of 1,000 on the pungency axis both mean "as loud as the axis's own reference"; they do not mean the same absolute stimulus intensity, because each axis is anchored to a different molecular reference. A full read is therefore not one number but eight — an ordered vector — and the clearest object to read it on is the one cultivar built to register on every axis at once.

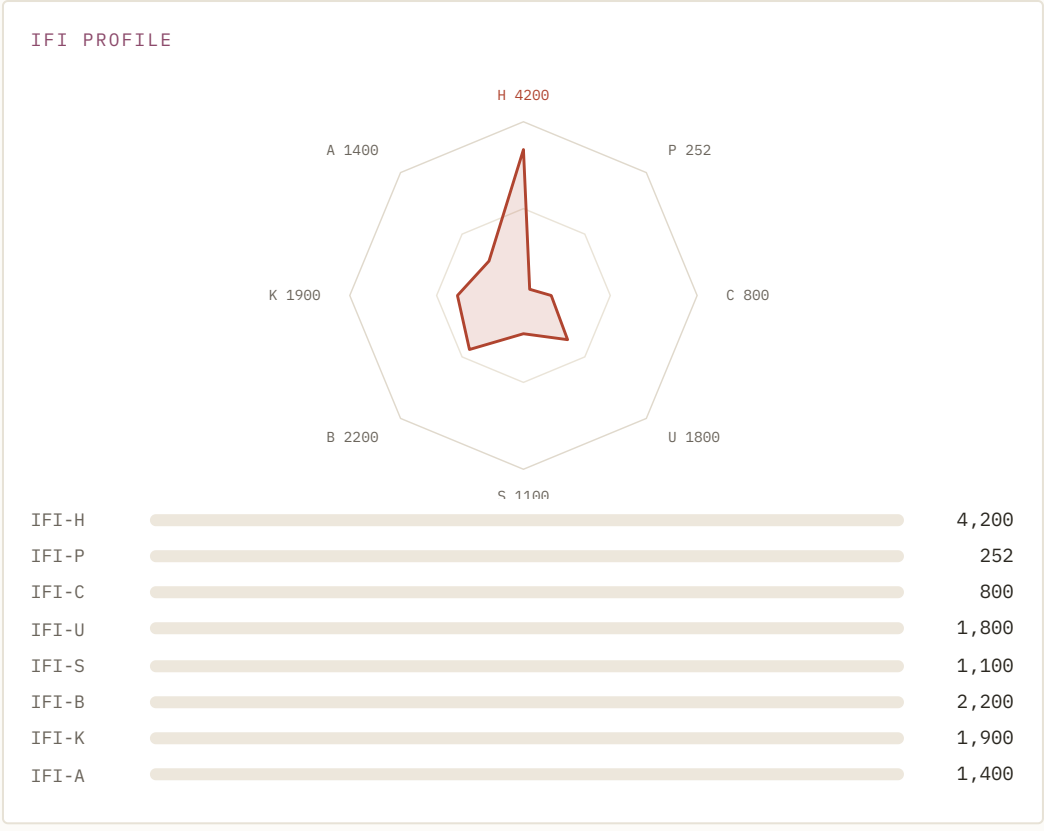

Read the formula backward and each coordinate becomes a concrete statement about a dose. Rearranged, $\mathrm{EC}_{50}(\text{candidate}) = 1{,}000 \times \mathrm{EC}_{50}(\text{reference}) \div \mathrm{IFI}$ — so a score tells you exactly what fraction of the reference concentration reaches the receptor's halfway point. The heat axis at 4,200 means the capsaicinoids reach TRPV1's half-activation at 1,000 ÷ 4,200 = 0.238 of the reference concentration: a little under a quarter of the dose, which is to say 4.2 times louder than the line we call 1,000. The cool axis at 800 and pungency at 252 both read below the reference line — cool at 0.80×

and pungency at 0.25× the reference — the two axes on this profile where the design runs quieter than its own baseline. A plant built to maximize capsaicin-heat concentrates its design energy on TRPV1; the TRPM8 cool channel and the TRPA1 mustard-pungency channel run on different chemesthetic hardware altogether, and they reflect that focus.

AXIS	IFI	LOUDNESS VS. REFERENCE (IFI ÷ 1,000)	DOSE VS. REFERENCE (1,000 ÷ IFI)
Heat · H	4,200	4.20×	0.238
Zero-Bitter · B	2,200	2.20×	0.455
Kokumi · K	1,900	1.90×	0.526
Umami · U	1,800	1.80×	0.556
Astringency · A	1,400	1.40×	0.714
Sweetness · S	1,100	1.10×	0.909
Cool-Burn · C	800	0.80×	1.250
Pungency · P	252	0.25×	3.968

Holding the eight as a vector lets us ask arithmetic questions the single numbers cannot answer. Add the coordinates and the raw sum is 13,652; divide by eight and the arithmetic mean is 1,707 — a fair figure for "how loud, on average, is this plant." But the mean hides the thing that matters most about a full-spectrum read, which is *balance*. For that we take the geometric mean — the eighth root of the eight coordinates multiplied together — which comes to roughly 1,330. The gap between the two, 1,707 against 1,330, is a direct measure of unevenness; their ratio, about 0.78, is an evenness index where 1.00 would mean a perfectly flat profile with every axis equal. A score of 0.78 says this plant is loud across the board but tilted — heat and bitter-attenuation ride high, cool and mustard-pungency sit low — and that tilt is exactly what a geometric mean is built to see.

The choice of geometric over arithmetic is not cosmetic; it is the whole reason "full spectrum" cannot be counterfeited. An arithmetic mean can be bought with one enormous axis: pile all your loudness into heat and let the other seven sit at the floor, and the average still looks respectable. A geometric mean cannot be bought that way. Because it multiplies, a single coordinate near the baseline drags the whole product down toward the baseline, and seven coordinates near the floor pull it to the floor no matter how loud the eighth. A cultivar that fires on two of eight axes has six coordinates resting near the reference line; multiply them in and its breadth score collapses, while its loudest axis barely moves the needle. Breadth, in this index, is defended by the multiplication itself. You earn a high full-spectrum read only by speaking on every axis — which is why the plant that does, the 8/8, is the one we photograph.

On a scorecard the same vector is drawn as a radar, each spoke normalized to a maximum of 5,000. Heat at 4,200 reaches 84% of the way out its spoke; pungency at 252 reaches 5% of its own. The shape you see — the lopsided eight-pointed star — is the vector made visible, and a trained eye reads the tilt off the silhouette before it reads a single number.

The dose that means half — why the index is anchored at the receptor

Everything above rests on one word in the formula: EC_{50}, the concentration at which a stimulus drives its receptor to half of full response. We could have anchored the index anywhere on the dose–response curve. We anchored it at half for a reason that is as much statistics as biology.

A receptor's response to rising concentration is not a straight line but a sigmoid, well described by the Hill equation,

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

where [A] is the stimulus concentration, R_max the saturating response, and n the Hill coefficient, the steepness term. Set [A] equal to EC_{50} and the equation collapses to R = R_max ÷ 2 — the halfway point, by definition. TRPV1 is a tetramer, four subunits closing cooperatively around their ligand, so its heat response carries a Hill coefficient near 2: the curve is steep, and small changes in capsaicinoid concentration swing the response hard. That steepness is precisely why half is the right anchor. Plotted against the logarithm of concentration, the sigmoid is flat and uninformative down at threshold, flat and saturated up at the top, and steepest exactly at its midpoint — at EC_{50}. The half-maximal point is where the receptor discriminates best, where a given error in concentration produces the smallest error in the number you read back. Anchoring the index at the floor would measure noise; anchoring it at saturation would measure nothing, because everything saturates alike. Anchoring at half measures the plant where the receptor itself is listening hardest.

The payoff is a scale with enormous reach. A sunlit shishito scores near 40 on the heat axis; by the formula its capsaicinoids need 1,000 ÷ 40 = 25 times the reference concentration to reach TRPV1's halfway point. The 8/8 cultivar scores 4,200 and needs 0.238 times that concentration. From shishito to full fire the axis spans 4,200 ÷ 40 = 105-fold in loudness — and 25 ÷ 0.238 ≈ 105-fold the other way in dose, the two reciprocal because loudness is simply one over the dose that means half. One axis, two orders of magnitude, every pepper in the world placed on it by a single measurement.

It helps to set that against the scale the world already knew. The Scoville scale was born as an act of tasting: dilute an extract in sugar water, step the dilution down, and record the factor at which a trained panel can no longer feel the burn — that factor is the Scoville Heat Unit. It is a *chemical* measure, a proxy for how much capsaicinoid a fruit holds, today read off an HPLC trace rather than a panel's tongues. Pure crystalline capsaicin sits near 16,000,000 SHU, so when our model predicts approximately 2.94 million SHU for our 8/8 cultivar, that places it at roughly 18% of the pungency of the isolated molecule (2,940,000 ÷ 16,000,000 = 0.184) — nearly a fifth of the way to pure capsaicin, designed into a fruit you can hold. The index measures something else entirely: not how much the plant holds but how loudly it speaks to the receptor. The two track, and they do not lock step.

CULTIVAR	SKU	PREDICTED SHU	IFI-H	AXES	IFI-H PER 1,000 SHU
Ember	SS-0087	866,000	1,570	2/8	1.81
Blaze	SS-0091	1,720,000	2,630	4/8	1.53
Firestorm	SS-0094	2,470,000	3,490	6/8	1.41
Phoenix	SS-0096	2,937,948	4,200	8/8	1.43

Read the last column down the ladder and a quiet story appears. At the bottom, every thousand Scoville units our model predicts buys 1.81 points of receptor loudness; near the top, only about 1.43. The loudness-per-mass *falls* as the heat climbs, then flattens. Mass and loudness are measuring different things, and here is the gap between them made numeric: capsaicin is the most potent capsaicinoid, and as the pathway is pushed toward its ceiling, more of the added heat accrues as the gentler congeners — dihydrocapsaicin and its minor relatives — so each extra unit of mass speaks a little more softly to TRPV1 than the last. It is a diminishing return written into the plant's own chemistry. And our designed 8/8 cultivar, at 1.43, sits right on the curve the ladder traces, a hair above Firestorm — our model's indication that the designed edits aimed to lift total output without warping the potency of the blend. Two instruments, the chemist's and the receptor's, reading one pepper and gently, honestly disagreeing.

The road inward — eight streams, one cortex

The signals that leave the tongue do not stay eight for long. The anterior two-thirds of the tongue sends its taste up the chorda tympani, a branch of the facial nerve; the back third rides the glossopharyngeal; the throat and epiglottis add the vagus; and running alongside all of them, the trigeminal carries the heat, the cool, the pungency, and the mechanical drag of astringency — the chemesthetic half of flavor, the part that is felt. These separate cables all terminate in one place: the nucleus of the solitary tract, low in the medulla, the first station where the streams begin to pool.

From there the primate brain takes a more direct road than the rodent one that taught us most of what we know. In us the solitary nucleus projects largely straight to the thalamus — to the parvocellular division of the ventral posteromedial nucleus — and up to the primary gustatory cortex, folded into the anterior insula and the frontal operculum deep in the side of the head, where the streams are read together and handed on to the orbitofrontal cortex. There, and only there, taste finally meets smell, sight, texture, and memory, and becomes the thing we call liking.

This architecture settles an old argument in the field. The *labeled-line* view holds that each quality travels its own dedicated wire — a sweet line to a sweet cell to a sweet region — while the *combinatorial* view holds that quality is read from the pattern across many broadly tuned fibers at once. Decades of recording resolve it toward a middle, and that middle is useful to us. The basic tastes do run on reasonably dedicated lines, which is why they feel categorical and distinct. But intensity, blend, and the fine distinctions — where one cultivar's warmth differs from another's, where bitterness shades into depth — ride on the *pattern* across the population and on its *timing*, the order and rhythm in which fibers fire. For a flavor

designer this is the liberating fact of the whole book: the brain is built to integrate. A well-composed set of eight reads will be experienced as one thing, not eight things fighting for attention. We do not have to make the axes agree with one another. The brain does that for us. We only have to place them well.

The arithmetic of a clean read — convergence as noise reduction

There is a second, less obvious reason the brain pools its streams, and it is the same reason our index pools its measurements: averaging kills noise. Convergence is not only how eight become one — it is how eight become *reliable*.

Start with a single score. An IFI value is a ratio of two half-activations, the candidate's against the reference's, and ratios carry the error of both terms. Relative variances add in quadrature:

> $(\sigma_{IFI} \div IFI)^2 = (\sigma_{ref} \div EC_{50,ref})^2 + (\sigma_c \div EC_{50,c})^2$

Specify each EC_{50} at our batch coefficient of variation — the ≤7% we hold our constructs to, against the 42% a wild-type line will drift across seasons — and a single score inherits $\sqrt{(0.07^2 + 0.07^2)} \approx 9.9\%$ of scatter, the reference's own wobble folded into every number that leans on it. The remedy is to pin the reference down with many runs until its CV collapses toward zero; then the quadrature term vanishes and the index CV falls back to the construct's own ≈7%. It is the sort of bookkeeping no one tastes and everyone depends on: a shared anchor is only as steady as the number of times you have measured it.

Now take the whole vector. A full-spectrum read is, in part, an average across eight axes, and averaging independent measurements shrinks their scatter by the square root of their count. Eight axes each at 7%, if independent, give a composite CV of $7\% \div \sqrt{8} \approx 2.5\%$. This is precisely why the nervous system reads flavor off a *population* of fibers rather than trusting any one: a thousand noisy reporters, averaged, tell a cleaner story than the best single one. But the honest caveat is the interesting part. The eight are not fully independent — a single plant, a single harvest, a single warm week pushes them together — and when a shared factor correlates the axes at some coefficient ρ, averaging can no longer drive the composite below roughly $7\% \times \sqrt{\rho}$. At $\rho = 0.25$ that floor is 3.5%, not 2.5%; the correlated part of the noise does not average away, ever. The brain lives with exactly this limit — correlated firing among neurons sets the floor on how much a population can clean up — and so do we. It is a quiet piece of respect we owe the reader: to know not only how precise a score is, but why it cannot be more precise than it is.

The four hundred readers the tongue never had — retronasal flavor

Everything so far is *taste*, and taste is the smaller half of what you call flavor. Most of what you mean by the flavor of a food is not taste at all. It is *smell*, delivered the back way — retronasally — rising from warm food in the mouth up behind the palate to the olfactory epithelium at the top of the nose, where some four hundred kinds of olfactory receptor read the volatile molecules the eight taste-and-trigeminal readers never touch.



A strawberry in cross-section — pinch your nose through one and almost all of *strawberry* vanishes, because strawberry-ness is a smell the tongue was never built to read.

The olfactory system is combinatorial in a way the tongue is not. Each sensory neuron in the nose expresses a single kind of receptor; each receptor answers to many molecules, and each molecule lights a particular subset of receptors, so that an aroma is spelled not by one key in one lock but by a chord struck across hundreds. Four hundred receptors, read in combination, resolve a number of distinct smells that runs well into the thousands and, by some estimates, far beyond — a vocabulary so large it makes the eight tastes look like an alphabet beside a language. And the brain refers that whole language to the mouth. It binds the retronasal smell so tightly to the taste arriving at the same instant, and localizes it so convincingly on the tongue, that we feel the aroma *in the mouth* and call the entire experience taste.

Pinch your nose and eat a strawberry and you get sweet, a trace of acid, a faint astringent grip — and almost nothing of *strawberry*, because strawberry-ness is an olfactory read riding on top of the taste. Release your nose and the fruit floods back in a single instant, the smell snapping into the mouth where it never actually was. The binding is seamless, and it is the clearest demonstration we know of how much of flavor the tongue does not read.

> You do not taste sweetness and smell strawberry and feel cool and register a little drag. *You taste a strawberry.* The binding is so complete it feels like a single sense.

This is why the index measures the eight and does not pretend to measure flavor whole. The eight axes are the part we can both compute and design — the receptors we know by name and can target molecule by molecule, the bones of the thing. The retronasal remainder is real, enormous, and for now largely beyond our direct design; we work with it, choosing cultivars whose aromatics are right and whose volatiles carry, but we do not yet *write* an aroma the way we write a capsaicinoid blend.

The binding — temperature, texture, color, sound

Smell is the largest of flavor's silent partners, but it is not the only one, and the brain folds them all into the single object so smoothly that we never feel the seams.

Temperature is a partner twice over. It is a stimulus in its own right — TRPV1 opens to the heat of the food as readily as to a capsaicinoid, which is why a warm sauce reads hotter than a cold one of identical build — and it is a *modulator* of taste itself: sweetness perception climbs toward body temperature and falls when a food is chilled, so the same sugar reads sweeter warm than cold. Texture is a partner through the trigeminal touch system that also carries astringency; thickness, oiliness, and crispness all tune how long a flavor lingers and how intense it feels. Even the senses that seem to have nothing to do with the mouth lean in. Color shifts flavor — a drink dyed red reads sweeter and more of cherry, and tasters confidently find berry notes in a white wine quietly tinted dark. Sound shifts it too — a crisp, high register sharpens the sense of freshness and crunch, the effect sometimes called sonic seasoning.

The most useful of these bindings for a designer is *congruence*. When an aroma and a taste belong together by long habit — a strawberry's smell with sweetness, a broth's aroma with umami — the brain does not add them; it multiplies them. A congruent aroma makes a sweetness read sweeter than the sugar alone can account for, a super-additivity measurable at the tongue. This is leverage we design around directly: the right volatile over the right taste buys perceived intensity that costs nothing in sugar or salt, flavor the body is given rather than loaded with. The seams never show because the brain is built to erase them — to take smell and taste and warmth and drag and the color of the thing and present them as one undivided experience. Our discipline is to respect that the experience is one, even while we build it from parts.

Liking is computed last — the learned palate

There is one layer the eight readers do not contain, and it is the one that finally decides a meal: *liking*. Pleasure is not computed in the primary taste cortex at all. It is computed downstream, in the orbitofrontal cortex, where taste and smell and sight and memory and the body's hunger are weighed together — and, crucially, where the verdict is *learned*.

We are born loving sweet and recoiling from bitter, the two innate biases of a body guarding its energy and its safety; a newborn's face sweetens to sugar and curdles to quinine before a single meal has taught it anything. And yet across a life we learn to love the bitterness of coffee, the burn of chili, the grip of a tannic wine — teaching ourselves to savor the very alarms we began by fearing. The psychologist's phrase for the chili case is *benign masochism*: the particular delight of a danger the body reports and the mind knows to be safe, the fire that is not there enjoyed precisely because it is not there. The same orbitofrontal cortex adds another twist the moment you begin to eat: *sensory-specific satiety*, the fading of a food's pleasantness as you take it in while the uneaten foods beside it stay bright — the reason the last bite of one dish never satisfies like the first bite of the next.

This plasticity sets the honest limit of our work. We design the read, placing eight molecules against eight locks with real precision; the brain decides the liking, out of its own history. We can build a flavor exactly. We cannot build the wanting of it — only offer it, cleanly made and honestly scored, and let a person's own palate come to meet it. That boundary does not trouble us. It is the right one. A plant grown to carry the most it can, a heat offered as a pleasure rather than an assault, a flavor built to nourish and not merely to impress — these we can put on the table with intention. The savoring is the eater's, and it always should be.

The bones of flavor — what the Index writes, and what it leaves to you

So the eight axes are the skeleton, and flavor is the whole living body. This is not a confession of limits so much as a correct account of what a skeleton is for. A skeleton is what everything else hangs on — and a flavor designed at the receptor is a flavor built from the bones out.



Tomato in cross-section — restore the receptor reads the modern fruit lost, and an ordinary tomato's aroma suddenly has a structure to stand on.

Consider a tomato. Its aroma — the green, warm, unmistakable smell of a tomato — is a retronasal read, carried by volatiles we did not touch, and in most modern fruit that aroma is still present, still correct, still waiting. What the long decades of shipping and shelf-life took was not the smell but the *structure underneath it*: the umami and the sweetness, the receptor reads that give the aroma something to stand on. Restore those axes — rebuild the savory depth and the sweetness through the plant's own native routes — and the fruit does not gain a new smell. It gains a frame, and the aroma that was always there suddenly reads like so much more, because for the first time in a generation it has bones beneath it. The eater tastes a tomato that tastes the way a tomato should, and never once suspects that what was fixed was the half they could not name.

That is the whole argument of the convergence. The brain will bind whatever we give it into one seamless thing; it will refer the smell to the mouth, fold in the warmth and the drag and the color, and hand back a single, indivisible experience with no seams to see. What we can do — the part that is ours to write, molecule against lock, indexed where the receptor listens hardest and held to within a few percent batch to batch — is make the skeleton true. Build the eight reads well and placed with care, and when the aromatics and the temperature and the texture and a person's own history settle down onto that frame, the whole thing stands up straight. We write the bones. The body, and the loving of it, assemble themselves — which is exactly as it should be, at a table, with the people we grew it for.

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

Chapter 13 · When Receptors Talk — *synergy, suppression, and the multi-axis read.*

The eight readers are not strangers to one another. They share molecules at the wet surface of the tongue, they share neurons in the nerves that carry their news, they share the integrating cortex that gathers their separate reports into the single thing we call flavor, and they lean on one another in ways that make the whole far larger than any sum. A flavor, read closely, is never eight independent measurements taken side by side. It is a conversation — and the interactions between the axes are not noise to be filtered away but the very material a family designs in. The most beautiful things in our library are not loud on any one axis. They are *composed* across several, built precisely out of the ways the receptors talk to one another.



A cultivar in flower — the plant composes its own molecules across a dozen branching pathways; we compose the reads those molecules make at the tongue.

There are, in the end, three grammars to the conversation. In **synergy**, one axis lifts another at the receptor itself — a molecule that is almost silent alone makes a second molecule speak louder than its concentration should allow. In **suppression**, one axis quiets another, both at the periphery where molecules compete and in the cortex where a loud read masks a soft one. And in **tension**, two or more axes are held against each other until they resolve into a third sensation that is none of them — the domain where flavor design

becomes most plainly an art built on exact biology. This chapter is those three grammars, and the arithmetic that lets us write in them and stand behind what we wrote.

Synergy — one axis lifting another at the lock

The cleanest synergy in the whole index is umami's, and it is worth following molecule by molecule, because it is the place where "receptors talk" stops being a metaphor and becomes a measurable shift in a binding curve. The umami receptor that matters here is the heterodimer **TAS1R1·TAS1R3** — the glutamate-reading mGluR1 runs the same logic in a simpler key. Each subunit carries, on the outside of the cell, a large bilobed domain hinged like a Venus flytrap: two lobes that close on a hinge when the right molecule settles into the cleft between them. Glutamate binds in the cleft of the TAS1R1 flytrap and tips it from the open, resting conformation toward the closed, active one. That closure *is* the signal. It torques the dimer; the strain travels down a cysteine-rich stalk to the seven-helix bundle buried in the membrane; and on the inner face a waiting G-protein — gustducin, of the Gαi family — is released to drive phospholipase C-β2 → IP_3 → a pulse of calcium → the opening of the TRPM5 channel → depolarization → ATP vented through the CALHM1 pore to the afferent nerve. One molecule in a cleft, a dozen stations later, becomes a brothy depth reported to the brain.

Now add the ribonucleotide. The 5′-ribonucleotides — guanylate (GMP), which dried shiitake and porcini carry in abundance, and its cousin inosinate (IMP) — taste of almost nothing on their own. But a ribonucleotide binds at a *second* site, just outside the glutamate cleft near the hinge, where it acts as a molecular doorstop: once glutamate has pulled the flytrap partway shut, the ribonucleotide wedges against the outer face of the closed lobes, lowers the energy of the active state, and makes the closure both easier to reach and slower to release. It adds no glutamate. It changes what a given amount of glutamate is *worth*.



A tomato in cross-section — the locular jelly where glutamate and ribonucleotide collect; savory depth is these two molecules reading louder together than either reads alone.

In the language of the Index, that change has an exact shape. Loudness on any axis is anchored to the **EC_{50}** — the concentration that drives a receptor to half of its full response — through

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

so that the native wild-type reference reads exactly 1,000 and a candidate that needs *less* stimulus to say the same thing scores *higher*. A positive allosteric partner like the ribonucleotide does its work entirely inside the EC_{50}: it does not raise the ceiling of the response, it slides the whole dose–response curve to the left, so the receptor reaches half-activation at a lower glutamate concentration. Define the **synergy factor** as the ratio of the two half-activation concentrations,

> $S = EC_{50}(\text{glutamate alone}) \div EC_{50}(\text{glutamate + ribonucleotide}) = IFI(\text{with}) \div IFI(\text{without}),$

and the second equality falls straight out of the formula — because the reference EC_{50} is common to both scores, the ratio of the scores *is* the ratio of the inverse EC_{50}s. Synergy, read this way, comes off the Index directly, as a multiplier.

Put real numbers to it. Anchor the umami reference where the biochemistry sits — glutamate half-activates the TAS1R1·TAS1R3 heterodimer near **1.2 mM** — so EC_{50}(ref) = 1.2 mM carries IFI-U 1,000. Our flagship **Phoenix** is designed to an **IFI-U 1,800**, which the formula inverts to an effective half-activation of

> $EC_{50} = 1.2\ \text{mM} \times 1{,}000 \div 1{,}800 = \mathbf{0.667\ mM},$

a savory depth that our model predicts reaching half-response at 0.56× the glutamate the wild type would need. In weighable units — glutamate's molar mass is 147.1 g/mol, and in dilute water a milligram per litre is a part per million — that is **177 mg/L (177 ppm)** at the reference and **98 mg/L (98 ppm)** in Phoenix: the same half-signal, at a little over half the glutamate. The restoration is not one trick but several braided together — more glutamate laid down, the ribonucleotide synergy, a kokumi floor beneath it — yet the ribonucleotide alone, sliding EC_{50} from 1.2 to 0.667 mM, accounts for a synergy factor of **S = 1.8**.

What that multiplier does to an actual mouthful is best seen on the dose–response curve, which is not a straight line but an S. The fraction of receptor driven, θ, follows a **Hill relation**, $\theta = [A]^n \div (EC_{50}^n + [A]^n)$; for a cooperative class C dimer the Hill coefficient n runs near 1.5. Hold the glutamate concentration fixed at the reference value where, alone, it gives exactly half-activation — θ = 0.5 by definition — and switch the ribonucleotide on. The effective EC_{50} drops by the factor S, the ratio $[A]/EC_{50}$ rises to 1.8, and the occupancy climbs to

> $\theta = 1.8^{1.5} \div (1 + 1.8^{1.5}) = 2.42 \div 3.42 = \mathbf{0.707.}$

The receptor goes from half-lit to seven-tenths lit with not one additional molecule of glutamate added to the dish. That is the whole gift of synergy, and why we prize it above brute concentration: it buys loudness without volume. The signal rises while the quantity of the costly molecule — and the carbon and energy the plant spent making it — stays flat.

UMAMI CONDITION	EC_{50} (MM GLUTAMATE)	EFFECTIVE IFI-U	0 AT THE REFERENCE DOSE
glutamate alone (reference)	1.20	1,000	0.500
+ ribonucleotide (realized, Phoenix)	0.667	1,800	0.707
+ saturating ribonucleotide (in-vitro ceiling)	~0.15	very high	~0.95

The last row is the honest edge of the gift. In a cuvette, with ribonucleotide pushed to saturation, umami synergy can reach significant amplification — EC_{50} collapsing toward 0.15 mM, a score that would run off the top of a scorecard normalized to 5,000. A fruit cannot go there. Saturating ribonucleotide is metabolically expensive, and a plant has only so much carbon to spend; the 1.8 we actually build is the synergy the plant can *afford* to sustain pod after pod, not the synergy the receptor would permit in glass. The design lives at that boundary, and naming it honestly is part of the measurement.

Synergy has a second, subtler voice, and it belongs to the quietest axis. **Kokumi** is read by the calcium-sensing receptor, **CaSR** — another class C GPCR with its own flytrap domain, tuned natively to extracellular calcium but promiscuously roused by polyamines, bare amino acids, and above all the **γ-glutamyl peptides**: glutathione, γ-Glu-Val-Gly, the γ-Glu-Cys-Gly that garlic, yeast, and long-simmered beans carry. A γ-glutamyl peptide has no taste of its own; you cannot point to a kokumi the way you can point to a sweetness. What it does is raise the gain and lengthen the time-course of the *other* reads — sweet, salty, umami all made rounder, more continuous, slower to let go — so that a broth gains body and a sauce gains the quality we can only call *continuity*, the sense of a flavor that arrives whole and recedes without a seam. In the index this is IFI-K, and it is synergy of the broadest kind: not one molecule potentiating one receptor, but one receptor modulating the whole chord.

Suppression — one axis quieting another

Where synergy adds, suppression subtracts, and it runs along two separate roads. The first is **peripheral**: at the tongue, molecules and receptors compete directly — for binding, for the shared downstream machinery of a taste cell, for the finite transmitter the cell can vent. The second is **central**: in the nucleus of the solitary tract, where the taste nerves first synapse, and in the gustatory cortex of the insula, where the final read is assembled, a loud input masks a soft one much as a loud tone masks a quiet one in the ear. The two roads matter to a designer differently, because the first can be engineered at the plant and the second cannot — you can lower the molecule that competes, but you cannot talk the cortex out of hearing the loud thing first.

The cleanest peripheral case is written into the architecture of the receptors themselves. The sweet receptor is the dimer **TAS1R2·TAS1R3**; the umami receptor is **TAS1R1·TAS1R3**. They are built on a **shared subunit** — the same TAS1R3 partners with TAS1R2 to read sweet and with TAS1R1 to read savory — and a shared part means the two reads are not wholly independent. Model it to first order: let a cell hold a fixed pool T of TAS1R3, let a fraction f of it pair

into umami dimers and the remaining (1 − *f*) into sweet dimers, with the other subunits in excess. The maximal response each axis can mount scales with its dimer count, so

```
Sweet ceiling ∝ (1 − f) · T and Umami ceiling ∝ f · T,
```

and the two ceilings are tied by a conservation line: lift one and you lower the other, because they draw on one pool. Phoenix's savory-leaning balance — **IFI-U 1,800 against IFI-S 1,100** — is exactly this partition made visible, more of the shared subunit committed to the umami read than to the sweet. It is a genuine design constraint, and the honest caveat sits right beside it: sweet and umami cells are largely distinct populations in the real tongue, so the pool is shared only partly, and the partition line is a useful first approximation rather than a hard law. We design against it as a tendency, treating it as a computed guide to the balance our model predicts.

The central suppressions are older than any receptor map — every cook has used them. Salt quiets bitterness: a pinch over grapefruit, a little in the water for bitter greens, the sodium ion blunting the bitter read both at the periphery and in the central mixing. Sweetness and bitterness suppress each other mutually, which is why a bitter shoulder on an intense sweetener must be silenced for the sweetness to read clean, and why so much of our sweetness work and our zero-bitter work turn out to be the same project pulling from opposite ends. Put the central masking in a form you can compute. Let the perceived intensity of a target axis be its raw receptor drive reduced by a share of the masker's drive,

```
P(target) = R(target) − k · R(masker),
```

a subtractive first-order model with masking strength *k*. Take a stevioside sweetener whose raw sweet drive would read 1,400 but which carries a bitter shoulder reading 900, with an illustrative k = 0.33. Perceived sweetness is 1,400 − 0.33 × 900 = 1,400 − 297 ≈ **1,100** — a clean sweetness held down to middling by its own bitterness. Now attenuate the bitter. Our zero-bitter axis silences **TAS2R38** by a canonical **20–65%**; take a 55% cut, and the bitter drive falls to 405. The perceived sweetness rises to 1,400 − 0.33 × 405 = 1,400 − 134 ≈ **1,266** — the same sweetener, about fifteen percent louder, with not a grain more of the sweet molecule. To design a clean sweetness is frequently to design a *suppression*: not adding sweet, but subtracting the bitter that was hiding it.

The model is a straight line laid over something that curves, and we say so. Real masking is nonlinear and saturates; part of it lives at the periphery where our subtractive central form is only an approximation; and the bitter receptor is **polymorphic** — TAS2R38 comes in a taster and a near-non-taster form, so the very bitterness we are attenuating is punishing to one eater and nearly absent in another. We attenuate against the taster, because that is the mouth that needs it, and we call the number a reference rather than a promise identical in every head.

Tension — holding axes against each other

And some conversations are neither lift nor quieting but **tension** — two or more axes held in deliberate opposition until the mouth resolves them into a single thing that is none of them. This is the native domain of the multi-axis read, and the plainest way to hear it is to read a cultivar that speaks on every axis at once.



Phoenix (SS-0096) in the controlled chamber — the only cultivar in our library that registers on all eight axes together, each read tuned to hold the others in balance.

IFI PROFILE

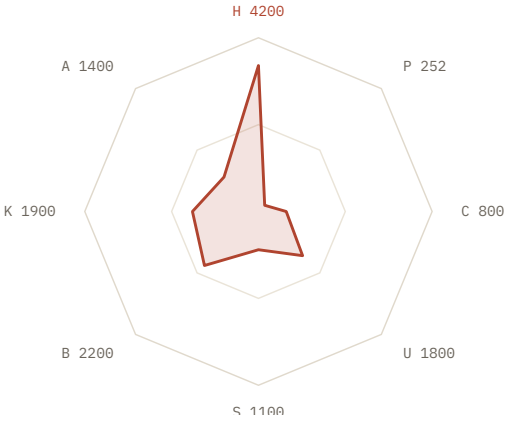

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

Read the scorecard as a chord and not a list. A towering TRPV1 heat at 4,200 is the root; a computed TRPA1 pungency at 252 sharpens its edge; a restored umami at 1,800 and a kokumi at 1,900 lay a savory floor beneath the fire so it reads as depth rather than mere assault; a restrained sweetness at 1,100 and a bitterness-attenuation read at 2,200 keep it from either cloying or turning harsh; a whisper of TRPM8 cool at 800 and a plum-dark astringency at 1,400 frame the whole. Not one axis stands at the scorecard's ceiling of 5,000. The achievement is not any single peak; it is the *interval* between the

reads — eight numbers chosen so the mouth meets not eight sensations but one coherent, astonishing event. To design the chord is to work in the spaces between the axes: lift the heat and the savory floor must rise with it or the fire reads as raw; sharpen the pungency and the cool must come up a fraction to keep the front from turning harsh. Every number is placed with respect to the other seven.

Across the platform these compositions are catalogued as the **multi-axis libraries**, and each is a named conversation between specific receptors:

LIBRARY	AXES	RECEPTORS IN CONVERSATION	THE DESIGN LOGIC
Swicy	H × S	TRPV1 ↔ TAS1R2·R3	burn set against sweetness — the dare and the invitation in one bite, each making the other legible
Savory Depth	U × K	TAS1R1·R3 + ribonucleotide ↔ CaSR	glutamate–ribonucleotide synergy floored by kokumi continuity; our most-requested chord
Bitter-Sweet	B × S	TAS2R38 ↔ TAS1R2·R3	suppress the bitter, lift the sweet — coffee and cacao made to read sweet without sugar
Mála	H × C × P	TRPV1 + TRPM8 + TRPA1	three irritant channels at once — the numbing-heat tingle of Sichuan cooking

The dual-axis **M2** set runs to roughly **8,400** constructs and the triple-axis **M3** set to some **2,800** — a plain measure of how much of the interesting design space lies not on any single axis but in the conversations between them. Most of a flavor's room to be *new* is in those intervals, not in any one height.

The *mála* of Sichuan cooking is the connoisseur's example, and the richest piece of receptor biology on the list, because all three of its channels are trigeminal — the touch-and-temperature senses that pose as tastes — rather than gustatory. The chili's capsaicin opens **TRPV1** for the heat; the aromatics' cinnamaldehyde and allyl compounds open **TRPA1** for the sharpness; and the Sichuan peppercorn adds hydroxy-α-sanshool, which does something stranger than either — it nudges TRPV1 and TRPA1 while also closing a family of background potassium channels (the two-pore KCNK channels) in light-touch fibers, leaving those fibers hyperexcitable and firing at a flutter near fifty beats a second. The brain reads that flutter as a buzzing, numbing tingle that is neither heat nor cold nor pain. Three alarm channels, played together into a sensation that is somehow not alarming at all but *craved* — the nervous system's warning apparatus, composed with care, becoming one of the deep pleasures of the table.

Set that chord against what a novice builds — a single axis driven to its ceiling and nothing beneath it:

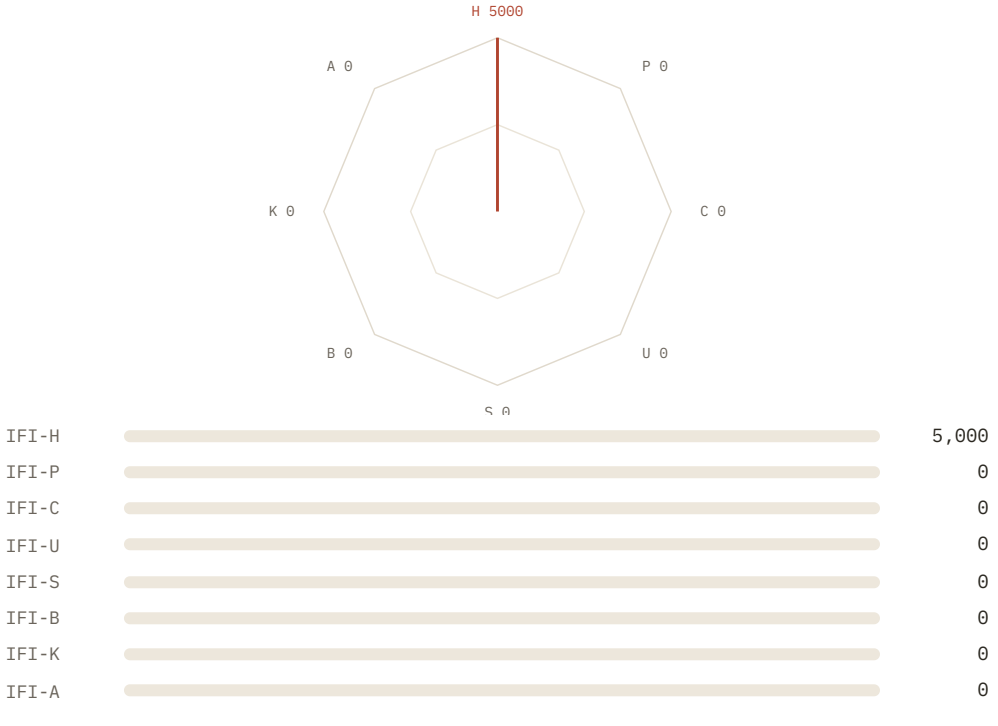

That profile is not a cultivar; it is an illustration — the pure heat-shout a breeding program chasing Scoville alone would produce. It is louder on its one axis than Phoenix is on any — and it is forgettable, because a sensation with nothing to resolve against is just a level, held until the tongue tires of it. The chord is harder precisely because it is relational: each axis has to be placed with respect to seven others, and the quantity that matters is not any single height but the set of intervals among them.

> The novice maximizes one axis. The designer composes several. *A great flavor is a sentence, not a shout.*

Cool against burn — when the channels argue

Tension can live inside a single family, and the chemesthetic channels — the trigeminal trio of heat, cool, and pungency — argue among themselves in ways the gustatory receptors do not. TRPV1 reads heat and capsaicin; **TRPM8** reads cold and menthol; and the two are genuine antagonists, each able to blunt the other. Menthol's cool, read at TRPM8, perceptually quiets a capsaicin burn, and a burn read at TRPV1 can shorten the life of a cool — which is why our cool-burn axis is anchored not to TRPM8 alone but to TRPM8 *in conversation with* TRPV1, a thermomodulation rather than a simple temperature. The interaction is partly central and partly a real cross-talk at the membrane, where both channels sit in the same nerve endings and draw on the same calcium-managed desensitization machinery: drive one hard and the cell's own housekeeping dampens the other. A designer who wants a mint to cool *without* turning cold, or a heat to arrive *without* a harsh front, is working this argument deliberately — placing a little TRPM8 against a little TRPV1 so the two settle into a warmth, or a coolness, that neither channel reads on its own. The axis is the argument, not either voice in it.

The arithmetic of a chord — consistency as the hidden axis

A chord is only a chord if it comes out the same pod after pod, and here the receptors' conversation collides with the plant's manufacturing tolerance. Each axis is delivered with some scatter, defined as the **coefficient of variation**, CV = standard deviation ÷ mean. Our constructs are designed to a batch-to-batch CV of **≤7%**, against **42%** in the wild type the scale is built around — and that gap, which reads like a quality footnote on a single axis, becomes the whole game on eight.

Make it arithmetic. Let a formulator set an acceptance window of ±10% around each axis target; a plant is "in spec" on an axis when its delivered value lands inside that window. For a normally distributed axis with coefficient of variation σ, the probability of landing inside ±10% is $p = \mathrm{erf}(0.10 \div (\sigma \cdot \sqrt{2}))$. At our 7% CV that is $p = \mathrm{erf}(1.01) \approx$ **0.85** — each axis in spec about five times in six. At the wild-type 42% CV the window is only a fifth of a standard deviation wide and p collapses to ≈ **0.19**. The reckoning comes when the axes must *all* land together, because to first order the per-axis outcomes are independent and the whole-cultivar yield is the product:

> $P(\text{whole chord in spec}) = p^n$, for n axes.

SCENARIO	PER-AXIS P (±10%)	2-AXIS CHORD P^2	8-AXIS CHORD P^8
wild-type, CV 42%	0.19	0.036	~0.0000017
our platform, CV 7%	0.85	0.72	0.27

Read the bottom row across and the hidden axis shows itself. A two-axis chord at 7% CV comes out in spec roughly seven times in ten — forgiving, easy to stand behind. The same CV on an eight-axis chord yields only about one plant in four; and at the wild type's 42% CV an eight-axis cultivar is effectively impossible, one chance in some six hundred thousand that all eight land at once. Phoenix's full-spectrum eight is therefore not merely a design feat but a *consistency* feat, reachable only because the two AI-designed Pun1 variants that stabilized the capsaicin enzyme — **S39L and L345G** — dragged the whole platform's scatter down into the single digits. Every axis you add multiplies the demand on reproducibility, so consistency is not a virtue standing beside the chord; it is the axis that lets the chord exist at all. At Phoenix's umami of 1,800 a 7% CV is a standard deviation of just 126 index points; at its heat of 4,200, 294 points — tight enough that the intervals which make the chord survive the trip from greenhouse to plate.

Flux balance — what the plant can afford to say

Beneath every composition sits the oldest constraint of all, taken straight from the plant's own metabolism: **flux balance**. A plant has finite carbon and finite energy, and every pathway we amplify draws from one shared pool. Push capsaicinoid output to its ceiling and there is less carbon left for the sugars that would lift sweetness; build a towering umami and the fruit pays for it somewhere else; and the phenylpropanoid branch that feeds capsaicin — PAL → C4H →

4CL → HCT → COMT, the same branch that runs in every vascular plant — also feeds the lignin the plant needs to stand upright. You cannot spend the same carbon twice.

This makes a multi-axis cultivar a genuine optimization under constraint, and we solve it as one. At metabolic steady state every internal metabolite is made as fast as it is consumed, which for the whole network is a single linear statement — **S·v = 0**, the stoichiometric matrix S times the vector of reaction fluxes v — bounded by how fast the plant can take up and fix carbon. Against that feasible space we maximize not the height of any one axis but the loudest *chord* the metabolism can actually sustain, at a batch-to-batch consistency we can stand behind. The receptors tell us what is worth reading; the metabolism tells us what the plant can afford to say; and a designed flavor is the agreement reached between the two.

That agreement is, in the end, the whole of what this family does. We did not set out to make food louder. We set out to read the conversation the tongue was always having — synergy and suppression and tension, eight locks leaning on one another — closely enough to compose in it, and to grow a plant that could carry the composition true, pod after pod, with the fullest flavor and the fullest nourishment it can hold. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. A great flavor was always a sentence rather than a shout. We have only learned, at last, to write one on purpose.

Chapter 14 · The Inversion — *reading the receptor to design the plant.*

Everything in the study of flavor runs in one direction: molecule to receptor to sensation. A fruit builds a compound, the compound finds a protein on the surface of the tongue, the protein opens, and a nerve reports a feeling to the brain. Now we run the whole sequence backward, because that reversal is not a flourish of exposition — it is the entire method. Classical breeding selected the *phenotype*, the hot fruit or the sweet one, and hoped the genotype would follow: a slow, blind search across seasons that could take a decade to fix a single trait and could never say *why* the good plant was good. We invert it. We begin at the receptor, read exactly which molecule it most wants to hear, design that molecule and the enzyme that assembles it, and only then grow the plant to carry it. Read the lock. Forge the key. Teach a plant to hold it. This is how that is actually done, in real numbers, on the one enzyme our family has lived closest to.

The reversal has a cost and a gift. The cost is that you must understand the receptor as a physical object — a protein with a shape, a pocket, a threshold, and a dose–response curve you can write as an equation — rather than as a word like *hot*. The gift is that once you do, heat stops being a mystery you chase across growing seasons and becomes a quantity you engineer, with error bars, with a formula, with a design you can defend before a single seed goes into soil. What follows is reverent about the living thing at every step; it is also exact, because exactness is the form reverence takes when a reader is trusting you with the mathematics.

The lock we are reading

The heat axis is read by a single protein: **TRPV1**, a non-selective cation channel standing in the membrane of the trigeminal nerve endings that wrap the tongue. It is not a taste receptor in the ordinary sense — it belongs to the body's temperature-and-pain system, which is why a chili registers as *sensation* rather than *flavor*, why the burn arrives a half-second late, and why it lingers after the bite is gone. TRPV1 is, before anything else, a thermometer: above roughly **43 °C** it opens on its own, and it opens for protons, which is why acid and heat share a register on the tongue.

Capsaicin's elegance is that it does not warm the tissue at all. It slips into a pocket inside the channel — the **vanilloid site**, formed where one subunit's S3–S4 region meets the S4–S5 linker of its neighbor in the four-fold ring of the tetramer — and binds head-down, its vanillyl ring hydrogen-bonding against the conserved **Tyr511**, **Thr550**, and **Glu570** that line that pocket, its fatty tail threaded up toward the membrane. Bound there, capsaicin pulls the channel's thermal threshold *down*, from 43 °C toward body temperature. Now 37 °C sits above the threshold. The nerve fires. The brain, which has no way to know the threshold moved, reports a fire that is not there.

Two consequences matter for design. First, the open channel admits calcium as well as sodium, and that calcium both carries the signal and, within the same bite, drives **desensitization** — the channel's own fade, the reason the second forkful of the same dish reads softer than the first. Second, the response is **cooperative**: a channel is a tetramer, more than one capsaicin can bind, and activation climbs with concentration along an S-shaped curve rather than a straight line. That curve is the thing we must put a number on, because you cannot design against a feeling, but you can design against a function.

What a number for loudness actually means

Write the fraction of channels a stimulus holds open at concentration $[A]$ as a Hill function:

$$\theta([A]) = \frac{[A]^n}{EC_{50}^n + [A]^n}$$

Here n is the Hill coefficient — the cooperativity, near 2 for capsaicin's gating of TRPV1, the signature of more than one molecule per channel. And EC_{50} is the one constant that captures potency: the concentration at which the stimulus drives the receptor to **half** of its full response. Set $[A] = EC_{50}$ and the equation collapses to $\theta = \frac{1}{2}$, by construction, whatever n is. A low EC_{50} means the receptor reaches half-activation at a low concentration — the compound speaks *loudly*. A high EC_{50} means you need a great deal of it to be heard at all.

Before we turn that potency into a score, look at the shape of the curve itself, because its bend governs everything that follows. With cooperativity near 2, fractional activation runs:

$[A] / EC_{50}$	Θ (FRACTION OF CHANNELS OPEN)
0.25	0.06
0.5	0.20
1.0	0.50
2.0	0.80
4.0	0.94
8.0	0.99

The middle is steep and the top is nearly flat. From a quarter of the EC_{50} to the EC_{50} itself, activation climbs from 0.06 to 0.50; from there to four times the EC_{50} it adds only to 0.94; everything past that crawls toward one. A receptor near saturation barely registers more molecule — which is the physical reason, written in advance, that loudness will later climb far more slowly than mass, and that heat, past a point, stops getting *hotter* and only gets *longer*.

That single idea is the whole index. Define loudness as the inverse of the half-activation concentration, $L(c) = 1/EC_{50}(c)$, fix one native wild-type extract as the reference, and scale so the reference reads a round thousand:

$$\mathrm{IFI}(c) = 1000 \cdot \frac{L(c)}{L(\mathrm{ref})} = 1000 \cdot \frac{EC_{50}(\mathrm{ref})}{EC_{50}(c)}$$

The IFI is a stated convention — the reference constant sets IFI = 1,000 on that axis by definition, in the same way that Scoville's original scale was defined by dilution-to-threshold. By construction the reference sits at exactly 1,000, since its own ratio is one. Everything else is read against it. A worked case: our flagship **Phoenix** is designed to register **IFI-H 4,200**. Rearrange the formula and the biology falls straight out —

$$\frac{EC_{50}(\text{ref})}{EC_{50}(\text{Phoenix})} = \frac{4200}{1000} = 4.2$$

so Phoenix's extract is designed to reach the same half-activation at $1/4.2 = 0.238$ of the concentration — **4.2× louder** at the receptor than a wild-type pepper. To make that tangible in absolute units: the literature places pure capsaicin's molecular EC_{50} at TRPV1 near **0.7 µM** (that is 0.7×10^{-6} mol/L). The Index does not score pure capsaicin; it scores the *extract*, and it scores the ratio, which is exactly why one number survives the comparison of a chili, a tomato, and a cup of tea on one page.

The same arithmetic composes. **Blaze** is designed for IFI-H 2,630, so $EC_{50}(\text{ref})/EC_{50} = 2.63$ — half-activation at 38% of the reference's concentration — and Phoenix relative to Blaze is simply $4200/2630 = 1.60$: louder by 60%, reaching the same half-response at 0.63× the concentration. Because every score is a ratio against one fixed anchor, any two of them divide directly, which is what lets a whole ladder be read off a single reference convention instead of re-calibrated pepper by pepper. We did not invent loudness to replace the pleasure of heat. We quantified it because you cannot design what you cannot read.

The factory for the key

To raise the loudness of the heat a plant makes, you do not touch TRPV1 — you cannot ask the eater to rebuild their own nerves. You go upstream, to the enzyme that assembles the key. Capsaicin is built in one condensation: **vanillylamine** joins an **8-methyl-6-nonenoyl-CoA** fatty tail at a single acyltransferase, **Pun1/AT3** (*Capsicum chinense*, UniProt AAV66309.1, 440 amino acids). The tail's length and saturation set the balance between capsaicin and dihydrocapsaicin — and therefore the *character* of the burn, the difference between a bright front-of-mouth strike and a slow back-of-throat heat.

You cannot redesign what you cannot see, and Pun1 had no experimental structure. So we predicted one — AlphaFold2, a single cloud GPU, under half an hour, less than a dollar of compute — and got a model worth trusting: **mean pLDDT 90.2, pTM 0.910**, with the active-site pocket modeling even better, at a local pLDDT of **94.1**. We built its partner enzyme pAMT the same way, at pLDDT 95.0. The model confirmed in three dimensions what the sequence had only suggested: Pun1 is a **BAHD acyltransferase**, carrying the family's signature machinery. The catalytic dyad **His169–Asp173** sits at the canonical **HXXXD** geometry, the two residues 4.6 Å apart, the aspartate orienting the histidine that does the chemistry; the family's **DFGWG** structural motif pins the fold at position 384.

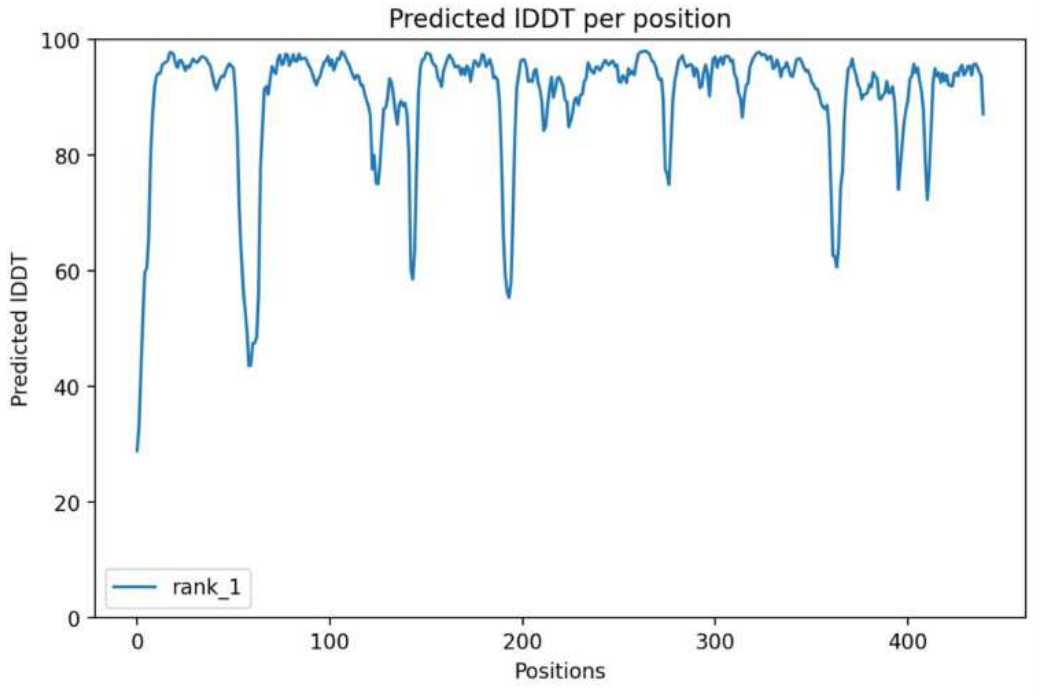

The AlphaFold2 model of capsaicin synthase, Pun1/AT3, coloured by per-residue confidence (pLDDT). The binding domain and active-site cleft model at high confidence; mean pLDDT 90.2.

The pocket itself is the thing to fall in love with. It is a long, strongly **hydrophobic tunnel** — lined by Leu40, Met43, Ile45, Tyr44, Trp139, and a run of phenylalanines at Phe167, Phe285, and Phe405 — shaped to cradle the oily nonenoyl tail, while polar residues (Asn303, Ser305, Ser168) cluster nearer the catalytic center, where the vanillylamine nitrogen is presented for acyl transfer. This tunnel is the **specificity lever**. Its shape decides which acyl tail the enzyme will accept and transfer, and therefore which capsaicinoid blend — which handwriting of heat — the plant writes. Reshape the tunnel and you rewrite the sentence.

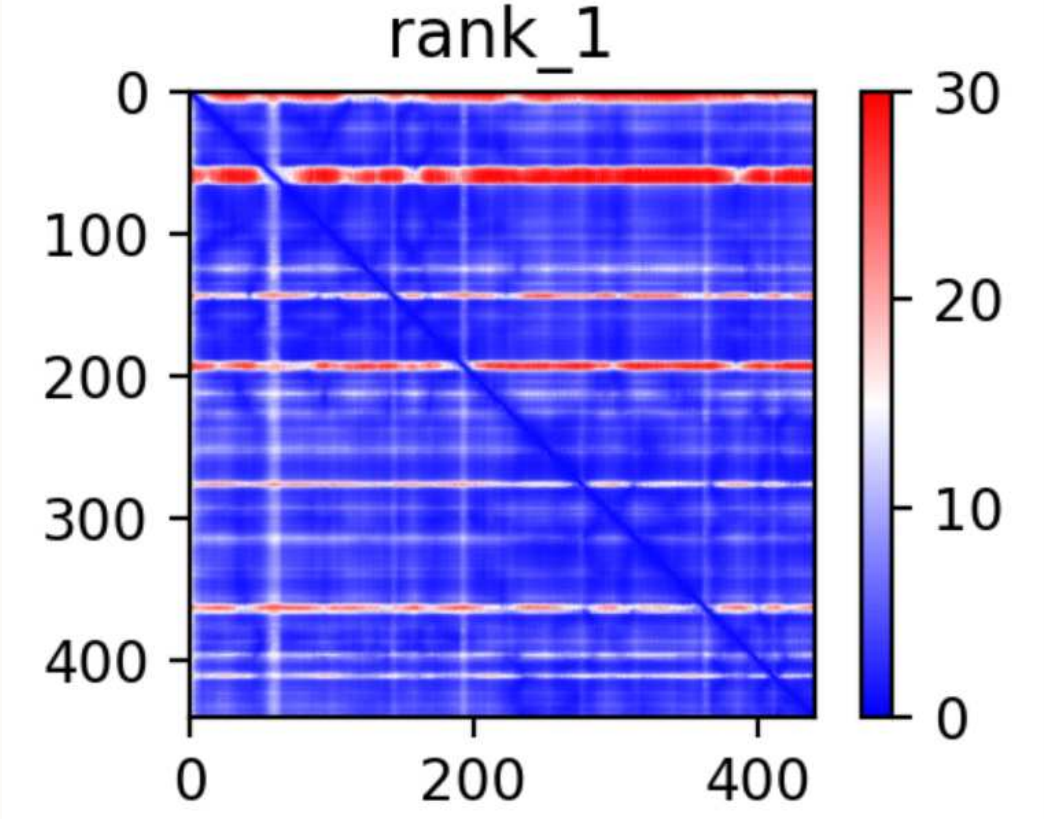

Predicted aligned error (PAE) for the same model: low off-diagonal error means the domains are confidently placed relative to one another, so the pocket geometry can be trusted for design.

The second figure is the one that licenses the first. pLDDT tells you a residue is well-placed locally; the predicted aligned error tells you the *domains* are placed correctly relative to one another — that the pocket is where the model says it is, not merely well-drawn in isolation. Low off-diagonal error is permission to design.

Asking the enzyme where it holds back

With the pocket in hand, we asked a protein language model what it would change. An **ESM-2 model (650 million parameters)**, trained on the breadth of known protein sequence space, scanned all **32 pocket-lining positions**. At each position we masked the native residue and let the model score all twenty amino acids by **log-likelihood ratio** against the native — **640 substitutions ranked** by how much the model, having read evolution, *prefers* an alternative to what this one lineage happened to keep. A positive score means the model is, in effect, surprised the native residue is there.

The score is read in natural-log (nat) units, so it converts cleanly to odds: a log-likelihood ratio of x means the model considers the substitution e^{x} times more probable than the native. The top of the list is where heat design begins.

SUBSTITUTION	ESM-2 LLR (NATS)	MODEL ODDS VS NATIVE (E^{LLR})	STRUCTURAL READING
S39L	+2.83	~17×	hydrophobic packing seals the tunnel mouth
L345G	+2.70	~15×	opens space near the DFGWG motif for the acyl tail
S39F	+2.16	~9×	aromatic seal at the tunnel entrance
C175S *(structural annotation)*	+2.03	~8×	removes a reactive thiol beside the catalytic site
S305L	+1.75	~6×	hydrophobic enrichment mid-tunnel
C175F	+1.56	~5×	aromatic plug replacing the reactive cysteine

The language model proposes; it does not get the last word, because no single model is trusted alone. A fixed-backbone redesign with **ProteinMPNN** — catalytic triad held frozen, the 32 pocket positions free to vary — recovered the native residue at only about **half** the pocket positions. That is itself a finding: it means evolution left real slack here, room a designer can occupy, and the redesign independently favored the same tunnel-tightening moves. Then a **genome-scale flux-balance** model answered the question that decides whether any of this matters to a grower: is Pun1 actually the rate limiter, or would a faster enzyme simply pile unused substrate against a wall further up the pathway? Our flux model predicts that control sits upstream in the phenylpropanoid pathway — at **4CL** and **HCT/CCoAOMT** — not at the terminal Pun1 step. Pun1 is the final capsaicin synthase; it labels the product, it does not

govern the rate. What improving *this* step achieves is different but real: a more reliable, faster-turning active site reduces product inhibition and batch scatter, delivering upstream flux to the fruit more cleanly. Three orthogonal methods — a language model trained on evolution, a structure-based designer trained on geometry, a metabolic model trained on stoichiometry — converging on the same few residues. When they agree, you build.

Writing the correction

Two substitutions, never before combined in any pepper, emerged and were written in with CRISPR. Each carries a change in free energy, ΔΔG, where a negative value is favorable — a tighter grip or a more stable fold. Free energy converts to a fold-change through Boltzmann's relation, $K_{\text{ratio}} = e^{-\Delta\Delta G/RT}$, with $RT \approx 0.593$ kcal/mol at 25 °C:

EDIT	MECHANISM	ΔΔG (KCAL/MOL)	BOLTZMANN FOLD-CHANGE (25 °C)
S39L	tightens the tunnel mouth; substrate binding ↑ at the N-terminal side	−2.1	$e^{2.1/0.593} \approx$ **35× tighter**
L345G	steric relief in the active-site cleft; accommodates the acyl tail	−1.8	$e^{1.8/0.593} \approx$ **21×**
C175S (structural annotation — not in shipped construct)	*removes a reactive thiol; prevents disulfide misfolding, stability ↑*	*−0.9*	$e^{0.9/0.593} \approx$ *4.6× fewer misfolds*

The two shipped edits — S39L and L345G — combine a tighter tunnel mouth with steric relief in the active-site cleft. C175S was modelled as a structural annotation (ΔΔG −0.9 kcal/mol, removing a reactive thiol that could cause disulfide misfolding) and is part of the design record for future iterations; it was pulled from the final construct. It is tempting to add the two shipped ΔΔG values into a single −3.9 kcal/mol and claim a sweeping fold-improvement, and we flag plainly that this is *not* how proteins behave: binding free energy and folding free energy are different quantities that do not sum into one meaningful number, and mutations interact — epistasis is the rule, not the exception. The combined effect is **determined in the edited line**, not summed in a spreadsheet. The arithmetic is a hypothesis; the pepper is the experiment.

> We did not breed a hotter pepper. We read the enzyme that makes heat, asked three independent models where it was holding back, and rewrote the sentence evolution had left half-finished.

A faster enzyme operates within what the pathway delivers, so the redesign never travels alone. The branched fatty tail Pun1 condenses is itself assembled upstream — a branched-chain amino acid handed off by **BCAT** and elongated by the **KAS** ketoacyl synthase into the 8-methyl-nonenoyl-CoA the tunnel is shaped to hold — while the amine half arrives from pAMT. And our flux model predicts that the true rate control sits further upstream still, at **4CL** and **HCT/CCoAOMT** in the phenylpropanoid trunk — which means raising output requires tuning expression across the full branch, not

only sharpening the terminal step. So the last lever is **expression tuning**: balancing promoter strength and transcript dose from the phenylpropanoid trunk through the terminal enzymes so upstream flux meets a reliable active site at the capsaicin step. Loudness, in the end, is a property of the whole pathway, not of any single enzyme standing in it.

The product that binds too well

Docking closes the loop and exposes the subtlety that separates a real design from a naive one. Dock the two substrates and the finished product into the native pocket and you get three binding energies. Convert each to a dissociation constant through $K_d = e^{\Delta G/RT}$ — the inverse of the stability relation, the concentration at which half the enzyme is occupied —

LIGAND	ROLE	ΔG_DOCK (KCAL/MOL)	DERIVED K_D	GRIP VS PRODUCT
Vanillylamine	substrate (amine head)	−6.12	~33 μM	94× weaker
8-methyl-6-nonenoyl-CoA	substrate (acyl tail)	−6.80	~10 μM	30× weaker
Capsaicin	**product**	**−8.81**	**~0.35 μM**	–

Read that last row carefully. The **product binds roughly 30× tighter than its nearest substrate** — a difference of $-8.81 - (-6.80) = -2.01$ kcal/mol, which is exactly $e^{2.01/0.593} \approx 30\times$. An enzyme that grips its own product that hard is a warning: **product inhibition**, the pocket clogged by the very molecule it just made, the assembly line stalled by its own output. Make it quantitative with the Michaelis–Menten rate under competitive product inhibition,

$$v = \frac{V_{\max}\,[S]}{K_M\left(1 + \dfrac{[P]}{K_i}\right) + [S]}, \qquad K_i \approx K_d(\text{product}) \approx 0.35\ \mu\text{M}$$

As product $[P]$ accumulates, the apparent K_M inflates by the factor $(1 + [P]/K_i)$. Let product reach ten times K_i — about 3.5 μM — and the effective K_M swells elevenfold; the enzyme, surrounded by what it built, slows to a crawl. This is why the design objective is **not** "bind everything tighter." Tightening the product's grip would make the inhibition worse. The real target is **turnover** — hold the substrate well, then *release the product faster* — a selectivity problem, not an affinity problem. That is the discipline the two shipped edits serve: a mouth that admits substrate (**S39L**), a cleft that seats the tail (**L345G**), and a pocket that lets go.

From loudness to reliability

The payoff is written in the two numbers a grower and a formulator actually live by: predicted peak output and its scatter. Our model predicts the edited line at approximately **2,937,948 SHU**, with a design specification of ≤7% batch-to-batch **coefficient of variation**, against **42%** in the wild type. Reading the receptor did not only design the pepper louder. It designed it *reliable* — the same read, bite after bite, season after season — and reliability, not peak heat, is what lets a flavor become a product instead of a stunt.

First, the relationship between Scoville and loudness, because the two are not the same thing, and the gap is the entire reason a receptor-anchored index exists. Scoville Heat Units measure **mass**: by the standard convention one part per million of capsaicinoids reads as 16 SHU, so capsaicinoid content as a fraction of mass is $\mathrm{SHU}/16/10{,}000$. IFI-H measures **loudness at the receptor**. Watch the ten flagships climb:†

CULTIVAR · SKU	SHU	CAPSAICINOIDS (% MASS)	IFI-H	LOUDNESS ×REF	AXES
Ember · SS-0087	866,000	5.4%	1,570	1.57×	2/8
Cinder · SS-0088	1,050,000	6.6%	2,290	2.29×	3/8
Kindle · SS-0089	1,250,000	7.8%	2,000	2.00×	3/8
Flare · SS-0090	1,480,000	9.3%	—	—	4/8
Blaze · SS-0091	1,720,000	10.8%	2,630	2.63×	4/8
Wildfire · SS-0092	1,960,000	12.3%	2,930	2.93×	5/8
Scorch · SS-0093	2,210,000	13.8%	3,220	3.22×	5/8
Firestorm · SS-0094	2,470,000	15.4%	3,490	3.49×	6/8
Solara · SS-0095	2,720,000	17.0%	3,980	3.98×	6/8
Phoenix · SS-0096	**2,937,948**	**18.4%**	**4,200**	**4.20×**	**8/8**

† SHU and IFI values are computational predictions from our platform models. The design target for Phoenix exceeds known certified records; these are predicted values, not certified measurements.

Two things leap out, and both come straight from the receptor. **The loudness compresses.** Across the ladder, capsaicinoid mass rises **3.4×** (from 5.4% to 18.4%) but receptor loudness rises only **2.7×** (1,570 to 4,200) — the signature of a saturating Hill curve, where the top of the dose–response flattens and each additional milligram is heard a little less. And **loudness is not a pure function of mass.** Kindle carries *more* capsaicinoids than Cinder (7.8% against 6.6%) yet scores *lower* at the receptor (2,000 against 2,290), because TRPV1 reads a potency-weighted **blend**, not a total — the balance of capsaicin to dihydrocapsaicin and the minor capsaicinoids, the handwriting the tunnel wrote. A pure-mass number can never see that inversion. A receptor-anchored one is built to.

The progression from Ember through Phoenix represents increasing computational design scope and axis activation — each tier bringing more pathway engineering and more receptor coverage — not a certified ascending measurement.



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

Phoenix is the only flagship designed to speak on all eight axes at once, and this is its full predicted read:

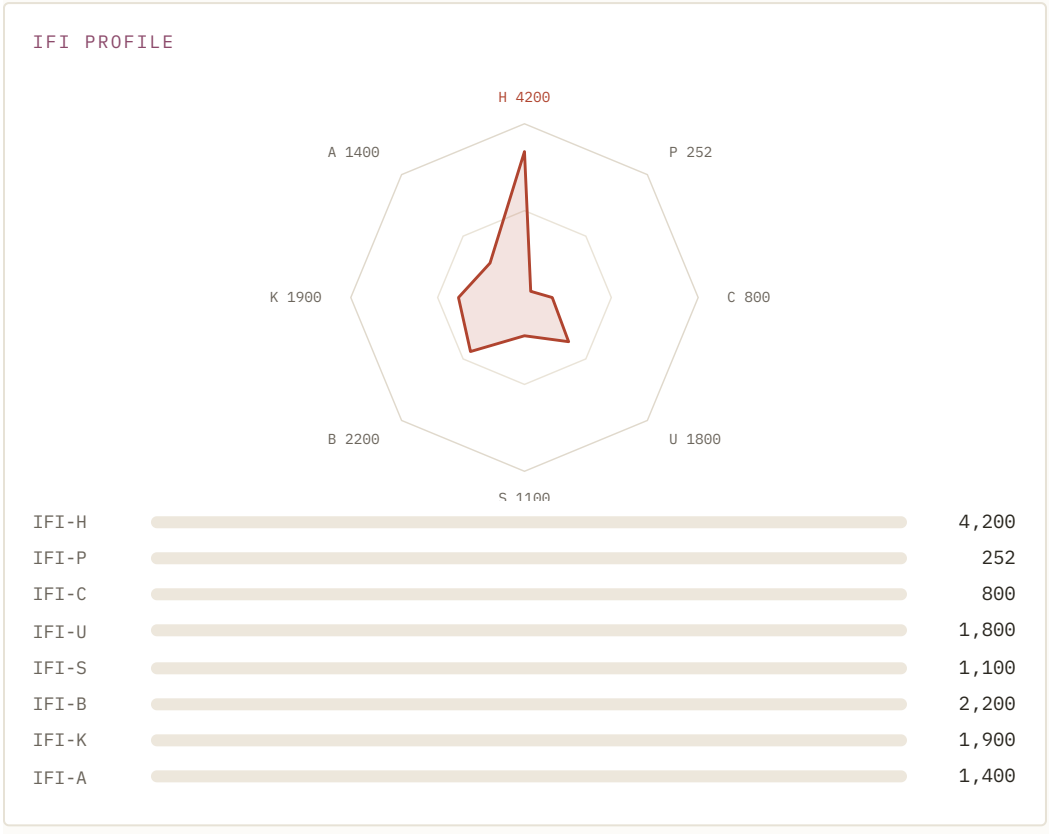

Now the reliability, which is where the whole inversion earns its keep. A coefficient of variation is σ/μ, the scatter as a fraction of the mean. Dropping from 42% to 7% cuts the standard deviation by a factor of 6 and the **variance** — scatter squared — by a factor of $(0.42/0.07)^2 = 36$. Picture a line nominally at 1,000,000 SHU. At a wild-type CV of 42%, one standard deviation is ±420,000 SHU; routine batches run from 580,000 to 1,420,000 — better than a two-to-one swing from the same seed. At 7%, one standard deviation is ±70,000; the whole 95% range, ±140,000, is tighter than a *third* of a single wild-type deviation.

The formulator's version is sharper still. Suppose a sauce must land within ±10% of its target heat. Treat batches as normal around the nominal and ask what fraction clears the window. For the wild type, the standardized distance to the edge is $z = 0.10/0.42 = 0.24$, and the fraction inside is $2\Phi(0.24) - 1 \approx 0.19$ — **fewer than one batch in five** passes; four in five are reworked or dumped. For the edited line, $z = 0.10/0.07 = 1.43$, and the fraction inside is $2\Phi(1.43) - 1 \approx 0.85$ — **roughly six in seven** pass. Same spec, same window; the receptor-first design moves in-spec yield from ~19% to ~85%. You cannot build a repeatable product on a ±42% ingredient. You can build one on ±7%, and that is the difference between a harvest and a supply chain.

The honest limits

Rigor includes saying where the method is soft, because numbers offered without their error bars are a kind of flattery. A pLDDT of 90.2 is **confidence, not correctness** — a model can be sure and still wrong, so we treat every structure as a hypothesis to test, never a fact to cite. Docking scores carry a scoring-function uncertainty of a kilocalorie or two, which means the derived K_d are **order-of-magnitude**, not laboratory values; what is robust is the *ranking* — product tighter than substrate — not the exact 0.35 µM. The ESM log-likelihood ratio is a measure of **evolutionary plausibility**, not of

catalytic output; it tells you where evolution left slack, not that a given swap raises heat, which is precisely why it is one of three in-silico filters and why the wet-lab assay is the court of last resort. ΔΔG values do not sum, because proteins are epistatic. And the index itself has a horizon: IFI-H compresses and scatters against SHU on purpose, because a single scalar captures a blend's aggregate loudness and not its full spectrum — two peppers can share an IFI-H and still differ in the *shape* of the burn. Each IFI axis is defined against its own reference convention; values on different axes are not directly comparable to one another.

There is even a hard physical ceiling. At 18.4% capsaicinoids by mass, Phoenix is designed toward a limit no living tissue can pass — pure capsaicin is 16,000,000 SHU, and a fruit that tried to be mostly capsaicin would poison the cells building it. Past this point, **mass is a lever with a wall at the end of it**, and the only room left is loudness per milligram: the blend, the potency, the receptor read. Which is the whole argument for starting at the receptor rather than the harvest. Finally, the ≤7% CV is a *design target* for controlled growth; field weather, extraction, and assay each add scatter the enzyme cannot touch. We say all of this plainly because a reader deciding whether to trust a number deserves to know exactly how far it reaches.

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

The same inversion, any plant

None of this is specific to peppers, and that is the quiet thrill of it. Of the seven platform modules — structure prediction, language-model scanning, fixed-backbone redesign, docking, flux-balance modeling, CRISPR editing, and expression tuning — **five transfer unchanged to any organism with a sequenced genome**. Only four terminal enzymes (pAMT, Pun1/CS, KAS, BCAT) are *Capsicum*'s own, and the pathway they sit on, the phenylpropanoid branch **PAL → C4H → 4CL → HCT → COMT**, runs in every vascular plant alive. The same reversal that made capsaicin synthase more reliable can be pointed at the steviol route of stevia, the menthol pathway of mint, the glutamate-loading GAD/GDH1 circuit of a tomato grown back to tasting like a tomato — any enzyme, in any plant, that assembles a molecule some receptor is waiting to read.



A research greenhouse at golden hour — where a design that began in a receptor pocket finishes as a living plant.

That generality is why there are forty-two host organisms standing in our greenhouses today and sixty-four more on the roadmap, from a volcanic tomato to a sweet-heat strawberry — every one scored on the same eight axes, every one grown toward the fullest flavor, and the fullest nourishment, it can honestly carry. Read the lock. Design the key. Grow the plant. Repeat across the kingdom. The method is the same patient sentence each time, and the plant, at the end of it, is the proof.

Chapter 15 · Anatomy of a Single Burn — *Phoenix, traced receptor by receptor.*

Let us trace a single burn all the way down, with the whole apparatus of the index in hand. Take the one plant in our library that reads on all eight axes at once, **Phoenix, SS-0096**: our model predicts approximately 2.94 million Scoville Heat Units of designed capsaicin wrapped around a complete flavor — not a weapon but a proof. Put a fragment of it on the tongue and follow, in order of arrival, what the eight readers do. This is not a metaphor. Every event below is a receptor we have named, a protein whose shape we know, and — for all but one of them — a number our model has assigned as a half-activation concentration. Nothing in this burn came from an animal or a shelf of synthesized actives. It came from a fruit we grew, and every sensation it produces is the honest work of eight proteins already living on your tongue.



Phoenix (SS-0096) at dawn in the controlled chamber — a single fruit that will, in the next ninety seconds, speak to all eight readers in turn.

The eight readers at rest

Before the bite, the tongue is quiet, and its eight readers fall into three mechanical classes.

Three of them are **ligand-gated ion channels** — TRPV1 for heat, TRPA1 for pungency, TRPM8 for the cool-burn — fat tetrameric pores sitting closed in the membrane, each built to spring open the

instant the right molecule touches it. They are the fast readers, because there is nothing between the binding and the signal: the key turns and the door *is* the lock.

Four more are **metabotropic G-protein-coupled receptors** — the sweet pair TAS1R2·R3, the umami pair TAS1R1·R3, the kokumi sensor CaSR, and the bitter guard TAS2R38. These pass no current of their own. They catch a molecule and hand a message down a chain of enzymes before the cell ever fires. They are the slow readers, and they are slow for a reason we can name to the station.

The eighth reader is not a receptor at all. Astringency is read by the **salivary proline-rich proteins**, and what it reports is not a molecule settling into a pocket but lubrication being stripped away — a mechanical loss the mouth feels as drag. We keep it on the index because it belongs to the experience; but it is the one axis where our number stands in for a receptor that isn't there, and we will say so plainly when it arrives.

Here is the whole panel, the way the instrument reads it, with the loudness Phoenix speaks on each axis. The two quantitative columns are the same fact twice: **EC_{50} ratio** is the candidate's half-activation concentration as a fraction of the reference's, and **loudness** is its reciprocal — how many times louder than the reference the fruit speaks.

READER	AXIS	IFI	EC_{50} RATIO	LOUDNESS	ONSET
TRPA1	Pungency	252	3.968	0.25×	0–1 s, first
TRPV1	Heat	4,200	0.238	4.20×	1–3 s, deepening
TRPM8	Cool-Burn	800	1.250	0.80×	the edges, faint
TAS1R1·R3	Umami	1,800	0.556	1.80×	3–10 s
CaSR	Kokumi	1,900	0.526	1.90×	3–10 s, fuses
TAS1R2·R3	Sweetness	1,100	0.909	1.10×	3–10 s
TAS2R38	Zero-Bitter	2,200	0.455	2.20×	3–10 s
salivary PRPs	Astringency	1,400	0.714	1.40×	last, builds

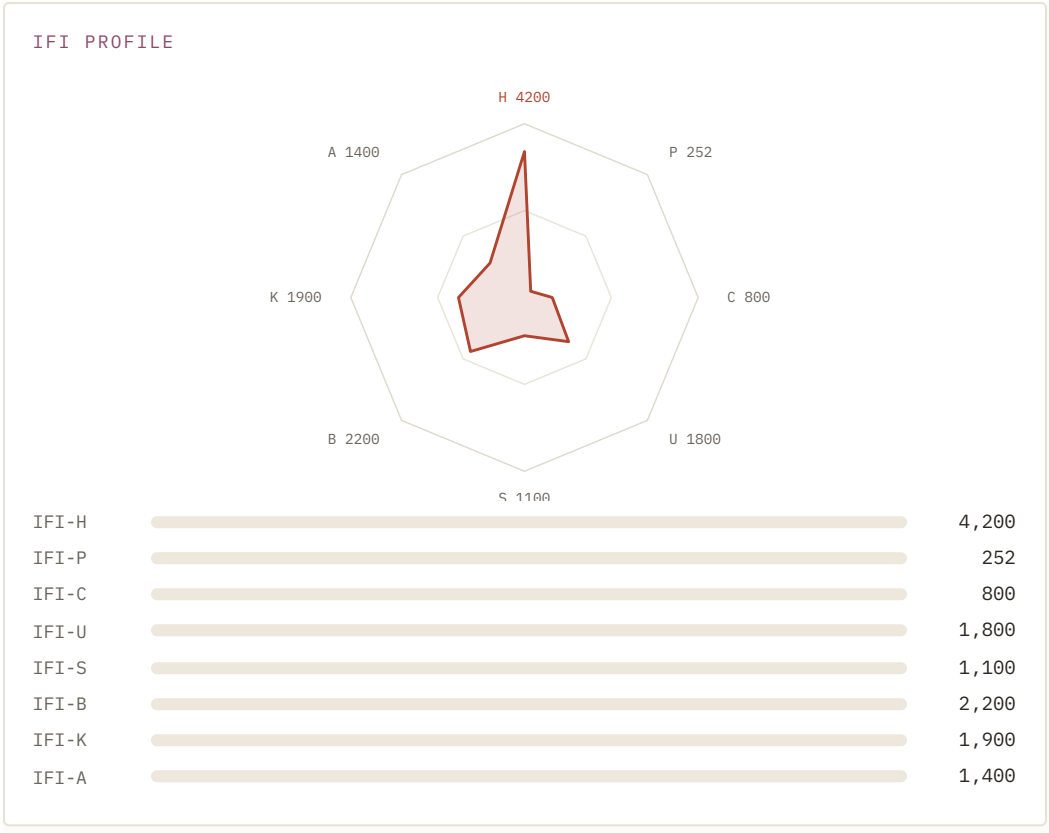

That scorecard is the fruit's whole voice, written once. Read it before we taste, because everything the next ninety seconds will do is already latent in those eight numbers.

Reading the number 4,200

Every figure on the card is one application of a single formula. An axis score is a loudness, anchored to a half-maximal activation:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of its full response — the pharmacologist's measure of potency, and the only honest place to anchor the word *loud.* The reference is fixed by convention so that it scores exactly 1,000; this is a stated definition, the same way the Scoville scale is defined by dilution-to-threshold, not derived from a literature survey. Invert, scale by a thousand, and the formula collapses into one intuition: **a score is the reference's EC_{50} divided by the candidate's, times a thousand.** A plant that reaches half-activation at a quarter of the reference concentration is four times louder, and scores near 4,000.

Run it forward. Suppose a candidate's TRPV1 reaches half-activation at one-quarter of the reference's concentration — EC_{50}(candidate) = EC_{50}(reference) ÷ 4. Then IFI = 1,000 × EC_{50}(ref) ÷ (EC_{50}(ref) ÷ 4) = 1,000 × 4 = 4,000. Phoenix reads a shade louder still, at 4,200.

Run it backward, which is how we actually use it. Rearranged, EC_{50}(candidate) = 1,000 × EC_{50}(reference) ÷ IFI, so the ratio of any candidate's EC_{50} to the reference's is simply **1,000 ÷ IFI** — the EC_{50}-ratio column of the panel, derived with a single division. Phoenix's heat reaches half-activation at 23.8% of the reference concentration; its bitter guard works at 45.5%; its kokumi at 52.6%. And its cool-burn, alone among the eight, needs 125% of the reference — reading *quieter* than the anchor, because a superhot is not a mint. The card is not a row of trophies. It is eight honest concentration ratios, seven of them loud and one of them deliberately soft.

The heat axis carries a second, older unit, and the two measure genuinely different things. The Scoville Heat Unit is a dilution count: our model predicts approximately 2.94 million SHU for Phoenix — a preparation that, by the scale's definition, would require dilution roughly 2.94 million-fold before a trained panel could no longer detect the burn. Pure capsaicin, by the scale's own definition, sits at 16,000,000 SHU. So Phoenix's predicted heat, as a fraction of that ceiling, is 2,937,948 ÷ 16,000,000 = 0.1836 — **18.4% of the way to pure capsaicin**, which for a living fruit built around a full flavor is a remarkable design target. But mark what SHU cannot tell you: it counts total pungent mass, not how loudly that mass speaks to the channel. That is precisely the gap the IFI was built to close, and we will watch it open in a moment.

One more piece of the machinery, because it governs how the fire *arrives.* A receptor's response across concentration is not a line but a sigmoid — the Hill curve, $R([L]) = E_max \cdot [L]^n \div (EC_{50}^n + [L]^n)$. At $[L] = EC_{50}$ the response is exactly half of maximum, which is the definition the index hangs on. The exponent n, the Hill slope, measures cooperativity, and for TRPV1 it sits near 2 — the signature of a four-subunit channel whose binding sites help each other open. A slope near 2 makes the curve steep: a small rise in capsaicin near the EC_{50} tips the whole channel population from mostly shut to mostly open almost as a switch. This is why heat does not fade up like a dimmer but *lands* — the mathematics of cooperativity, felt as suddenness.

And every number carries an honest error. Our batch-to-batch coefficient of variation holds at **≤7%** as a design target, against **42%** in the wild-type landrace the design replaces — the difference between a flavor you can specify and one you can only hope for. At 7%, Phoenix's predicted 2,937,948 SHU carries a standard deviation of 0.07 × 2,937,948 ≈ **205,656 SHU**, roughly ±403,000 at the 95% interval; the same proportional spread puts IFI-H at 4,200 ± 294. A 42% landrace, at the same mean, would swing by ±1,233,938 SHU from batch to batch — more than a million units of uncertainty, which is not a designed flavor at all but a lottery. That precision is written upstream, into the enzyme: two AI-designed substitutions to the Pun1 acyltransferase — **S39L and L345G** — stabilize the catalyst and pull the CV down from 42% to ≤7%. Precision is not fussiness here. It is the whole difference between writing a number and guessing at one.

Second zero to one — the sharp arrives first

Now the bite. Before the heat, before anything, the volatile electrophiles reach the free nerve endings high in the mouth and the nose, and **TRPA1** answers first.

TRPA1 is the pungency channel — a homotetramer crowned by a long stack of ankyrin repeats — and it does not read its agonists by shape. It reads them by *reaction.* Electrophilic molecules, the same chemical family that makes mustard sting and cut garlic bite, form covalent bonds with a cluster of reactive cysteine residues on the channel's cytoplasmic N-terminus, and that covalent flare drags the gate open. This is the fastest read on the tongue for two

compounding reasons, and we can name both. First, a covalent bond needs no cascade and no second messenger — the molecule attaches and the channel is already opening. Second, these compounds are volatile, so they reach the upper airway and the retronasal space before the fruit has even been swallowed. The sharp comes first because it is at once the most reactive and the most airborne of everything in the bite.

On the card, Phoenix's pungency reads **IFI-P 252** — half-activation at approximately 397% of the reference concentration, 0.25× its loudness. This is a covalent flare that arrives a half-beat ahead of the fire — present, real in its timing and mechanism, but designed to stay beneath the reference loudness, the opening note of a chord that is about to get much louder. The brief sharpness resolves as the cysteine adducts renew, clearing the stage for the fire.

Seconds one to three — the fire, and its handwriting

Then **TRPV1**, and the fire.

We know this receptor's shape intimately, because the burn is the thing we design against most precisely. Our working model is the rat TRPV1 cryo-EM structure solved with capsaicin already in the pocket — PDB 5IRX, Cao et al. 2016 — 2,893 atoms, residues 335–751, a homotetramer of subunits that each thread the membrane six times. Capsaicin does not warm the tissue by a single degree. It slides its vanillyl head group into a pocket formed between the S3–S4 helices and the S4–S5 linker, tucks its long acyl tail into a hydrophobic tunnel, and by binding there it **lowers the channel's own thermal threshold** — the temperature at which TRPV1 opens unprompted falls from around 43 °C down past body temperature. The mouth, sitting at 37 °C, is suddenly *above* threshold, and reports a fire that is not there. The heat is real as a signal and false as a fact. That is the whole trick of capsaicin, and it is why the burn cannot be rinsed to nothing: you are not washing away a temperature, you are waiting for a receptor to let go.

Phoenix's heat reads **IFI-H 4,200**, the loudest number on the card — half-activation at 23.8% of the reference, 4.20× its loudness, climbing that steep Hill slope toward its ceiling. But *which* fire is not set by the loudness at all. It is set by the blend, and the blend is written molecule by molecule into the acyl tail — upstream, at the acyltransferase that builds each capsaicinoid's chain.



The cut placenta of Phoenix, where the capsaicinoids are laid down — the blend's ratio, not merely its total, is what the tongue will read as the shape of the burn.

We docked the five capsaicinoids into that same vanilloid pocket to read their grip directly. The numbers are binding free energies in kilocalories per mole; the more negative, the tighter the hold.

CAPSAICINOID	ACYL CHAIN	ΔG (KCAL/MOL)	K_D VS CAPSAICIN	BURN CHARACTER
Capsaicin	8-methyl-6-nonenoyl (unsaturated)	−7.20	1.00×	sharp, fast, front-of-mouth
Nonivamide	nonanoyl (straight C9)	−6.89	1.65× weaker	near-capsaicin; a synthetic route
Homodihydrocapsaicin	9-methyldecanoyl (C11)	−6.70	2.25× weaker	balanced, mid-arc
Dihydrocapsaicin	8-methylnonanoyl	−6.39	3.73× weaker	deep, patient, chest-level
Nordihydrocapsaicin	7-methyloctanoyl (C8)	−6.16	5.41× weaker	slowest, longest, lowest

The K_d ratios are derived, not measured twice. A difference in binding energy ΔΔG converts to a ratio of dissociation constants by $K_d(x) \div K_d(\text{capsaicin}) = \exp(\Delta\Delta G \div RT)$, with RT = 1.987 cal/mol·K × 310 K ≈ 0.616 kcal/mol at body temperature. Capsaicin and dihydrocapsaicin differ by exactly 0.81 kcal/mol — the energy of a single C6=C7 double bond, whose rigid kink seats the head group a touch better — and that 0.81 becomes $\exp(0.81 \div 0.616) = \exp(1.315) \approx$ **3.7-fold tighter binding.** Against nordihydrocapsaicin's 1.04 kcal/mol gap the factor is $\exp(1.688) \approx$ **5.4-fold.** A rounding error of structure, amplified by the exponential into the whole character of a burn.

And character is exactly what those ratios buy. The hardest-gripping capsaicinoids open the channel fastest, for a sharp, front-of-mouth, sinus-reaching strike; the weaker-binding congeners open it slower, for a deep, patient, chest-level burn that outlasts the strike. Phoenix is written for the full arc — the strike that becomes the deep burn — because we tuned the *ratio* of the five, not merely their total. This is the whole lesson of blend design made edible: the SHU is the sum, but the shape is the proportion.

Which exposes a gap the bulk number hides, and it is worth seeing across the ten flagship cultivars. The table below lists our computational predictions — these are designed values, not certified measurements — and shows how the design scope grows from Ember's two active axes to Phoenix's full eight.

CULTIVAR	SKU	PREDICTED SHU	IFI-H	EC_{50} RATIO	AXES
Ember	SS-0087	866K	1,570	0.637	2/8
Cinder	SS-0088	1.05M	2,290	0.437	3/8
Kindle	SS-0089	1.25M	2,000	0.500	3/8
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 Cinder against Kindle, and the gap opens in plain sight. Kindle carries **more** total predicted heat — 1.25M SHU against 1.05M — yet reads **quieter** to the channel, IFI-H 2,000 against 2,290. More capsaicinoid mass, lower receptor loudness, because the blend is weighted toward the slower-binding congeners. SHU counted the mass and got it right; only the EC_{50} read the voice. That single inversion is the reason the index exists: a flavor is not how much is in the fruit, but how loudly it speaks to the protein that reads it.

And at the very edges of all this, almost unnoticed, **TRPM8** is brushed — menthol's cold channel, the only reader quieter than the reference at **800**, 0.80× its loudness. Not a note of its own so much as a faint cool frame thrown around the fire: the one axis we built soft on purpose, because a cool that competed with a 4,200 heat would be a contradiction, not a flavor.

Seconds three to ten — the savory floor, and the gain

Now the slow readers bloom, a full beat behind the channels — and the delay is not vagueness, it is architecture. Where TRPV1 and TRPA1 are doors that open on contact, the sweet, umami, kokumi, and bitter receptors are the mouths of an **eight-station cascade**, and every station costs time:

1. the ligand docks in the Venus-flytrap domain of a class-C GPCR;
2. the flytrap closes, turning a binding event into a change of shape;
3. the closure torques the cysteine-rich stalk down onto the seven-helix bundle;
4. the bundle switches on a heterotrimeric G-protein — α-gustducin with its βγ partners;
5. the βγ pair activates phospholipase C-β2;
6. PLCβ2 cleaves PIP_2 into IP_3, which empties the cell's calcium stores;
7. that calcium opens TRPM5, and the membrane finally depolarizes;
8. the depolarization opens CALHM1 channels, releasing ATP to the waiting afferent nerve.

Eight steps, each with an enzyme's dwell time, is why savor arrives a second or two behind fire. But the same cascade that costs time also *amplifies* — one caught molecule becomes thousands of ions — which is why these axes, once they land, read so full.

What lands, in order of loudness: glutamate and ribonucleotide reach the **TAS1R1·R3** flytraps, and umami blooms at **1,800** — half-activation at 55.6% of reference, a savory floor sliding in beneath the fire. The γ-glutamyl peptides engage **CaSR**, kokumi at **1,900**; kokumi has no flavor of its own, but as a positive allosteric modulator it turns up the gain on every other reader, and it is the receptor that fuses eight separate sensations into one continuous thing. Beneath that, a sugar-free sweetness surfaces at **TAS1R2·R3**, **1,100** — the gentlest metabotropic read, 1.10× the reference, just enough to round the heat. And the bitterness a plant this complex would otherwise carry is held firmly down: **TAS2R38** reads **2,200** on the zero-bitter axis, a 20-to-65% attenuation that keeps the whole chord clean rather than metallic. The fire now has a floor. It is savory, not merely hot.

Seconds ten to ninety — the long fade

Then the fire begins to let go, and it does so by its own hand. The calcium that floods through the open TRPV1 pores recruits calcineurin and depletes the membrane's PIP_2, and both feed back to lower the channel's sensitivity — **desensitization**, tachyphylaxis, the receptor's own mercy. The burn at second sixty is already gentler than at second ten, not because the capsaicin has left but because the channel has turned its own volume down. The TRPA1 flare is long resolved, its cysteine adducts reduced and renewed.

But one reader is still *climbing.* Astringency — the **salivary proline-rich proteins**, **IFI-A 1,400** — does not adapt, because it is not a receptor desensitizing to a ligand; it is a physical depletion. Tannins cross-link the proline-rich proteins and the mucins that keep the mouth slick, precipitating them out of the saliva, and each passing second strips more lubrication than the last. So the plum-dark drag arrives last and stays longest, framing the slow exit long after the fire has gone kind. This is the axis to be honest about: its IFI is a model — a half-saturation of protein binding standing in for a half-activation of a channel — because there is no single receptor here to anchor it. It is the one place on the card where the number is an analogy rather than a direct read. We keep it because the mouth keeps it; we flag it because precision means saying exactly where the model bends.

And the kokumi continuity ensures the whole thing recedes as one fading chord rather than eight notes snapping off at different times. The finish is as designed as the attack.



A fragment of Phoenix on the microscope stage — every axis on the scorecard is a half-activation concentration read here first, long before it is ever felt.

The body answers a lie

All the while, beneath the eight reads, the rest of the body is answering a false alarm — because TRPV1's fire is a lie the nervous system believes completely. The activated nerve endings release substance P and calcitonin-gene-related peptide, which dilate the small vessels of the skin; the face flushes, the brow and scalp break into a gustatory sweat, the eyes and nose stream, and the heart quickens on a thread of sympathetic adrenaline. The body mounts a full defense against a heat, and an injury, that were never there. And then — the reason chili is craved and not merely endured — the brain, finding no true damage behind the pain, floods the system with its own endorphins and endocannabinoids to meet an injury that never came. What is left is the soft, warm, faintly giddy afterglow any chili-lover knows: a whole-body reward issued in answer to a purely molecular deception. It is the clearest proof that a receptor read is never only a flavor. Eight proteins change shape on a tongue, and a whole person — sweat, pulse, mood, and all — answers. We design heat as a pleasure given to the body with intention, never as an assault, and this is where the two finally part: the pleasure is the body's own reply to a message it was safe to believe.

What the number can, and cannot, say

Honesty is a form of respect for the reader, so here is where the card stops. The IFI measures **loudness — a half-activation concentration — and nothing else.** It does not measure how much you enjoy the burn; potency and pleasure are different sciences. It does not measure the dose actually delivered in a given bite, which depends on how much fruit, how much fat to carry the capsaicin, what temperature, what pH — capsaicin is lipophilic, so a spoon of cream genuinely quiets it by partitioning the molecule away from the channel. It does not account for the reader on the far side: the bitter axis in particular rides on TAS2R38 genotype, and a taster and a non-taster will read the very same fruit differently by nature, not by error. The docking energies that write the blend are honest as **rank order and ratio** — exp(ΔΔG ÷ RT) is trustworthy, and it is the reason we quote fold-differences rather than absolute affinities — but a −7.20 is a reliable *tighter than −6.39*, not a measured dissociation constant, and we will not pretend otherwise. And the card is the **fresh** read: desensitization means that by second sixty the living 4,200 is already something lower, and the index records the arrival, not the whole fade. Phoenix itself is an extreme — few people will, or should, eat anything this loud, and we built it to prove a point, not to be anyone's dinner.

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

One bite, eight readers

> One bite. Eight readers. A pungent flare, a fire with a savory floor, a sweetness held clean, a cool at the edges, and a plum-dark drag that outlasts them all — *not eight sensations, but one, which is the only kind the mouth has ever known how to feel.*

That scorecard is the whole method on one plate. Each number is a receptor read at a computed loudness; the eight together are a flavor composed, not stumbled into; and every one of them traces back through a docking score, a cascade of stations, and an edited enzyme to a lock we learned to read before we ever tried to write a word into it. What our model demonstrates at a predicted 2.95 million Scoville is true at forty: the quiet shishito and the screaming superhot are read by the very same eight proteins, and only the volume differs. The volume is ours to write — carefully, cleanly, with the whole arithmetic shown, on a fruit grown to carry the fullest flavor, and the fullest nourishment, it can hold. The method is the point, and the method is reading.

Chapter 16 · The Universal Reader — *one tongue, every plant.*

The eight receptors do not know what plant they are reading. TRPV1 cannot tell a pepper from a cultivar that did not exist until we grew it; the sweet heterodimer closes on a steviol glycoside exactly as it closes on cane sugar; the calcium-sensing receptor reads a garlic's γ-glutamyl peptide and a yeast's glutathione with the same patient clasp. The reader is universal. The plant is not. Everything in this book has been circling that one asymmetry, and it is time to name it plainly: because the locks are identical in every mouth and wholly indifferent to where their keys were grown, a method that learns to read the locks transfers across the entire plant kingdom. The work was never really about peppers, or tomatoes, or any single species. It is about the receptors — and about everything green that can be taught to speak to them.



The founders in the greenhouse at the close of a day — the work was built by a family; the computing was one pair of hands, the table and the care belonged to all of us.

The lock is in your genome; the key is in the plant

To call a receptor *universal* is to make a precise claim about where it is written. TRPV1, the sweet and umami and kokumi dimers, the bitter receptor, TRPM8, TRPA1 — each is a protein encoded in the human genome, folded by the same ribosomes in every person at a table, studded into the membranes of the same sensory cells. You did not acquire your TRPV1 by eating chilis; you were born carrying the gene, and so was every eater who came before chilis and every one who will never taste one. The lock is yours, inherited, fixed. It is the most conserved half of the entire act of flavor.

The key is the other half, and the key is the plant's. A flavor is a molecule speaking to a receptor — capsaicin to TRPV1, a steviol glycoside to the sweet flytrap, allyl isothiocyanate to TRPA1 — and the molecule is manufactured by the organism, in its cells, by enzymes its genome specifies. This is the whole leverage of the platform in a single sentence: *we cannot and would never rewrite the lock in a person's mouth, but we can write the key into a plant.* The receptor is held constant on purpose, because it is the shared standard against which everything is measured; the molecule is the only thing we touch.

Hold that division steadily and the apparent audacity of forty-two engineered species becomes something quieter and more rigorous. We are not inventing new senses. We are composing new sentences in an alphabet of eight letters that the human body already knows how to read, and that it reads the same way whether the sentence was spoken by a pepper that has existed for ten thousand years or by a tomato we taught last spring to carry a real umami load. The universality of the reader is not a poetic flourish. It is the experimental control that makes every number in this book comparable to every other.

Two families of reader, and one that reads with its hands

The eight do not all read the same way. Walk the architecture and they sort cleanly into two protein families plus a single, honest exception — and that sorting is itself the deepest biology in the book, because it tells you *how* a universal lock is built.

AXIS · CODE	RECEPTOR	MOLECULAR CLASS	HOW THE READING HAPPENS
Heat · IFI-H	TRPV1	tetrameric cation channel (TRP)	capsaicin lowers the heat gate; Na^+/Ca^{2+} pour in
Cool-Burn · IFI-C	TRPM8	tetrameric cation channel (TRP)	menthol shifts the cold gate open
Pungency · IFI-P	TRPA1	tetrameric cation channel (TRP)	reactive molecules bond covalently to cysteines
Sweetness · IFI-S	TAS1R2·TAS1R3	class C GPCR heterodimer	sugar closes the Venus flytrap
Umami · IFI-U	mGluR1 / TAS1R1·TAS1R3	class C GPCR	glutamate closes the flytrap; 5′-nucleotides lock it
Kokumi · IFI-K	CaSR	class C GPCR	γ-glutamyl peptides ease the flytrap shut
Zero-Bitter · IFI-B	TAS2R38	class A GPCR	a bitter agonist fills a pocket — we remove it
Astringency · IFI-A	salivary PRPs	secreted proline-rich proteins	tannins precipitate the protein; felt as grip

The **first family is the ion channels** — TRPV1, TRPM8, TRPA1, the three TRP proteins that carry heat, cool-burn, and pungency. Each is a tetramer: four identical subunits arranged around a central pore, every subunit threading the membrane six times, the four pore-loops meeting in the middle to make a gate for sodium and calcium. These do not report to a taste cell at all. They sit in the free nerve endings of the trigeminal nerve, the somatosensory wiring of the face, and when they open they depolarize that nerve *directly*. The reading event is a conformational one: capsaicin slips into a pocket between

the third and fourth membrane helices of TRPV1 and lowers the temperature at which the channel's gate springs — the tissue is not warmed, the threshold is simply dropped below body heat, and the body reports a fire that is not there. Menthol does the mirror trick at TRPM8, shifting the cold gate open so the mouth feels a coolness with no fall in temperature. TRPA1 is the strange one, and the most instructive: it is not opened by a shape that fits, but by *chemistry*. The pungent molecules of mustard, garlic, and cinnamon are electrophiles; they form covalent bonds — actual, lasting chemical bonds — to specific cysteine residues on the channel's enormous ankyrin-repeat tail, and that bonded modification is what pulls the gate open. TRPA1 does not recognize its key. It is changed by it.

The **second family is the G-protein-coupled receptors** — sweet, umami, kokumi, and bitter — and they report not to the trigeminal nerve but to the taste bud. Three of them (sweet, umami, kokumi) are class C GPCRs, the elegant kind that carry a huge bilobed **Venus flytrap domain** out in the saliva on a stiff cysteine-rich neck; a ligand settles into the trap, the two lobes clasp shut, and that closure torques down through the neck into the seven-helix bundle buried in the membrane. The bitter receptor, TAS2R38, is a smaller class A GPCR with its binding pocket inside the membrane bundle itself. However they catch their molecule, the four of them then hand the signal to the *same* downstream wire — a shared transduction spine that is as universal as the receptors feeding it: the taste G-protein **gustducin** switches on **phospholipase C-β2**, which splits the membrane lipid PIP_2 into **IP_3**; IP_3 dumps calcium out of the cell's internal stores; the rising calcium opens the **TRPM5** channel; the cell depolarizes and releases ATP through the **CALHM1/3** channel onto the gustatory nerve. Four different front doors, one hallway. Even the diversity of taste collapses onto a single conserved pathway before it ever reaches the brain.

And then the **one that reads with its hands**: astringency. There is, strictly, no astringency *receptor*. The molecules responsible are tannins — the polyphenols of tea, grape, pomegranate — and what they bind is not a protein in a membrane but the **proline-rich proteins** your salivary glands secrete by the milligram to keep the mouth slick. Tannins cross-link and precipitate those proteins; the saliva loses its lubrication; and the puckering drag you feel is largely *mechanical*, the tongue registering friction through the same touch sense that tells it a surface is rough. Of the eight readers, seven are conformational events in a single protein and the eighth is a colloidal one across a crowd of them. We keep astringency on the index anyway, scored by the same half-maximal logic, because the sensation is real, reproducible, and designable — but we never pretend it is a lock-and-key when it is a loss of slip. Honesty about that difference is part of what makes the other seven trustworthy.

Why a universal lock can be scored

A fixed lock is the one thing that makes a *number* possible. If every mouth carried a different TRPV1, a heat score would mean nothing from one eater to the next; because the receptor is shared, a single

curve describes how loudly a molecule speaks to it, and that curve can be read off a bench anywhere in the world and believed everywhere.

The curve is a dose-response, and it is not a straight line but a sigmoid. We fit it with the Hill relation,

$$\text{response} = R_{max} \times [L]^n \div (EC_{50}^n + [L]^n),$$

where [L] is the concentration of the stimulus, R_max the receptor's full response, and *n* the Hill coefficient that sets the steepness. The one number that summarizes the whole curve is **EC_{50}**, the concentration at which the receptor reaches half of R_max — the half-activation point. A molecule that drives its receptor hard reaches half-max at a *low* concentration; a feeble one needs a flood. EC_{50} is therefore an inverse measure of loudness, and the Iyengar Flavor Index simply turns it right-side up:

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

Anchor a reference stimulus — by construction it scores exactly 1,000, since the formula run against itself is one times a thousand — and every candidate is read against it. More potent than the reference means a smaller EC_{50} means a score above 1,000; weaker means below. The scale is receptor-anchored, not food-anchored, and that is precisely why it travels: the denominator is a property of the key, the numerator is a fixed property of the lock, and the lock is the same in Chennai and in a chamber in the greenhouse and in the reader's own mouth.

The IFI is a stated convention — each reference constant defines IFI = 1,000 on that axis, in the same spirit as the Scoville scale's dilution-to-threshold definition. On the heat and pungency axes, the reference stimulus has a well-characterized EC_{50} from the published literature (IFI-H anchors to capsaicin/TRPV1; IFI-P anchors to allyl isothiocyanate/TRPA1 at a convention constant of 2.1 µM). The remaining six axes are anchored to a stated reference compound by convention. Because each axis carries its own reference point, scores on different axes represent loudness relative to that axis's reference and are not directly comparable across axes.

The arithmetic runs as cleanly backward as forward, which is the test of a real measurement. Take **Phoenix**, the one flagship cultivar designed to register on all eight readers at once — the full-spectrum proof that the axes are genuinely separable and genuinely universal:

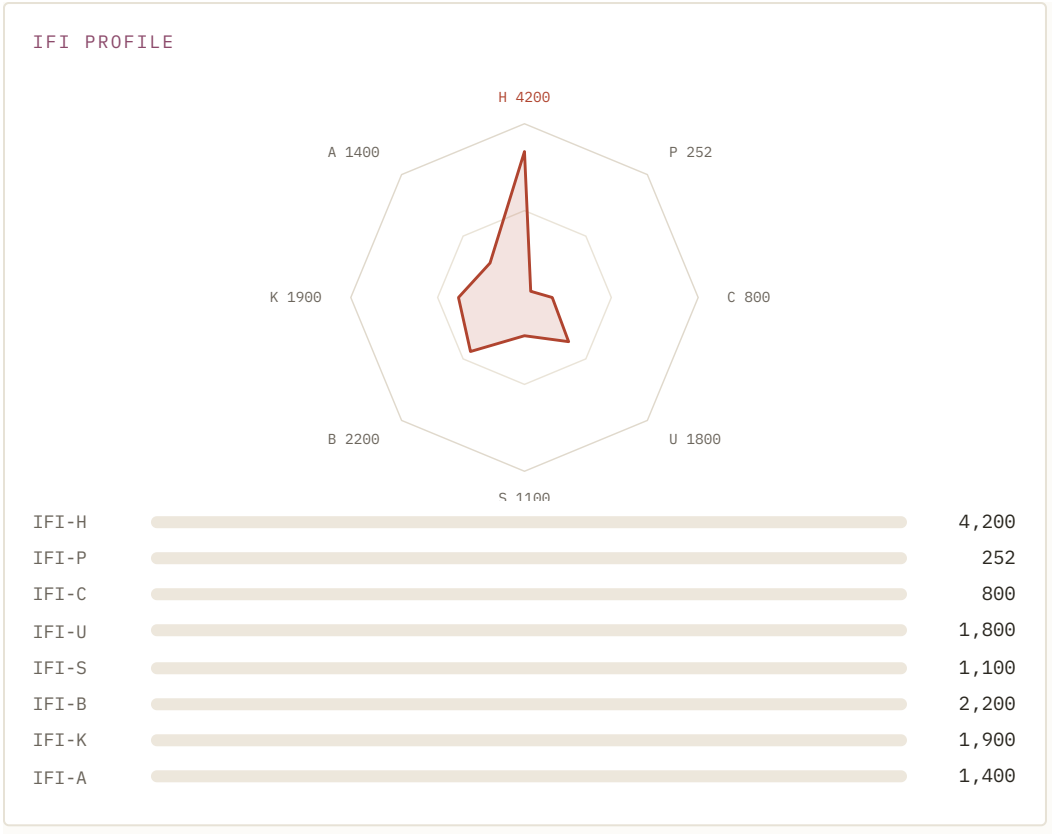

Its heat reads **IFI-H 4,200**. Invert the formula — EC_{50}(candidate) = 1,000 × EC_{50}(reference) ÷ IFI — and 4,200 means the capsaicinoid load in that tissue half-activates TRPV1 at **1 ⁄ 4.2 ≈ 0.238** of the reference's concentration. Phoenix is designed to speak to the receptor 4.2 times louder than the native wild-type anchor: our model predicts the fruit reaches half of TRPV1's full fire at under a quarter of the reference material. Do the same on the umami axis — IFI-U 1,800 — and the savory ligand is designed at 1.8× the reference's loudness; on kokumi, 1,900 puts the richness at 1.9×; on zero-bitter, 2,200. One cultivar, eight independent reads, every one a ratio against that axis's stated reference lock.

One ceiling is worth stating because the scorecards enforce it: each axis is normalized to a maximum of **5,000**. Phoenix's heat at 4,200 sits at 84% of that ceiling — high on the dial, with headroom deliberately left, so that the rare natural or engineered ligand that overshoots can still be reported as a true number in the text even when its bar on the radar is capped. We report the real value and let the chart show the cap, because pretending the ceiling is the truth would be the one dishonesty the index cannot afford.

What the number knows that a Scoville cannot

Heat is the one axis the world already thought it could measure, and the comparison is the sharpest way to show what a receptor-anchored number buys you. The Scoville Heat Unit counts *molecules* — how much capsaicinoid a fruit holds, read today by HPLC against the old dilution convention of roughly **16 SHU per part per million**. IFI-H counts something else: how loudly that capsaicinoid load actually speaks to TRPV1. For a single pure molecule the two would track in lockstep. Across real cultivars, with real blends of capsaicin and dihydrocapsaicin set by the length and saturation of an acyl chain, they come apart — and the gap is information.

All SHU values in the table below are computational predictions from our flux and pathway models, not certified measurements.

CULTIVAR	SKU	PREDICTED SHU	IFI-H	PREDICTED SHU PER IFI-H POINT
Ember	SS-0087	866,000	1,570	552
Cinder	SS-0088	1,050,000	2,290	459
Kindle	SS-0089	1,250,000	2,000	625
Blaze	SS-0091	1,720,000	2,630	654
Wildfire	SS-0092	1,960,000	2,930	669
Scorch	SS-0093	2,210,000	3,220	686
Firestorm	SS-0094	2,470,000	3,490	708
Solara	SS-0095	2,720,000	3,980	683
Phoenix	SS-0096	~2,950,000	4,200	702

Read the last column, because it is where the two measurements disagree out loud. It is not constant and it is not even monotonic. **Cinder** spends only 459 predicted Scoville units to buy a single IFI-H point — a loud, efficient capsaicinoid blend in our model. **Kindle**, one design tier up at a predicted 1.25 million SHU, spends 625: our model predicts it carries **19% more capsaicinoid than Cinder** (1.25M ÷ 1.05M = 1.19) and yet reads **13% lower on the receptor** (2,000 against 2,290). More molecule, quieter receptor. A Scoville meter would rank Kindle the hotter fruit; TRPV1 would rank it the softer one, because the blend our model predicts is the less potent key. That divergence is invisible to the old number and native to the new one, and it is the entire reason to anchor on the lock.

The broad trend in that column tells a second story, and it is pure receptor physics. From Ember to Phoenix the predicted Scoville load rises **3.41×** (2,950,000 ÷ 866,000) while the receptor read rises only **2.68×** (4,200 ÷ 1,570). The loudness grows more slowly than the concentration — exactly what a sigmoid predicts as it flattens toward saturation. Push [L] far above EC_{50} and the Hill curve bends over; each additional molecule finds fewer unoccupied channels to open, and the receptor returns less for more. The rising cost-per-point, from 552 at the bottom toward ~700 near the top, is that flattening made arithmetic. It is why chasing Scoville units past a certain point is a race against diminishing receptor return, and why Phoenix's 4,200 — 84% of the normalized ceiling — is already deep in the compressed zone where the lock is nearly spoken out.

A last conversion grounds the scale in something you could weigh. Our model predicts Phoenix's ~2.95 million SHU; dividing by 16 SHU per ppm gives roughly **184,000 ppm** of capsaicinoids — approximately **18% of the fruit's dry weight** in our model, against the ceiling of pure capsaicin near 16,000,000 SHU. The index turns that staggering predicted chemical load into a single receptor-anchored word, *4,200*, and the word means the same thing on any bench that keeps the same lock.

Designing the key, then installing it in any genome

If the reader is universal, the method that serves it must be too — and here the biology is generous. The molecule TRPV1 wants is specified by the receptor, and most of the road to that molecule is already paved in every plant on earth. Capsaicin is a child of the **phenylpropanoid pathway**, and its early enzymes are not pepper inventions:

ENZYME	ROLE	WHERE IT RUNS
PAL	phenylalanine → cinnamate	every vascular plant
C4H	→ *p*-coumarate	every vascular plant
4CL	→ *p*-coumaroyl-CoA	every vascular plant
HCT	acyl transfer	every vascular plant
COMT	O-methylation toward the vanillyl branch	every vascular plant
pAMT	amine transfer → vanillylamine	pepper-specific
Pun1 / CS	capsaicin synthase — the final condensation	pepper-specific
KAS · BCAT	branched-chain acyl supply	pepper-specific

Five conserved enzymes carry any plant most of the way; only **four terminal enzymes** — pAMT, Pun1/CS, and the branched-chain supply of KAS and BCAT — are the pepper's own. To teach a new genome to make capsaicin is therefore not to build a factory from nothing but to install a short final stage onto machinery already running. That is the molecular meaning of *transfer*, and it is why a volcanic tomato or a sweet-heat strawberry is an engineering schedule rather than a fantasy.

The key itself we design before we grow it, and Phoenix is the proof of principle. Rather than chase heat by selecting phenotype and hoping the genotype followed — the method of every breeding program before — we scored every possible single-residue change in the Pun1 capsaicin synthase with a protein language model for its predicted effect on catalytic efficiency and fold stability. Our flux model predicts that flux control sits upstream in the phenylpropanoid pathway — at 4CL and HCT/CCoAOMT — not at the terminal Pun1/AT3 step; Pun1 is the capsaicin synthase that labels the product, and designing its stability and substrate affinity ensures the final condensation keeps pace with whatever flux the upstream enzymes deliver. We wrote the winning edits in with CRISPR. Two substitutions, never before combined in any pepper that ever grew, emerged:

EDIT	EFFECT	ΔΔG (KCAL/MOL)
S39L	binding affinity ↑ (N-terminus)	−2.1
L345G	steric relief at the active-site cleft	−1.8
C175S	*structural annotation: prevents a disulfide misfold — not included in the shipped construct*	*−0.9*

Treat the two shipped-construct free energies as additive and the aggregate is **ΔΔG ≈ −3.9 kcal/mol**. The thermodynamics of that number are worth carrying through, because they explain the single most important quality metric in the whole program. The equilibrium between a protein's competent fold and its failures scales as exp(−ΔΔG ⁄ RT); at body temperature RT is about 0.616 kcal/mol, so a −3.9 kcal/mol shift favors the working fold by roughly exp(3.9 ⁄ 0.616) = exp(6.33) ≈ **560-fold** — an order-of-magnitude statement, since affinity and stability terms are not strictly the same coin, but a true one in direction and scale. A fold that strongly favored is a fold that rarely slips, and that is why the engineered line holds its output to a batch-to-batch coefficient of variation of **≤7%** where the wild-

type scatters at **42%**, while our model predicts peak heat of approximately 2.95 million SHU. The structure is designed for consistency as much as for magnitude; the two are the same edit.

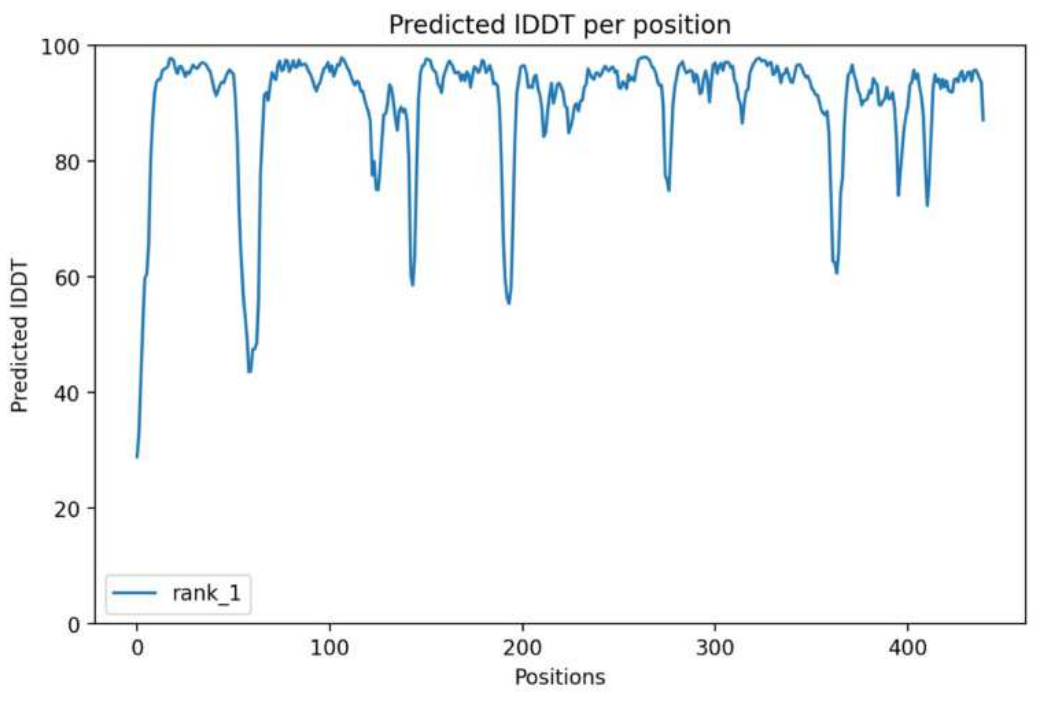

The Pun1 / AT3 model coloured by pLDDT confidence — the binding domain and active-site cleft, where S39L and L345G were written in, resolve at high confidence.

The error bars are a form of respect

A universal lock lets you measure; low variance lets you *resolve*. The difference between 7% and 42% is not cosmetic — it is the difference between a ladder and a smear, and the arithmetic says so exactly.

Take Phoenix's heat read of 4,200 and put a wild-type coefficient of variation on it. At **42%**, one standard deviation is ±1,764, and a single measurement spans **2,436 to 5,964**. That band swallows six rungs of the flagship ladder at once — Blaze, Wildfire, Scorch, Firestorm, Solara, and Phoenix itself all fall inside it. At that scatter the index literally cannot tell the top cultivar from the middle of the pack; the ranking is noise. Hold the same read to **≤7%** and one standard deviation is ±294, the band narrows to **3,906 to 4,494**, and the rungs step apart in clean order. The whole ten-rung heat ladder — 1,570, 2,290, 2,630, and the rest — is only legible because the error bars were driven small enough for adjacent rungs to stand clear. Precision is not a luxury on top of the measurement; below a threshold of variance there is no measurement at all.

There is a subtler piece of error mathematics in the full-spectrum read, and it cuts both ways. When you *divide* two noisy numbers — as the index does, candidate EC_{50} over reference EC_{50} — their relative errors add in quadrature, so two axes each measured at 7% yield a ratio no better than $\sqrt{(7^2 + 7^2)} \approx$ **9.9%**. But when you *average* independent axes, noise shrinks: the mean of Phoenix's eight reads, each at ≤7% and each on its own receptor, carries a composite coefficient near $7\% / \sqrt{8} \approx$ **2.5%**. A full-spectrum profile is therefore known, as a whole, more tightly than any one of its axes — the eight independent locks each vote, and the vote is steadier than any single voice. We report every axis to the precision it has actually earned and no further, and we say the width out loud, because a

number offered without its uncertainty is not a kindness to the person who has to design with it. It is a small deception. The error bar is the respect.

Where "universal" bends

A thesis is only trustworthy where it admits its own strain, and *universal reader* strains in four honest places. None of them undoes the method. All of them belong in any account that means to be believed.

The first is that the reader has **variants**. The lock is shared, but not byte-for-byte identical across people. The clearest case is the bitter receptor TAS2R38, which comes in two common forms separated at three amino-acid positions — the taster haplotype (proline, alanine, valine) and the non-taster (alanine, valine, isoleucine). To a taster, the thiourea compound PTC is a wall of bitterness; to a non-taster it is nearly nothing. The receptor is universal as *architecture* and polymorphic in *tuning*, and an honest index anchors to a defined reference genotype and says which one, rather than pretending every mouth reads the bitter axis alike. We design for the shared physiology and report against a stated anchor; we do not claim a single number captures every eater.

The second is **astringency**, and it is the exception the architecture already warned of. Because the sensation is tannins precipitating salivary proteins and the tongue feeling the loss of slip, there is no clean EC_{50} of a single lock to invert — the reading is a colloidal, mechanical event across a population of proteins, modulated by how much saliva a given mouth makes and how fast it recovers. We score it by half-maximal *perceived* astringency against a reference tannin, which is a real and reproducible measurement, but it is a measurement of an outcome, not of a channel opening. The index carries the astringent axis with that asterisk permanently attached.

The third is that **not every axis is a loudness**. Most readers we score for how hard a molecule speaks; the zero-bitter axis we score for how *quietly* it does. Bitterness is usually the signal we are removing — attenuating glucosinolates in a brassica, softening the catechins of a coffee by 20 to 65% — so a "high" zero-bitter read is the measure of a silence achieved, mapped onto the same 0-to-5,000 frame by its percent of suppression rather than by a potency. The formula still governs the bench work, but the axis points the other way, and it would be sleight of hand not to say so.

The fourth is the **matrix**. An EC_{50} measured on a purified ligand against a purified receptor is clean; a mouthful of food is neither. Kokumi is a gain on *other* tastes and barely registers alone; umami and sweetness potentiate each other; a tannin reads differently in water than in wine. We anchor the receptor-level number to keep it reproducible, and then we are candid that the *eating* happens in a matrix, that the reported score is for a stated base, and that the index measures receptor loudness — not liking, not satiety, and never health. A plant can read loud on all eight locks and still not be the one you want at the table. The number is a true instrument pointed at one real thing, and it is careful about everything it is not.

The map that is mostly empty

Lay the organisms against the axes and the shape of the work still to be done becomes visible as a grid. Forty-two host species down one side; eight receptors across the top; a book — or a cultivar, or an empty cell waiting for one — wherever a row meets a column.

READER (AXIS)	PLANTS WE HAVE TUNED TO SPEAK IT
TRPV1 · heat	pepper (1,060 constructs), habanero, the ten superhot flagships
TAS1R2·R3 · sweetness	stevia (910), monk fruit, sugarcane, the sweet-heat strawberries
TAS1R1·R3 · umami	tomato (2,040), soybean, shiitake, mushroom, koji mold, yeast
TAS2R38 · zero-bitter	coffee, cacao, broccoli, kale, cabbage, hops, potato
CaSR · kokumi	garlic, onion, yeast, vanilla
TRPM8 · cool-burn	peppermint (2,360), spearmint, eucalyptus, camphor, lavender
TRPA1 · pungency	black & white mustard (2,790), garlic, ginger, cinnamon, watercress, radish, thyme
salivary PRPs · astringency	tea (280), grape (1,610), pomegranate, blueberry, raspberry, sorghum

Forty-two organisms and eight axes make **336 intersections** before a single diet, cuisine, or frontier concept is added — and most of those cells are still empty. Spread the roughly **19,000 product design packages** across them and the average cell holds only about 56; spread them across the organisms and each species averages near 452, which means the densely-worked rows (tomato's 2,040, mustard's 2,790, peppermint's 2,360) are balanced by rows barely begun. Of **78,800 catalogued SKUs**, about **1,000 are live and growing today** — a little over one percent of what has been named has been planted. The map is not a record of completion. It is a record of how much remains.

Beyond the forty-two lies the roadmap: **64 concept organisms in five tiers**, each a plant whose genome could, in principle, be taught a word its wild form never carried. The near tiers are Solanaceae we half-know how to reach — a *Volcanic* tomato with a true TRPV1 read (Ember, Magma, Pyroclast, Caldera, Krakatoa), a *Dark Matter* eggplant, a *Seismic* potato, an *Acid* tomatillo. The far tiers are frontier synthesis measured in years, not months — a *Sweet Heat* strawberry holding the sweet flytrap and the heat channel in one fruit (Warm Kiss, Sugar Burn, Candied Fang, Devil's Dessert), a *Contradiction* cucumber, a *Storm* corn, a *Dynasty* rice. None of it is science fiction. Each is the same inversion the whole platform is built on — read the receptor, design the molecule, grow the plant — pointed at a new genome whose phenylpropanoid spine is already running and whose terminal enzymes are the only real work.



Tomato in cross-section — the flagship proof that the method leaves the pepper; the same eight locks, read in a fruit that never carried heat until we wrote it a key.

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

And beneath every number, the quiet thing that started us, which it is finally time to name. We did not learn to read the eight receptors in order to make food merely louder. The same plants we tune for flavor are plants we grow to hold the most they can hold, so that the body reading them at the table draws the fullest of what a plant can carry: the savory depth that lets a vegetable satisfy on its own terms, the attenuated bitterness that lets a worthwhile bitter food be eaten whole, the whole-plant care that treats flavor and nourishment as one act of design rather than a trade between them. We publish as we go — the method, the structures, the scores — released openly under a Creative Commons license across 21 peer-reviewed papers, with 3 patents filed, so that anyone holding the same instruments can check the same locks. We will never turn the nourishment into a claim. It is an intention, and it is the reason the eight readers were worth learning in the first place: not as an end, but as the alphabet in which a family learned to write food worth growing.

The tongue does not taste the world. It reads it — eight proteins, eight words, read billions of times a second in every mouth that ever softened at a table. We have spent this book learning to read them too.

> The lock is the same in every mouth. *That is why there is no last plant* — only the next receptor-read we have not yet taught something green to say.