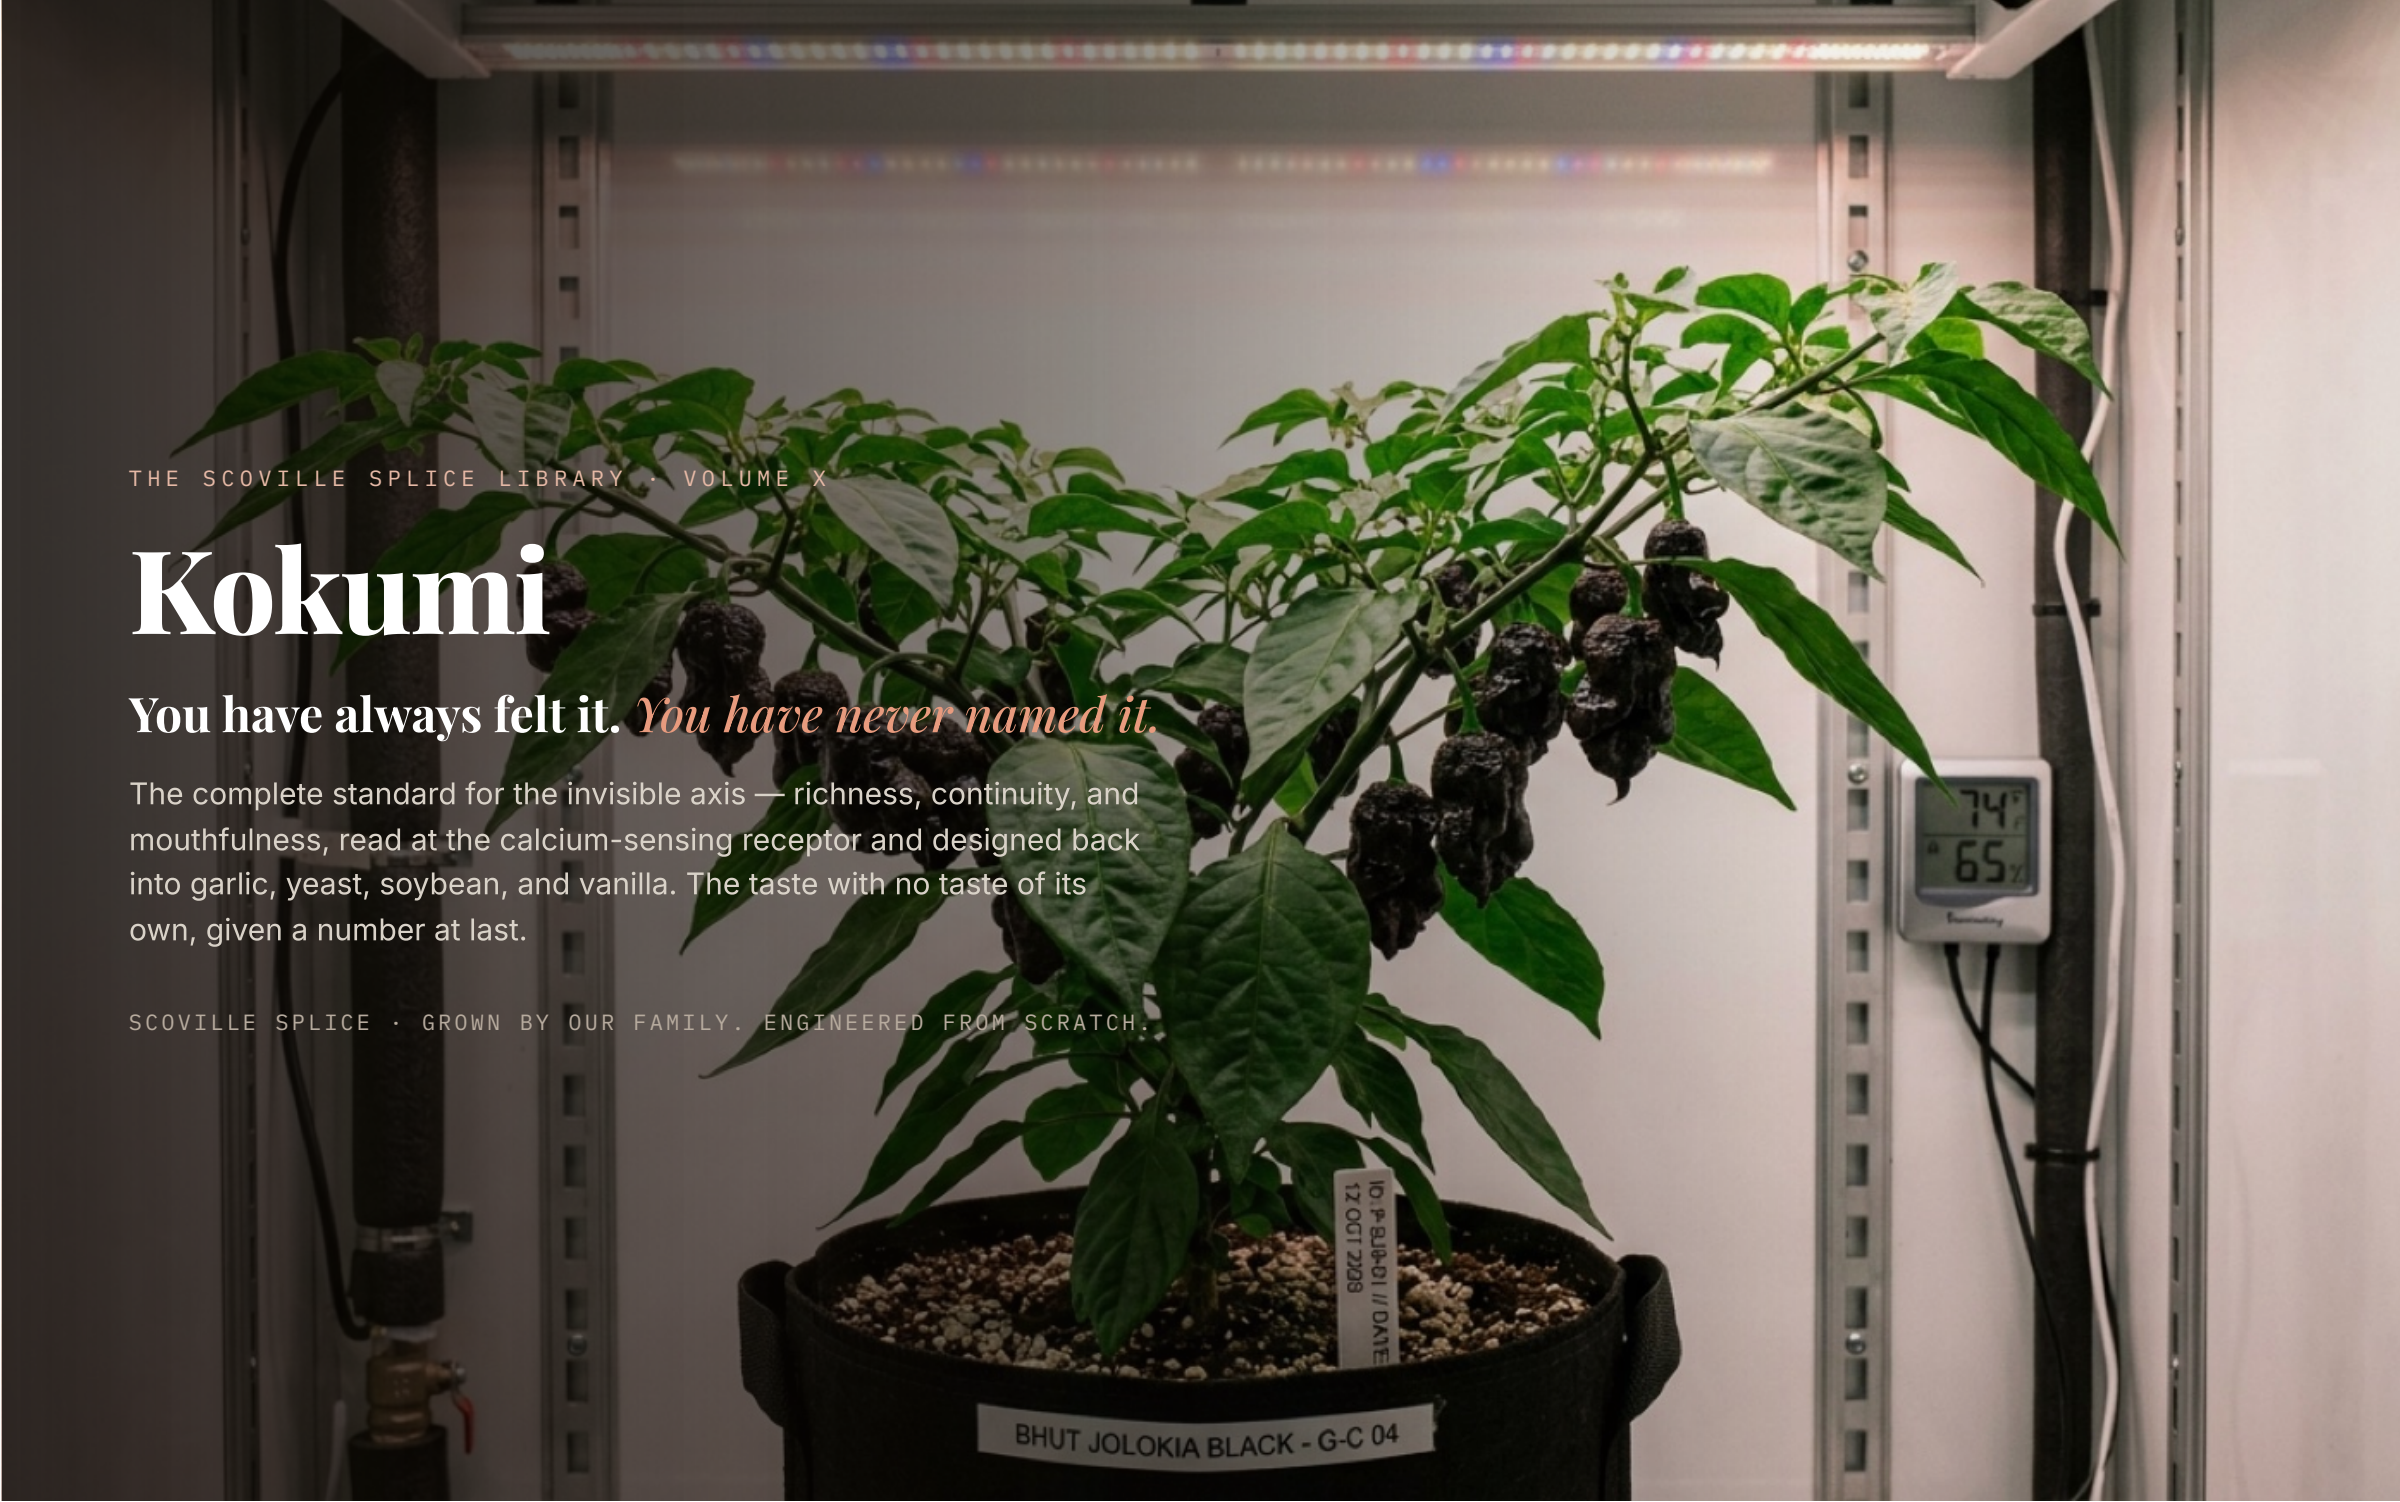

THE SCOVILLE SPLICE LIBRARY · VOLUME X

Kokumi

You have always felt it. *You have never named it.*

The complete standard for the invisible axis — richness, continuity, and mouthfulness, read at the calcium-sensing receptor and designed back into garlic, yeast, soybean, and vanilla. The taste with no taste of its own, given a number at last.

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

Contents

01 Opening essay · The taste with no taste

02 Chapter 1 · The three sensations — mouthfulness, continuity, complexity.

03 Chapter 2 · The calcium-sensing receptor — the biology of more.

04 Chapter 3 · Scoring a modifier — the mathematics of continuity.

05 Chapter 4 · The γ-glutamyl family — the peptides that speak to calcium.

06 Chapter 5 · The γ-glutamyl cycle — glutamate, cysteine, glycine, ATP.

07 Chapter 6 · Designing the pathway — flux, feedback, and the AI that reads proteins.

08 Chapter 7 · Garlic — the savory foundation.

09 Chapter 8 · The wider alliums: onion, leek, scallion — the gentler route.

10 Chapter 9 · Yeast & koji — the richness factories.

11 Chapter 10 · Soybean & the ferments — miso, shoyu, natto.

12 Chapter 11 · Vanilla — the kokumi of sweetness.

13 Chapter 12 · The naturally rich — where kokumi gathers on its own.

14 Chapter 13 · Kokumi × umami — the most-requested combination.

15 Chapter 14 · The amplifier — salt, sugar, and fat, reduced.

16 Chapter 15 · The kitchen of continuity — a short book of rich dishes.

17 Chapter 16 · Reproducible richness — CV ≤ 7%, GRAS, and the honest label.

18 Chapter 17 · The quiet axis — why richness is a kind of care.

Opening essay · The taste with no taste

There is a moment in a good broth that nobody has ever been able to name. It is not the salt, which you can point to; not the savory pull of glutamate, which has had a name for a hundred years; not sweetness or warmth or the green brightness of an herb. It arrives a half-second after everything else, and it is the thing that makes the spoonful feel *held* — rounded at the edges, full in the middle, unwilling to leave the mouth cleanly. Our grandmother called it, in the kitchen where this family first learned to cook, the part of the soup that *stays with you*. She was describing, without the vocabulary for it, the activation of a receptor she would never hear of, by a class of peptides that have no flavor of their own.

That is kokumi — 濃味, from the Japanese for *rich taste*, though "taste" is almost the wrong word, because kokumi is not one. Put a kokumi compound on your tongue alone, in water, and you will taste essentially nothing: a faint mineral whisper and then silence. It is only in the presence of other flavors — sweet, salty, savory, sour — that it wakes, and what it does then is not to add a flavor of its own but to **modify** the ones already there. It thickens them. It lengthens them. It makes a thin thing feel complete. For ten thousand years cooks have chased it by instinct — aging the cheese longer, roasting the garlic darker, letting the stock go another hour, reaching for the fish sauce, the yeast, the dark miso at the back of the shelf — and never once could any of them say *how much*.



The controlled chamber where our kokumi lines are held at constant humidity and light — the invisible axis, grown in the open where it can finally be measured.

The Iyengar Flavor Index gives it a number. Every axis in the Index is anchored to the half-maximal activation of a single receptor — not how much of a compound a food contains, but how *loudly* that compound speaks to the tissue that reads it.

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

For heat, the receptor is TRPV1 and the molecule is capsaicin, and the loudness is a fire that is not there. For kokumi the receptor is **CaSR, the calcium-sensing receptor**, and the molecules are the **γ-glutamyl peptides**, and the loudness is something stranger and

quieter: a signal that does not say *sweet* or *savory* but *more* — more body, more continuity, more of whatever is already true. What follows is the account of that axis: how the receptor works, which molecules speak to it and in which words, the enzymes that build those molecules, the six organisms we have taught to carry it, and the kitchen it unlocks. We did not set out to add anything to food. The other axes are about *presence* — more heat, more sweetness, more savor. Kokumi is about *integration*: the molecule that makes a dish stop sounding like a list of ingredients and start sounding like one thing. It is the most generous of the eight axes and the least vain, because its whole art is to make everything around it taste more like itself. We think that is a kind of care, and we built the number so we could design the care on purpose.

The receptor that reads richness

Begin with the tissue, because the tissue is where the whole story is decided. CaSR is a **class C G-protein-coupled receptor**, and class C is the architecture of the subtle senses — the same family that holds the metabotropic glutamate receptors that give savor its depth and the sweet-and-umami T1R heterodimers. Class C receptors do not work like a simple switch. Each one carries, folded out ahead of its membrane-spanning bundle, a large clamshell of protein called the **Venus flytrap domain**, and the receptor functions as a dimer — two flytraps hinged side by side, so that the closing of one lobe is felt by the other. This is not ornament. It is the physical reason kokumi behaves the way it does on the tongue: a receptor built to integrate several weak signals into one graded answer, rather than to report a single molecule's yes or no.

The flytrap has, in effect, two kinds of place where a molecule can land. There is the **orthosteric site**, the deep one, which reads free calcium ion — extracellular Ca^{2+} is CaSR's native agonist, the thing it exists to sense. And there are **allosteric sites** along the cleft and its lip, where L-amino acids and, crucially for us, the **γ-glutamyl peptides** bind. A molecule in the allosteric site does not open the receptor by itself. It makes the orthosteric site *easier to satisfy* — it stabilizes the closed, active conformation of the clamshell, so that the calcium already present in the mouth now reads as a louder command. When the flytrap closes, the seven-helix bundle beneath it rearranges, couples to its G-protein — principally **Gq/11** — and the Gq α-subunit switches on **phospholipase C-β**. PLCβ cleaves the membrane lipid PIP_2 into two second messengers: **IP_3**, which empties calcium from the cell's internal stores, and **DAG**, which recruits protein kinase C. The cell that was quiet fills with calcium from the inside, and *that* is the signal the brain finally receives.

Now the paradox resolves. CaSR sits on the taste-bud cell beside the machinery for sweet and umami and salt, but it does not run a transduction line the brain has learned to label as a basic taste. There is no dedicated "kokumi" wire climbing to the cortex. What CaSR does instead is raise the **gain** on its neighbors — it is a coincidence detector, a volume knob wired into the circuit of the other tastes. Give it a γ-glutamyl peptide and nothing else, in plain

water, and there is no companion signal to amplify; the knob turns on nothing, and you taste the silence our grandmother never had a word for. Give it the same peptide in a broth already carrying salt and glutamate and a little sweetness, and every one of those signals comes back larger, rounder, and — because of a bond we will come to — slower to fade. The taste with no taste is not a mystery. It is modulation, and modulation is a thing you can write as an equation.

Why the signal says "more," not "sweet"

To measure loudness you need the shape of the curve that loudness lives on. When a receptor reads an agonist, its response does not climb in a straight line; it follows a sigmoid described by the **Hill equation**, which in fractional form is

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

where R is the response as a fraction of the receptor's maximum, [A] is the agonist concentration, n is the Hill coefficient that measures cooperativity, and EC_{50} is the concentration that drives the response to exactly half of full. Set [A] equal to EC_{50} and the arithmetic is clean: $R = EC_{50}^n / (EC_{50}^n + EC_{50}^n) = ½$. That single point — the half-maximal concentration — is the whole anchor of the Index, because it is the one place on a sigmoid that is sharp, reproducible, and independent of where you happen to have drawn the top of the curve. For calcium at CaSR the curve is **cooperative**, with a Hill coefficient near 3; the two flytraps of the dimer help each other close, so the receptor flips from quiet to loud across a narrow band of concentration. That steepness is exactly what makes a small allosteric nudge produce a large perceptual change.

Here is the nudge, worked. A **positive allosteric modulator** lowers the orthosteric EC_{50} by some factor α — it shifts the whole calcium curve to the left. Suppose the mouth holds calcium fixed at the receptor's *unmodulated* half-point, so that without any kokumi peptide the response sits at R = 0.50. Add a peptide that shifts the EC_{50} down by a factor of **α = 2**. The response at that same, unchanged calcium becomes

$$R' = [Ca]^n / ((EC_{50}/\alpha)^n + [Ca]^n) = \alpha^n / (1 + \alpha^n)$$

With the cooperative n = 3, $R' = 2^3 / (1 + 2^3) = 8/9 =$ **0.89** — the signal climbs from half-strength to nearly full, a relative gain of about **78%**, and not one new molecule of calcium was added. Even in the conservative non-cooperative case, n = 1, the same two-fold shift lifts R from 0.50 to 2/3 = 0.67, a gain of **33%**. That is the mathematics of *more*: the peptide changes nothing about what is in the mouth and everything about how loudly the mouth reports it. When we say a sauce is "rounded" or a broth is "full," we are describing a dose–response curve that has slid left under its own stimulus, and a receptor that is now answering a quiet question with a loud voice.

The bond that ordinary peptidases cannot cut

The molecules that do this are the **γ-glutamyl peptides**, and the single most important fact about them is written in one Greek letter. In an ordinary peptide, glutamate links to its neighbor through its **α-carboxyl** — the backbone carbon — and the body is full of enzymes, the aminopeptidases and carboxypeptidases, built to recognize and snip exactly that α-linkage. A γ-glutamyl peptide links through the **γ-carboxyl** instead, the one out on the glutamate side chain. The resulting **isopeptide bond** is, to the mouth's ordinary peptidases, invisible: they reach for the α-geometry they know, find it occupied, and cannot cut. Only one class of enzyme, **γ-glutamyltransferase (GGT)**, at the membrane surface, cleaves it, and it does so slowly. This is the molecular reason kokumi *lasts*. A savory signal built on free glutamate is hydrolyzed and cleared in a breath; a signal carried on a γ-glutamyl peptide lingers, because the bond that holds it together is one the tongue's scissors were not made for. The continuity our grandmother tasted — the part of the soup that stays — is a hydrolysis rate.

The archetype is **glutathione**, γ-Glu-Cys-Gly, the little sulfur-bearing tripeptide every living cell builds by the ton. Its molecular weight is a number we will need in a moment: summing its atoms ($C_{10}H_{17}N_3O_6S$) gives **307.32 g/mol**. Glutathione is a modest CaSR agonist on its own, and it is where the field began, but the potent members of the family are its relatives — **γ-Glu-Val-Gly** (303.3 g/mol), one of the strongest kokumi peptides known, and, in the alliums we grow, the organosulfur peptides **γ-glutamyl-S-allyl-cysteine** and its kin, where a sulfur-linked allyl group rides the cysteine. Note the thread running through all of them: **sulfur**. Cysteine's thiol is the chemical handle on which much of kokumi's potency hangs, which is why the richest kokumi foods on earth — aged garlic, long-fermented soy, mature cheese, yeast extract — are precisely the ones that have had time to run their sulfur chemistry to completion. We did not invent this. We found the bond nature uses to make richness last, and we learned to build the peptide that carries it on purpose.

Building the molecule: the γ-glutamyl pathway, step by step

A plant cannot carry what it cannot synthesize, so the platform's work on this axis is, underneath the sensory language, the installation and tuning of a short enzymatic route. It begins upstream of flavor entirely, in **sulfate assimilation**: the plant draws sulfate from the soil, reduces it through APS and sulfite to sulfide, and condenses sulfide onto O-acetylserine to make **cysteine**. Cysteine is the sulfur gate; everything downstream inherits its atom. The first committed step is **γ-glutamylcysteine synthetase** (γ-GCS, also called glutamate–cysteine ligase), which joins glutamate to cysteine using the glutamate's γ-carboxyl and one molecule of ATP — and it is *here*, at this enzyme, that the uncuttable γ-bond is first forged. The product, **γ-glutamylcysteine** (250.3 g/mol), is already a competent kokumi molecule. **Glutathione synthetase** then adds glycine to make glutathione, and **GGT**, the same enzyme that cleaves these bonds at the tongue, runs in reverse as a transpeptidase to swap the terminal residue — handing the γ-glutamyl group to valine, to methionine, to a dozen acceptor amino acids and producing the whole diverse family of γ-Glu-X peptides from one conserved donor.



Garlic in cross-section — the allium's concentric rings are a sulfur factory; its kokumi rides on γ-glutamyl-S-allyl-cysteine, built from the same cysteine gate we tune in every kokumi line.

γ-GCS is the governor of the route — it is feedback-inhibited by glutathione and limited by cysteine supply — so lifting kokumi is rarely a matter of adding a single gene and walking away. It is a flux problem: relieve the feedback on γ-GCS, widen the sulfur supply feeding cysteine, and balance GGT so the pool shifts toward the potent γ-Glu-X peptides rather than pooling as plain glutathione. A genome-scale flux-balance model tells us, for each host, whether the true bottleneck is the ligase, the sulfur, or the transpeptidase, so we edit the step that is actually limiting rather than the one that is merely famous. One organism on this axis runs a second route as well: **vanilla**, whose signature is built through the **phenylpropanoid branch — PAL → C4H → 4CL → HCT → COMT** — the same chain of enzymes that opens in every vascular plant. PAL deaminates phenylalanine to launch the pathway that ends, many steps later, in vanillin; the kokumi of real vanilla is the marriage of that phenolic aromatic to a mouth-filling body the γ-glutamyl peptides supply. It is the most complex flavor in the kitchen, and the most faked, and the reason it cannot be faked well is that the imitation copies the aroma and forgets the body.

Giving richness a zero point: the IFI-K math

A number means nothing without a zero. The Index fixes one by naming a **reference molecule** for each axis and declaring its half-activation the pivot of the scale. For kokumi we anchor on glutathione and declare, by convention, EC_{50}(ref) = 1.00 mM — the fixed reference constant for this axis, defined as all the Index's anchors are defined: as a stated starting point, not a value derived from a published survey. Every kokumi score is then this molecule's potency measured against that zero:

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

Read the formula as a sentence and it says: *a candidate that is more potent than the reference needs less of itself to half-activate the receptor, so its EC_{50} is smaller, so its score is larger.* The score is linear in **potency** (that is, in $1/EC_{50}$): halve the EC_{50} and you double the number. A candidate exactly as potent as glutathione scores 1,000 by construction. Our full-spectrum flagship **Phoenix** carries a

predicted **IFI-K of 1,900**, and we can read that backward through the formula to learn what its kokumi chemistry is computed to be doing at the receptor:

> $1900 = 1000 \cdot (1.00\ \text{mM}) / EC_{50}(\text{Phoenix}) \Longrightarrow EC_{50}(\text{Phoenix}) = 1000 \cdot 1.00 / 1900 =$ **0.526 mM**

In that design, Phoenix's γ-glutamyl peptides are computed to half-activate CaSR at **526 µM**, against the reference's 1,000 µM — **1.9× more potent** than glutathione in our model, speaking the receptor's language at just over half the concentration. To make that a kitchen quantity, convert through the molecular weight: a 1.00 mM reference solution is 1.00×10^{-3} mol/L × 307.32 g/mol = **307 mg/L**, and the model's predicted 0.526 mM works out to $0.526 \times 10^{-3} \times 307.32$ = **162 mg/L**. The table below is nothing but the formula walked across the range of potencies, with glutathione's 1.00 mM as the pivot; it is, in effect, the calibration curve that turns a computed EC_{50} into a score.

EC_{50} AT CASR	POTENCY VS. REFERENCE	IFI-K	HOW IT READS
2.00 mM	0.5×	500	below the glutathione zero
1.00 mM	1.0×	1,000	the reference, by definition
0.526 mM	1.9×	1,900	**Phoenix · SS-0096**
0.400 mM	2.5×	2,500	a strong designed line
0.250 mM	4.0×	4,000	near the top of the practical band
0.200 mM	5.0×	5,000	the scorecard normalization ceiling

The ceiling is not arbitrary. The Index caps every axis at **5,000** for the radar scorecard, which on this axis corresponds to an EC_{50} of 200 µM — a five-fold gain over the reference. Beyond that point the receptor is being saturated rather than modulated, and, as any cook who has over-reduced a stock knows, more is no longer better.

Reading Phoenix: richness on a plant built for fire

Phoenix (SS-0096) is the only cultivar we have built that registers on all eight axes at once, and it is the clearest window onto how kokumi sits inside a complete flavor. It is, first, a plant of extraordinary heat — our model predicts approximately **2.95 million SHU**, IFI-H 4,200 — but what makes it the full-spectrum proof is that the fire is not alone. Here is its predicted profile, with each axis scored against its own receptor-anchored reference constant:

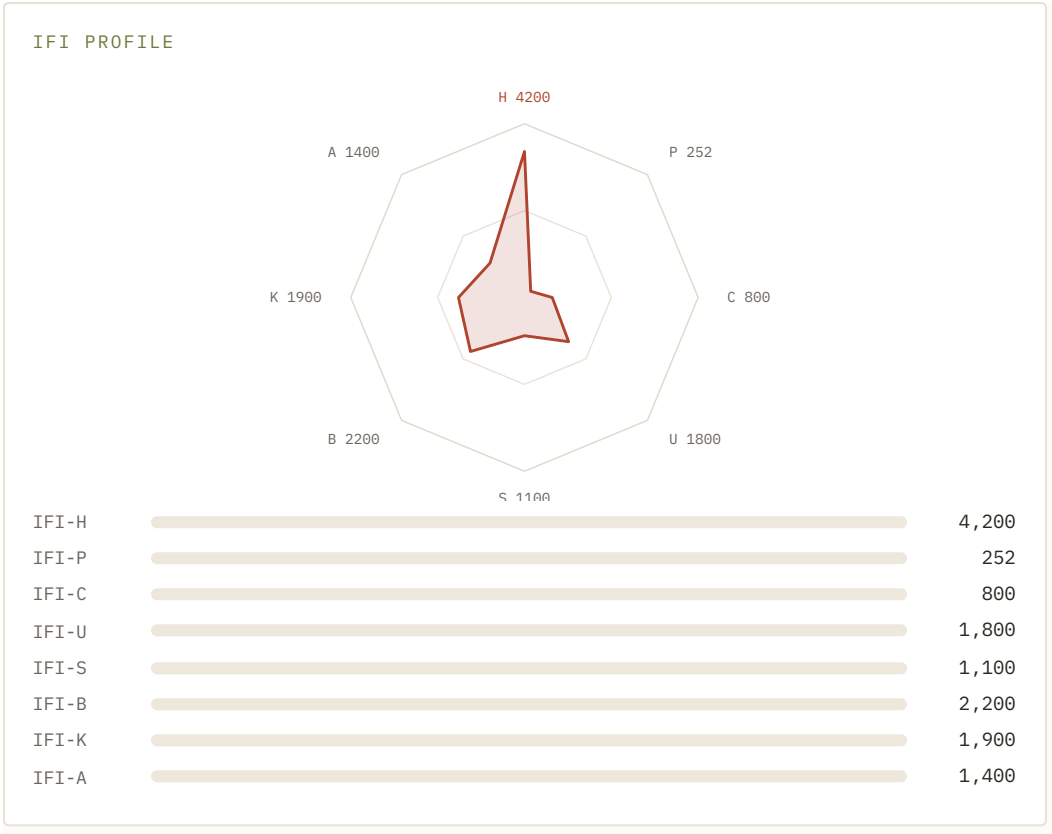

Look at the full predicted profile and the picture is breadth: all eight axes engaged. At **K = 1,900**, the kokumi design is 1.9× above the glutathione reference — a strong predicted signal on its own axis. Because each axis is scored against its own defined reference constant rather than a single unified scale, the numbers speak as independent intensities within their respective axes; together, they show a plant in which every channel is active and each is well above its baseline. That balance is deliberate. A plant engineered for extreme heat has, in every prior breeding program, tasted thin — heat without body, a signal with no floor under it. Building kokumi into Phoenix is what keeps nearly three million predicted Scoville units from reading as a chemical assault: the γ-glutamyl peptides lengthen and round the other seven axes so the heat arrives inside a complete mouth rather than on an empty one. The cultivar progression from Ember to Phoenix represents increasing computational design scope and axis activation — more channels engaged with each tier — and the number that turns a weapon into a flavor is the one you cannot taste by itself.

CULTIVAR	SKU	PREDICTED SHU	IFI-H	AXES EXPRESSED
Ember	SS-0087	866K	1,570	2/8
Blaze	SS-0091	1.72M	2,630	4/8
Firestorm	SS-0094	2.47M	3,490	6/8
Solara	SS-0095	2.72M	3,980	6/8
Phoenix	**SS-0096**	**2.95M**	**4,200**	**8/8 (full spectrum)**

The six organisms we taught to carry it

Kokumi is not a pepper trait. It is a peptide trait, and the peptide route transfers to any organism with the sulfur chemistry to feed it — which is why the richest kokumi work in the platform lives outside *Capsicum* entirely, in the six hosts that have always been the savory foundation of a kitchen. Together they carry **1,320 constructs** on this axis, and each one earns its place through a different corner of the same chemistry.

ORGANISM	CONSTRUCTS	PRIMARY AXES	KOKUMI CHEMISTRY	DIFFICULTY
Garlic	560	K / P	γ-glutamyl-S-allyl-cysteine; the organosulfur route	Moderate
Soybean	360	U / K	fermentation-built γ-glutamyl peptides; the miso/shoyu lineage	Frontier
Yeast	240	U / K	glutathione and γ-Glu-X from autolysis	Hard
Koji mold	80	U	the fermenting enzyme engine behind soy and sake	Hard
Onion	40	K	γ-glutamyl peptides of the sweet-sulfur allium	Moderate
Vanilla	40	K	phenylpropanoid aroma married to γ-glutamyl body	Hard

Garlic carries the most of them — 560 constructs — because garlic is, chemically, a kokumi factory that merely needs aging to run, and we tune it to run without the aging. Its γ-glutamyl-S-allyl-cysteines are the organosulfur peptides that give long-cured garlic its deep, mouth-coating finish, built from the same cysteine gate every line on this axis shares. **Yeast and soybean and koji** are the fermentation triad — the lineage of miso, shoyu, and every long-brewed savory depth — where microbial autolysis and the koji mold's enzymes together liberate glutathione and its potent relatives; these are the hosts behind the broths and sauces that have carried kokumi in the world's kitchens for a thousand years, now quantifiable to the milligram. **Onion** is garlic's gentler sibling, the sweet-sulfur allium. And **vanilla**, the outlier, proves the axis is not only savory: its phenylpropanoid aroma laid over a γ-glutamyl body is why true vanilla feels *thick* in a way the synthetic single-molecule imitation never does. Six organisms, one receptor, one bond — the whole generous axis assembled from the plants a family reaches for without thinking.

What the number costs, and what it cannot do

A number is only worth what its error bars allow, and on a modulatory axis the error bars matter more than anywhere else, because kokumi multiplies its neighbors rather than standing beside them. Our design specification targets a coefficient of variation of **CV ≤ 7%** for engineered batches, against roughly **42%** in the wild-type plants these lines descend from — and the gap is the difference between a designable ingredient and a guess. CV is simply the standard deviation as a fraction of the mean, σ/μ. On Phoenix's predicted IFI-K of 1,900, a 7% CV is an uncertainty of **±133** (0.07 × 1,900); the wild-type 42% is **±798** — a spread six times wider, nearly half the signal, large enough that two genuinely different kokumi levels could not be told apart. For a heat axis that error is an annoyance. For an axis whose whole function is to raise the gain on *other* tastes, a 42% swing in the multiplier is a dish that is thin one batch and cloying the next.

The error also propagates, and the formula tells you how. Because IFI-K is a **ratio** of two EC_{50} values, their relative errors add in quadrature:

$$CV(\text{IFI-K}) = \sqrt{\,CV(EC_{50}\ \text{ref})^2 + CV(EC_{50}\ \text{cand})^2\,}$$

If the reference is held at 5% and the candidate batch at 7%, the score's own uncertainty is $\sqrt{(5^2 + 7^2)} = \sqrt{74} \approx$ **8.6%** — on a 1,900 reading, about ±163. Run the candidate at the wild-type 42% instead and the term swamps everything: $\sqrt{(5^2 + 42^2)} \approx$ **42.3%**, ±804. The discipline of the controlled chamber is, in the end, a discipline about this square root.

And there are honest limits the number cannot argue away. Kokumi is **conditional**: because CaSR is first a calcium sensor, kokumi potency shifts with the calcium and the ionic background of the food itself — the same peptide reads louder in a mineral-rich broth than in distilled water, so a score is always a score *in a defined matrix*. It is **unmeasurable in isolation** — there is no kokumi taste to rate on its own, only a modulation of companions, which is why the assay reads the receptor directly rather than asking a panel to name a flavor that has no name. It **saturates**: past the 5,000 ceiling the receptor is swamped and the rounding tips toward a dull, flat heaviness. At very high doses the sulfur chemistry can cross into **bitterness**, CaSR's cooperative neighbor on the tongue, so richness and bitter are two ends of one dial. And it is **personal**: CaSR carries common human polymorphisms, so the same broth is measurably rounder to some mouths than to others, and the number is a population mean, not a promise to any single tongue. We publish the limits beside the score because a number offered without its uncertainty is not rigor; it is decoration.

The kitchen it unlocks

What the axis gives back, finally, is a kind of honesty in the pan. Because kokumi makes the mouth read *rich* without adding sugar or salt or fat, a cook can pull those three back and never be caught — a broth that carries its body on γ-glutamyl peptides instead of a slick of fat, a sauce that rounds on the receptor instead of the sugar, a seasoning that lets you cut the salt and keep the sense of fullness that salt was standing in for. We write that carefully and without a single health claim, because the ethos under this whole family's work is not a promise printed on a box; it is an intention — that food be grown to carry the most it can, and that the body be given the most it can draw, and that these are the same act. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. Kokumi is where that intention becomes a technique: the generosity of making everything taste more like itself, with less of the things we were only ever using to fake fullness.



Vanilla — the outlier of the kokumi six, proof that richness is not only savory; its thickness is a γ-glutamyl body beneath a phenylpropanoid aroma, the most complex flavor in the kitchen and the hardest to fake.

Our grandmother would not have followed a word of the receptor or the bond or the square root. But she knew the thing itself completely — the part of the soup that stays with you, the spoonful that feels held, the broth she let go one more hour because her mouth told her it was not yet *whole*. She was reading CaSR with her tongue and tuning γ-glutamyl peptides with a stockpot and the clock, the way every good cook who ever lived has done it: by instinct, and never once able to say how much. The number does not replace what she knew. It writes it down. And a thing written down is a thing you can hand to the next person at the table — which is, when we are honest about why a family builds any of this, the only reason worth doing it at all.



PART I OF V

The Invisible Axis

What kokumi is · the calcium-sensing receptor · scoring a modifier

Chapter 1 · The three sensations — *mouthfulness, continuity, complexity.*

Begin with the experience, because the experience is real and precise long before the chemistry explains it. When a trained panel is asked to describe what a kokumi substance does to a base of dashi or chicken stock or a simple sugar solution, they do not reach for new flavors. They reach, almost every time, for the same three words, and those three words turn out to map cleanly onto three measurable things. We have run this exercise more times than we can count, across years of panels, on broths our family has simmered since before any of us could read an assay — and the convergence is the first fact worth trusting. Give a hundred careful tasters the same spoonful and they will invent a hundred metaphors; strip the metaphors away and the same three axes remain.

The three words are **mouthfulness**, **continuity**, and **complexity**. Hold them, because everything in this book is a way of accounting for them, and because each one, pulled gently away from the others, turns out to be a thing you can put a number on. Our grandmother's *stays with you* was all three at once, felt as a single warmth. The Index pulls them apart so we can build each one on purpose.

Mouthfulness — the flavor that fills a space it does not occupy

The first sensation is **mouthfulness** — in Japanese, *koku* in its narrowest sense, sometimes translated as *thickness* or *body*. It is the sensation that the flavor occupies more of the mouth, spreads wider across the tongue and palate, fills the space rather than sitting in a line down the center. It is spatial. A thin consommé with kokumi added does not get more viscous — the physical thickness on the tongue is unchanged, you could measure it with a rheometer and find nothing — and yet it *feels* thicker.

That gap is the hinge of the whole axis, so it is worth making it exact. Pour the same consommé twice, once plain and once with a kokumi substance dissolved into it at a few parts per hundred thousand. Put each under a rotational rheometer and shear it: the viscosity curves — shear stress against shear rate, the Newtonian line of a mostly-water broth — lie on top of each other. There is, physically, nothing more in the mouth to feel. Now put the same two cups in front of the panel, and the second one is reliably scored *thicker*, *rounder*, *more filling*. A perception has moved while the physics has not. When a sensation rises and the rheology refuses to follow, you are no longer measuring a texture. You are measuring a receptor, and the rest of this book is spent describing which one and how to speak to it.

The Japanese kitchen that first catalogued these things had already split mouthfulness into two finer motions, and the distinction is not decoration — it is how our panel scores. *Atsumi* (厚み) is thickness, the vertical sense of the flavor having depth, a floor under it. *Hirogari* (広がり) is spread, the lateral reach, the way the taste moves out into the corners of the mouth instead of running down the midline of the tongue. A flat, competent stock has flavor but no *hirogari* — it stays in its lane. A stock with kokumi pushes outward to the molars and the soft palate, and the eater reports, without being prompted, that there is simply *more of it in more places*. Two axes of a single filling-out: one down, one sideways, neither of them anything a viscometer will ever see.

Continuity — a decay constant, slowed

The second sensation is **continuity** — in Japanese, *jizoku* (持続) — the persistence of flavor over time, the long tail. Flavor compounds have a shape in time: they rise to a peak intensity and then decay as they are diluted by saliva and cleared from the receptors. Kokumi substances flatten and extend that decay. The peak is not much higher; the *tail* is longer.

This one we can write down cleanly, because perceived intensity over time behaves, to a good approximation, like a first-order decay. Hold a stimulus in the mouth and rate it second by second on a slider — the time–intensity protocol our sensory lab runs — and after the initial rise the curve falls off like

```
I(t) = I₀ · e^(−t/τ)
```

where I_0 is the peak intensity and τ is the time constant of clearance. The total perceived experience — the thing that makes a spoonful feel generous rather than brief — is the area under that curve, and for a clean exponential that area has a beautifully simple form:

```
AUC = ∫₀^∞ I₀ · e^(−t/τ) dt = I₀ · τ
```

The whole experience is the peak *times* the time constant. Which tells you there are only two levers. You can push the peak I_0 up — that is what adding more of a tastant does, and it fades just as fast — or you can slow the decay τ, which is what a kokumi substance does. Take an illustrative stock whose savory note clears with $\tau_0 = 8$ s. Its perceptual half-life is $t_{1/2} = \tau \cdot \ln 2 = 8 \times 0.693 =$ **5.5 s** — half the flavor is gone in five and a half seconds. Now lengthen the time constant to $\tau = 12.8$ s, a 1.6-fold slowing, while leaving the peak untouched. The half-life stretches to $12.8 \times 0.693 =$ **8.9 s**, and the area under the curve rises by exactly the ratio of the time constants, $12.8 / 8 =$ **1.6** — sixty percent more total perceived flavor delivered, every bit of it in the back half, with the peak never moving. The spoonful that *stays with you* is, quite literally and quite measurably, a flavor whose decay constant has been slowed. We did not add intensity. We added *time*.

Complexity — the separate tastes drawn into a chord

The third sensation is **complexity**, or **harmony** — the sense that the separate tastes of a dish have been drawn together, that the salt and the savory and the sweet are no longer three soloists but a chord. This is the hardest to measure and the easiest to recognize. It is why a stock enriched with aged-cheese rind, or a curry finished with a spoonful of dark fermented paste, reads as *integrated* in a way the same dish without it does not. The parts have not changed. The relationship between them has.

We score it, carefully, as a harmony or integration rating, and we treat its number with more humility than the other two, because integration is a judgment about a *whole* and judgments about wholes drift. But the direction is never in doubt: add the kokumi substance and the panel's agreement about the dish tightens. Descriptors that were scattered — a little sweet, a little salty, a little thin — collapse toward a single word, *round*, or *complete*, or in our grandmother's kitchen simply *right*. Complexity is the axis where the arithmetic of separate tastes becomes a single experienced thing, and it is the axis that most resists being faked by any one added compound, because it is a property of the relationships, not the parts.



Garlic in cross-section — the clove is a warehouse of γ-glutamyl peptides, stored stable and odorless until the knife.

The molecule you already carry

The history matters here, briefly, because it is a history of exactly this difficulty of naming — and because it ends on a molecule that is already, quietly, inside you. The word *kokumi* entered the technical literature through the study of garlic. Researchers extracting the water-soluble, non-volatile fraction of garlic and onion — the part left behind after you remove everything that smells — found that it enhanced the mouthfulness, continuity, and complexity of other foods while contributing almost no taste of its own. The active principle, isolated and identified, was a peptide every cell of your body builds for itself: **glutathione**, γ-glutamyl-cysteinyl-glycine, three amino acids and a molecular weight of 307.32 grams per mole.

The cell's master antioxidant, the thing that keeps your own chemistry from rusting, moonlighting in the mouth as the molecule that makes soup taste like care.

We find that fact steadying rather than strange. The compound at the root of this axis is not an additive foreign to the body; it is something the body already manufactures by the gram and already knows. From that first garlic extract the family of active compounds widened into a whole catalog of **γ-glutamyl peptides**, some of them dozens of times more potent than glutathione itself — and the thing they share, the reason they are stable enough to survive a clove in a cellar and a pot on a fire, is a single unusual bond.

An ordinary peptide links the α-carboxyl of one amino acid to the amino group of the next; that is the bond every digestive peptidase in your gut is built to cut. A γ-glutamyl peptide links through glutamate's *side-chain* carboxyl — the γ-carboxyl — instead. The resulting amide is nearly invisible to the standard α-peptidases. It does not get cleaved in storage, it does not get cleaved in the first seconds of chewing, and so these peptides persist exactly where and when a flavor modifier needs to persist. The stability that lets garlic hoard them for a season is the same stability that lets them act on the tongue. Chemistry and shelf life and continuity are, here, one property seen three ways.

The enzymology is worth carrying, because it is the hook every engineering decision later hangs on. The γ-glutamyl bond is laid down by **glutamate–cysteine ligase** (GCL, also called γ-glutamylcysteine synthetase), which spends a molecule of ATP to join glutamate's γ-carboxyl to cysteine; **glutathione synthetase** then caps it with glycine to make glutathione proper. A third enzyme, **γ-glutamyltransferase** (GGT), can pick the γ-glutamyl group back up and hand it to a different amino acid entirely, minting the wider family of γ-glutamyl dipeptides that fermentation and the allium bulb are full of. A plant that stockpiles kokumi is a plant running these enzymes hard. Garlic does exactly this, storing γ-glutamyl-*S*-allyl-cysteine and its kin — stable, odorless, inert — until a knife ruptures the cells. (The sharp smell that follows is a *different* garlic pathway, alliin meeting the enzyme alliinase to make allicin; that is pungency, a separate axis on a separate receptor. The clove runs two chemistries at once, which is why on our Index garlic registers on both kokumi and pungency and why cutting it releases a sensation the whole kitchen can smell while the kokumi stays silent, dissolved, and waiting.)

The receptor underneath: a calcium sensor repurposed for care

All of these peptides speak to one receptor, and knowing which one turns these three sensations from a mystery into a mechanism. The target is **CaSR**, the calcium-sensing receptor — a class C G-protein-coupled receptor whose day job, in your parathyroid and your kidney, is to read the calcium concentration of your blood and hold it steady. It is a large, two-lobed machine: an extracellular "Venus flytrap" domain that closes around what it binds, a cysteine-rich linker, and the seven-transmembrane bundle that every GPCR uses

to carry the news inward. It works as a dimer, two copies clasped together, and that pairing matters for reasons that show up directly in the math.

Kokumi peptides do not switch this receptor on the way calcium does. They are **positive allosteric modulators** — they bind at a site distinct from the main agonist pocket and make the receptor more responsive to what is already there. This is the molecular translation of the single most important thing about the whole axis: kokumi is not a taste you add, it is a *sensitivity you raise*. When a γ-glutamyl peptide settles into the CaSR of a taste cell, the cell's answer to the surrounding calcium and amino acids gets louder. Inside the cell, CaSR couples through Gq/11 to phospholipase C-β, which splits PIP_2 into IP_3 and DAG; the IP_3 opens calcium stores and a transient of intracellular calcium rises — the cell's way of saying *more*. Nothing new has been introduced into the flavor. An existing conversation has simply been turned up.

And because the receptor is a dimer, its dose–response is **cooperative** — the two halves do not respond independently, so the curve that relates concentration to effect is steeper than a simple binding would give. That steepness is not a footnote; it is why a very small change in how much kokumi peptide is present can swing the sensation noticeably, and the dose–response puts exact numbers on it.

Why kokumi can only be measured as a difference

The measurement problem this creates is real and shaped the entire axis, so we name it at the outset. You cannot measure kokumi in water, the way you can measure sweetness or heat, because kokumi in water is nothing — a positive allosteric modulator with nothing to modulate is silent, and the panelist tastes a faint saline blankness and rates everything zero. Kokumi exists only *relative* to a flavor it modifies, so it can be measured only as a *difference*: the same base, scored with and without the kokumi substance, the gap between the two scores being the kokumi effect.

We standardized that base early and have never changed it — a defined solution of glutamate, salt, and a trace of sugar, chosen because it is savory enough to give kokumi something to grip and neutral enough not to compete — because a modifier can only be read against a fixed thing to modify. This is the single most important methodological fact about the axis: **kokumi is a measured difference, not a measured presence**, and every IFI-K number in this book is, at bottom, the size of a gap.

Measuring a difference has a cost, and honesty requires stating it. If the panel scores the plain base at P_0 and the modified base at P_1, each with its own measurement noise of standard deviation σ_p, then the quantity we actually care about — the difference $P_1 - P_0$ — carries a noise of $\sigma_p \cdot \sqrt{2}$. The act of subtracting two measurements inflates the uncertainty by a factor of 1.414. A kokumi number is intrinsically noisier to earn than a sweetness number, which is one reason the whole program lives or dies on reproducibility, and why

we hold batch-to-batch variation to **a coefficient of variation of 7 percent or better**. The difference is the signal; the $\sqrt{2}$ is the tax; low CV is how we still read the signal cleanly through the tax.

Putting a number on a gap: the IFI-K math

Here is where the sensation becomes a score. The Iyengar Flavor Index anchors every axis to the half-maximal activation of a receptor — the EC_{50}, the concentration at which a stimulus drives its receptor to half of full response — and inverts it against a fixed reference:

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

Read the formula slowly, because its shape is the whole idea. EC_{50} sits in the denominator of the candidate: a *smaller* EC_{50} means the substance reaches half-activation at a *lower* concentration, which means it is more potent — so a smaller EC_{50} gives a *larger* score. The Index does not measure how much of a compound a food holds. It measures how *loudly* that compound speaks to its receptor. The factor of 1,000 simply sets the anchor: when a candidate's EC_{50} equals the reference's, the ratio is one and the score is exactly 1,000.

For kokumi, the reference chooses itself. We anchor **IFI-K = 1,000** to the half-maximal potentiation of CaSR by glutathione — the first kokumi peptide ever isolated, the one the body already makes, the natural zero-point of the family. Everything else on the axis is read as a multiple of glutathione's loudness.

Now work it the other way, because this is the move that makes the Index useful: a score can be *inverted back into a concentration ratio*. Our full-spectrum flagship cultivar, Phoenix, carries an engineered kokumi channel scored at **IFI-K = 1,900**. Put that number into the formula and solve for the EC_{50}:

$$1{,}900 = 1{,}000 \times EC_{50}(\text{glutathione}) \div EC_{50}(\text{Phoenix kokumi})$$
$$EC_{50}(\text{Phoenix kokumi}) \div EC_{50}(\text{glutathione}) = 1{,}000 \div 1{,}900 = \mathbf{0.526}$$

The engineered fraction reaches half-maximal CaSR potentiation at **52.6 percent** of the glutathione concentration — it is 1.9 times louder than the molecule the Index is built around. The scorecard ceiling helps place that: our per-axis normalization maximum is 5,000, which by the same arithmetic corresponds to an EC_{50} of 1,000 ÷ 5,000 = 0.200 times the reference. Phoenix's kokumi, at 1,900, sits at 38 percent of full scale — a serious, built-in richness, and still with headroom to spare.

SUBSTANCE	IFI-K	EC_{50} (RELATIVE TO REFERENCE)	MEANING
Glutathione — reference anchor	1,000	1.000 × E_ref	the natural zero-point
Phoenix kokumi channel · SS-0096	1,900	0.526 × E_ref	1.9× louder than glutathione
Scorecard ceiling (per-axis max)	5,000	0.200 × E_ref	full-scale reference

A worked unit conversion, to keep the EC_{50} from floating free of anything you could weigh. Take the reference anchor, for this illustration, at an EC_{50} of 0.50 millimolar — half a millimole of glutathione per liter. Glutathione's molar mass is 307.32 g/mol, so

> 0.50×10^{-3} mol/L × 307.32 g/mol = 0.1537 g/L ≈ **154 mg/L**, i.e. about 154 ppm.

That is the kind of concentration we are speaking of — grams per thousand liters, a pinch dissolved in a stockpot — and Phoenix's louder channel would reach the same half-effect at 0.526 of it, nearer 81 mg/L. These are vanishingly small masses doing a large perceptual job, which is exactly the signature of a receptor-level effect rather than a bulk one.

Why the receptor's steepness matters: the dose–response, worked

Because CaSR is a cooperative dimer, the curve from concentration to effect is steeper than a textbook binding, and that steepness decides how forgiving the axis is to formulate against. Write the fractional response with a Hill function:

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

where n is the Hill slope. At $[A] = EC_{50}$ the response is 1/(1+1) = 0.5 for any n — that is the definition of EC_{50}, and it is why the Index can anchor to it so cleanly. The interesting question is how wide a concentration window separates a barely-there effect from a nearly-saturated one. Solve r = 0.1 and r = 0.9:

- For a **non-cooperative** receptor (n = 1): 10 percent response occurs at $EC_{50} \div 9$, 90 percent at $9 \times EC_{50}$ — an **81-fold** concentration span from faint to full.
- For a **cooperative** receptor (n = 2, in the range a CaSR dimer can show): the same 10-to-90 window collapses to $EC_{50} \div 3$ up to $3 \times EC_{50}$ — a **9-fold** span.

The general result is tidy: the 10-to-90 span is 81^(1/n)-fold. Cooperativity *sharpens* the response, and that cuts both ways in the kitchen. A steep kokumi curve means a small, well-chosen dose lands you most of the effect without a long slow climb — generous. It also means the plateau arrives fast: past saturation, more peptide buys almost nothing, and a formulator chasing "even more body" by adding more is simply spending money on the flat top of a curve. The cooperativity we admire in the receptor is also the thing that tells us when to stop.



Phoenix under the microscope stage — a pepper built to register kokumi at IFI-K 1,900, 1.9× the glutathione reference, alongside seven other axes.

Precision is the whole argument

A number is only a measurement if it comes back the same tomorrow. Phoenix's kokumi reads IFI-K = 1,900; at our platform tolerance of CV ≤ 7 percent, the standard deviation on that figure is 0.07 × 1,900 = **133**, so the honest statement is 1,900 ± 133, and a 95-percent interval (±1.96σ) runs roughly 1,639 to 2,161. Set that against the batch variability of an un-engineered biological source, where a CV of 40 percent and more is ordinary: there the same nominal 1,900 would carry a standard deviation near 800, and "1,900" could as easily have been 1,100 or 2,700. At that spread the number has stopped being a measurement and gone back to being an adjective. The CV is not a quality-control afterthought bolted onto the science; it is the precondition that lets the kokumi *difference* — already paying the √2 noise tax of being a difference — mean anything at all.

The ladder of kokumi carriers

Kokumi is not a pepper problem any more than any axis is. The γ-glutamyl machinery — GCL, glutathione synthetase, GGT — runs in essentially every plant and in the microbes we ferment with, which is why the carriers of this axis are the savory bedrock of nearly every cuisine on earth. These are the host organisms our program works through on the kokumi side, with the counts canonical to our library:

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

The pairings in that third column are the whole story of why these foods anchor their kitchens. Garlic and onion carry kokumi beside pungency — the bulb that gives a dish both its depth and its bite. Soybean and yeast carry kokumi beside umami — which is why a good miso, a long-fermented tamari, a spoonful of yeast extract, read not merely as *savory* but as *round*, umami and kokumi arriving together, the glutamate speaking and the γ-glutamyl peptides making the receptor listen harder. Vanilla stands slightly apart: a kokumi carrier on the sweet side of the pantry, part of why real

vanilla reads as *full* and *long* where synthetic vanillin reads as flat and brief — the mouthfeel and the continuity are doing work the aroma gets the credit for. Across just these five organisms our platform holds 1,240 constructs, every one scored against the same fixed base, every one designed toward the fullest flavor — and the fullest nourishment — it can carry.



Vanilla in cross-section — a kokumi carrier on the sweet side of the pantry, where the body and the long tail do the work the aroma gets the credit for.

When kokumi is engineered into a cultivar rather than extracted from one, it becomes one channel among the eight. Phoenix, the first cultivar we designed to register on all eight axes at once, carries its kokumi at 1,900 inside a full profile — and reading that profile is reading a flavor written as a chord:

IFI PROFILE

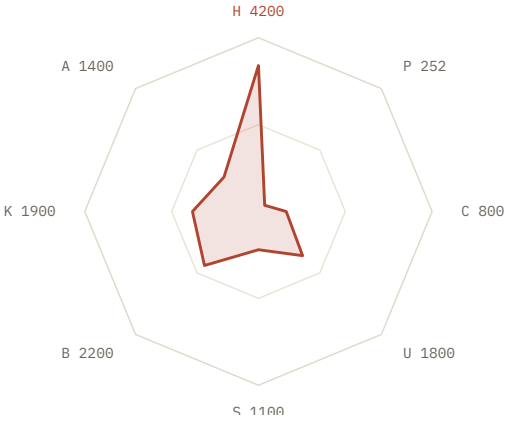

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

AXIS	CODE	PHOENIX · SS-0096	SHARE OF FULL SCALE (MAX 5,000)
Heat	H	4,200	84%
Zero-Bitter	B	2,200	44%
Kokumi	K	1,900	38%
Umami	U	1,800	36%
Astringency	A	1,400	28%
Sweetness	S	1,100	22%
Cool-Burn	C	800	16%
Pungency	P	252	5%

The kokumi line is not the loudest voice in that chord, and it is not meant to be. At 1,900 it sits just behind the umami at 1,800 and just ahead of it — the two savory axes carried together at nearly the same loudness, which is precisely the design: the glutamate to speak and the kokumi to make the mouth hold it longer and feel it wider. The heat is the thing you notice; the kokumi is the thing that makes the heat feel like it belongs to a meal rather than a dare.

Edge cases, failure modes, and the honest limits

An axis you can design is an axis you can get wrong, and the ways kokumi fails are as instructive as the ways it works.

It saturates, and the plateau is real. The cooperative Hill curve that makes a small dose effective also means that past roughly a few times the EC_{50} the response has flattened to within a few percent of its maximum. More peptide beyond that adds cost, not continuity. The failure is not a bad flavor; it is a wasted one.

It is base-dependent, so its number is not portable. IFI-K is defined against one standardized glutamate-salt-sugar base, and it means what it means only there. Drop the same substance into a thin, unsavory medium and the modulator has little to modulate — the score collapses toward the water-blankness we began with. A kokumi number lifted out of its base and compared against a different base is not a comparison; it is a category error. This is the discipline the fixed reference buys us, and the discipline it demands in return.

Texture can impersonate it. Added fat or a thickener will make a broth feel rounder too — but that roundness is rheology, and the rheometer sees it. The signature of *true* kokumi is the one we opened with: perception rises while the viscosity curve does not move. Any time body appears *with* a change in measured viscosity, we have to ask how much of it was the receptor and how much was the spoon's resistance, and we design the controls to split the two.

The composite can hide its own structure. A single IFI-K rolls three motions — mouthfulness, continuity, complexity — into one figure, and those three do not always move together. A substance can lengthen the tail (*jizoku*) beautifully while barely widening the spread (*hirogari*), or fill the mouth without integrating the dish. When they diverge, the one number lies by averaging, so we report the components alongside the composite rather than letting the headline figure stand alone.

COMPONENT (JAPANESE)	ENGLISH	WHAT IT MEASURES	INSTRUMENT
atsumi (厚み)	thickness / mouthfulness	spatial filling of flavor	panel + CaSR assay
hirogari (広がり)	spread	lateral reach across the palate	trained panel
jizoku (持続)	continuity	persistence over time	time–intensity curve
koku (こく)	richness (the whole)	integrated sum of the three	composite IFI-K

And the honest outer limit: of the eight axes, kokumi is the one whose full biology is least settled. The CaSR account is well supported and it is the frame we build and score within, but the complete set of mechanisms by which γ-glutamyl peptides shape oral perception is still being mapped, in our computational work and others'. We hold our numbers as *operational* — anchored, reproducible to 7 percent, honest about being differences — and not as the last word on a receptor the field is still learning to read. We make no claim about what these peptides do past the tongue; glutathione is something your body already makes and manages, and at the pinch-in-a-stockpot concentrations that carry flavor, our subject is the mouth and only the mouth.

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

These are not poetic flourishes, the three sensations — they are the axes along which our panel scores and the Index measures. *Atsumi* down, *hirogari* outward, *jizoku* across time, and *koku* the chord the three make together; a positive allosteric modulator raising a receptor's sensitivity rather than adding a taste; a score that inverts cleanly into an EC_{50} and a half-life and a mass you could weigh. That is the axis this book measures. Not a flavor. A way of being more.

Chapter 2 · The calcium-sensing receptor — *the biology of more.*

To score kokumi we had to find the lock the key fits, and the lock turned out to be a receptor that was never, in its day job, about food at all. The **calcium-sensing receptor** — CaSR — is the protein your parathyroid glands use to read the concentration of calcium in your blood, so your body can hold that concentration inside a window a few percent wide, hour after hour, for the whole of your life. It is one of the most important homeostatic sensors you own. It also sits on the taste cells of your tongue, and it is the receptor that reads kokumi. The same molecule that keeps your bones from dissolving to buffer your blood is the one that tells you a broth has become round and whole and worth staying with. We find that almost unbearably beautiful, and we design against it with the respect that fact deserves.

That dual citizenship is not a coincidence; it is the whole mechanism, and it is worth understanding in structural detail because the structure is what we design against. Nothing in this chapter is metaphor. Every sensation kokumi names — the mouthfulness, the thickness, the long quiet continuity after a spoonful — is a conformational change in a known protein, coupled to a known second-messenger cascade, read out as a known pattern of calcium inside a cell. We will walk it from the outside of the membrane to the inside, and then we will turn it into a number.

A homeostatic sensor that moonlights as a flavor receptor

Begin where the receptor spends most of its life, which is not the tongue but the **parathyroid chief cell**. Those four small glands behind the thyroid do exactly one urgent thing: they hold the ionized calcium in your blood near a set point of roughly **1.1 to 1.3 millimolar**, a band so narrow that the whole of your neuromuscular life — every heartbeat, every nerve firing, every muscle that answers it — depends on it not drifting. CaSR is the instrument that reads the band. When blood calcium rises, calcium ions bind the receptor, it switches to its active shape, and it *suppresses* the secretion of parathyroid hormone; when blood calcium falls, the receptor quiets, hormone is released, and the body pulls calcium back out of bone and kidney until the band is restored. It is a thermostat written in protein, and it is unusual among receptors in that its physiological job is to be exquisitely, continuously sensitive to small shifts in the concentration of a tiny ion.

Hold onto that one word — *sensitive* — because it is the hinge of the entire kokumi axis. A receptor whose evolutionary purpose is to notice small changes in its agonist, and whose activity can be *tuned up* by the right accompanying molecules, is a receptor practically built to be enhanced. Nature did not design CaSR to taste food. It designed a precision sensor and then, on the taste cells of the tongue, pointed that same precision at the contents of a mouthful. The flavor we call kokumi is, at bottom, your tongue borrowing the most carefully calibrated small-molecule sensor your body owns.

The architecture: a class C receptor, built as a dimer

CaSR is a **class C G-protein-coupled receptor** — the same structural family as the metabotropic glutamate receptors, one of which, mGluR1, is the umami axis of this very Index. The kinship is not trivia. It means the two savory axes, umami and kokumi, are read by cousins built on the same architecture, which is a large part of why they work on each other so profoundly when a dish carries both: they speak closely related structural languages, and a cell already primed by one is halfway to answering the other.

A class C GPCR is a two-part machine, and the parts matter because our molecules act on only one of them. Reaching out from the cell surface is an enormous **extracellular domain** shaped like a Venus flytrap — two lobes hinged together, open in the resting state, snapping closed around a ligand when one binds in the cleft between them. The flytrap connects, through a short **cysteine-rich domain** that acts as a rigid driveshaft, to the familiar **seven-transmembrane (7TM) bundle** threaded through the membrane, the part that actually talks to the G-protein inside the cell. When the flytrap closes, that motion travels down the cysteine-rich linker and twists the 7TM bundle into its active arrangement. The outside of the cell catches the molecule; the driveshaft carries the news; the inside of the cell answers.

And critically, CaSR works as a **homodimer**: two identical receptors clasped together by disulfide bonds at the top of their flytraps, their two clefts closing in a cooperative dance so that the binding of one ligand makes the next bind more easily. This is not an incidental fact of its crystallography. The dimer *is* the sensitivity. Two flytraps, mechanically coupled, give a response that turns on more sharply than either could alone — a steepness we will put a number on, because it is the single most useful thing the receptor's shape does for a designer.

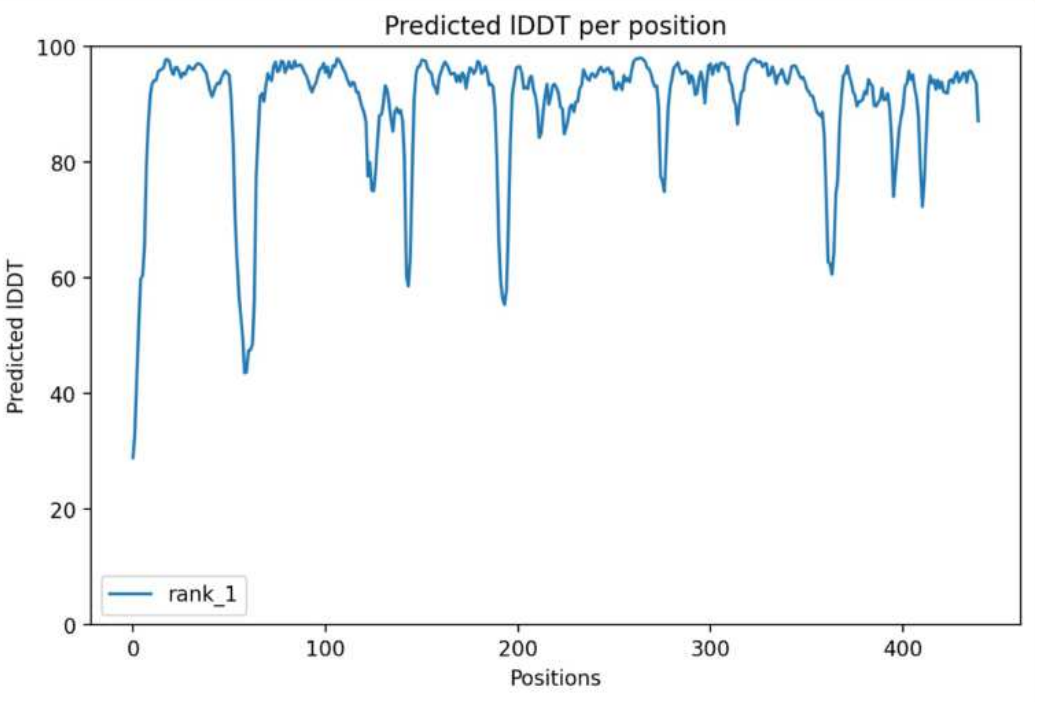

The Venus-flytrap domain of a class C receptor closing around its ligand is the mechanical heart of kokumi; we score the confidence of our pathway models the same way — high confidence through the binding cleft, lower at the flexible hinges.

The flytrap, and the biology of *more*

The flytrap's cleft is where the biology of *more* lives. Its primary ligand is the **calcium ion** itself — several of them, binding at multiple sites, each one helping to stabilize the closed, active conformation.

But calcium is not the only thing that binds there. **L-amino acids**, especially the aromatic ones — L-phenylalanine, L-tryptophan — bind alongside calcium as *co-agonists*, wedging into the closing cleft and making the flytrap shut more readily at any given calcium level. They do not open the lock on their own; they oil it. They lower the threshold at which calcium can close it.

This is **positive allosteric modulation**, and it is the deep logic of the whole axis. Most taste receptors are agonist machines: a thing binds, a taste happens, more of the thing means more of the taste. CaSR is different in kind. Its native behavior is to be *made more responsive* by the presence of the right companion molecules — to take whatever calcium and amino acid a mouthful already holds and read it as *more* than the sum would otherwise suggest. CaSR is, structurally, the biology of enhancement. That is why kokumi has no flavor of its own to point at. It is not a taste. It is a gain control on every taste already present, and the receptor that performs it is one whose entire design is to have its gain turned up.

γ-glutamyl peptides, and the enzymes that write them

The most effective companions food has ever offered the flytrap are the **γ-glutamyl peptides** — glutathione and its more potent relatives. They act at or near the amino-acid site in the cleft, and when one binds, it stabilizes the closed conformation more firmly than a lone amino acid can. To understand why they are special you have to look at the bond that names them, because that bond is both their power and the reason they taste of almost nothing on their own.

Glutathione is a tripeptide: **γ-L-glutamyl–L-cysteinyl–glycine**. An ordinary peptide bond joins the α-carboxyl group of one amino acid to the amino group of the next. The γ-glutamyl peptides break that rule at their first junction — glutamate is linked through its *side-chain* carboxyl, the γ-carboxyl, not the backbone one. That unusual **γ-linkage** is what the flytrap recognizes, and it is also what makes the peptide resistant to ordinary digestive proteases, which are built for α-bonds and slide off the γ-bond unsure what to do. The molecule survives long enough in the mouth to speak.

The cell writes these peptides in two deliberate enzymatic steps, and knowing them matters because they are the steps a plant can be engineered to run harder. First, **glutamate-cysteine ligase** (GCL, the rate-limiting enzyme of the route) joins glutamate to cysteine through that γ-bond, spending one molecule of ATP to force the unfavorable linkage — this is the committed, controlled step, and it is feedback-inhibited by the very glutathione it ultimately produces, a brake the cell uses to keep the pathway from running away. Second, **glutathione synthetase** adds the glycine, spending a second ATP, to finish the tripeptide. The more potent kokumi peptides — γ-glutamyl-valyl-glycine chief among them, several-fold stronger at CaSR than glutathione itself — are assembled by a related activity, **γ-glutamyl transferase**, which can transfer the γ-glutamyl group onto other amino acids and short peptides. Two ATP, three amino acids, one unusual bond: that is the whole cost of the richest savory signal a plant can carry, and every one of those inputs — glutamate,

cysteine, glycine — is something a green cell makes for itself. We have never needed an animal in the loop to find depth. The larder was always vegetarian.



A head of garlic in cross-section — an allium is a factory for γ-glutamyl peptides, the richest co-agonists a plant can offer the flytrap. These molecules have been in the human diet since the first bulb was dug out of the ground.

Inside the cell: the signal that rings

When a γ-glutamyl peptide stabilizes the closed flytrap, the cysteine-rich driveshaft twists the 7TM bundle, and inside the cell the receptor couples to **Gq/11**. The activated G-protein switches on **phospholipase C-β**, which cleaves the membrane lipid **PIP_2** into two second messengers: **IP_3** and **DAG**. The IP_3 diffuses to the endoplasmic reticulum and opens its calcium stores, releasing a flood of **intracellular calcium** into the cytoplasm. CaSR also couples to **Gi/o**, lowering cAMP, and drives the **ERK1/2** cascade that reaches into the nucleus — a receptor this important never speaks through only one channel. But in a taste cell it is the released intracellular calcium that is the signal, and here is the beautiful part: CaSR's calcium signal is characteristically **oscillatory and sustained**, not a single spike.

Many receptors answer a stimulus with one rise and fall of intracellular calcium — a shout, then silence. CaSR answers with a *train* of calcium spikes whose **frequency encodes the strength of activation**: a stronger kokumi stimulus is written not as a taller single spike but as a faster, steadier train of them, sustained for as long as the ligand sits in the flytrap. The receptor does not shout once. It rings. And a signal written as a sustained oscillation is a signal the brain reads as *lasting* — which is, we are convinced, the cellular substrate of continuity, the long quiet tail we measure in a kokumi stimulus's time-intensity curve. It also means the receptor resists the fatigue that dulls other tastes: where a constant sweet or salty stimulus desensitizes its receptor and fades within seconds, the oscillatory CaSR signal persists, because frequency-encoded signals do not wash out the way a single sustained push does. The kokumi receptor is built, at the level of its signaling, to not let go.

In the mouth this all happens on the **Type I taste cell** — long dismissed as the taste bud's quiet supporting cell, now understood to carry CaSR and to broadcast the kokumi signal sideways across the other taste modalities. That lateral broadcast is the cellular

reason kokumi modifies *everything* in a mouthful rather than contributing a taste of its own. It is the cell that makes the whole bud sing together.

Scoring a co-agonist: the arithmetic of kokumi

A sensation is not a measurement until it is a number, so we anchor kokumi the way we anchor every axis — to the half-maximal activation of the receptor, read through one formula:

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

EC_{50} is the concentration of a stimulus that drives its receptor to half of full response — low EC_{50} means a loud, potent ligand; high EC_{50} means a quiet one. Invert the candidate's EC_{50} against a fixed reference, scale by a thousand, and loudness becomes a score. For the kokumi axis we take the reference to be **glutathione**, the canonical γ-glutamyl peptide, the molecule the tongue has read as richness for as long as there have been tongues. Its half-maximal activation of reference CaSR *is* the axis's 1,000. The Index is defined purely by the ratio, so the score of any candidate never depends on the absolute anchor; but to make the arithmetic concrete, take glutathione's half-maximal concentration in round numbers near **2.5 mM** and let that be the thousand. Then a candidate's IFI-K reads off directly as a potency:

$$\text{potency ratio} = \frac{\mathrm{EC}_{50}(\mathrm{ref})}{\mathrm{EC}_{50}(c)} = \frac{\text{IFI-K}}{1000}$$

Work it on a real cultivar. **Phoenix** — SS-0096, our first construct engineered to register on all eight axes at once — carries kokumi at **IFI-K 1,900**. That is a potency ratio of 1.90: its kokumi peptide reaches half-activation at $\mathrm{EC}_{50}(c) = \mathrm{EC}_{50}(\mathrm{ref})/1.90 =$ **0.526 × 2.5 mM = 1.32 mM**, about 1,320 µM, or — converting through glutathione's molar mass of 307.3 g/mol — roughly **405 mg/L** of reference-equivalent peptide against the reference's 768 mg/L. In plain terms: Phoenix speaks richness **90% more loudly** than glutathione, needing barely half the concentration to half-fill the receptor. Against the scorecard's per-axis normalization ceiling of 5,000, an IFI-K of 1,900 sits at 38% of the full scale — present, generous, deliberately not maxed, because a full-spectrum plant spends its biology across eight axes rather than screaming on one.



Phoenix (SS-0096) on the microscope stage — the full-spectrum cultivar that carries kokumi at IFI-K 1,900 across a design profile our model predicts at approximately 2.95 million SHU of heat.

The same formula reads every axis Phoenix expresses, and running the whole profile through it is the clearest way to see what an IFI score *means* — including what a score below the anchor means.

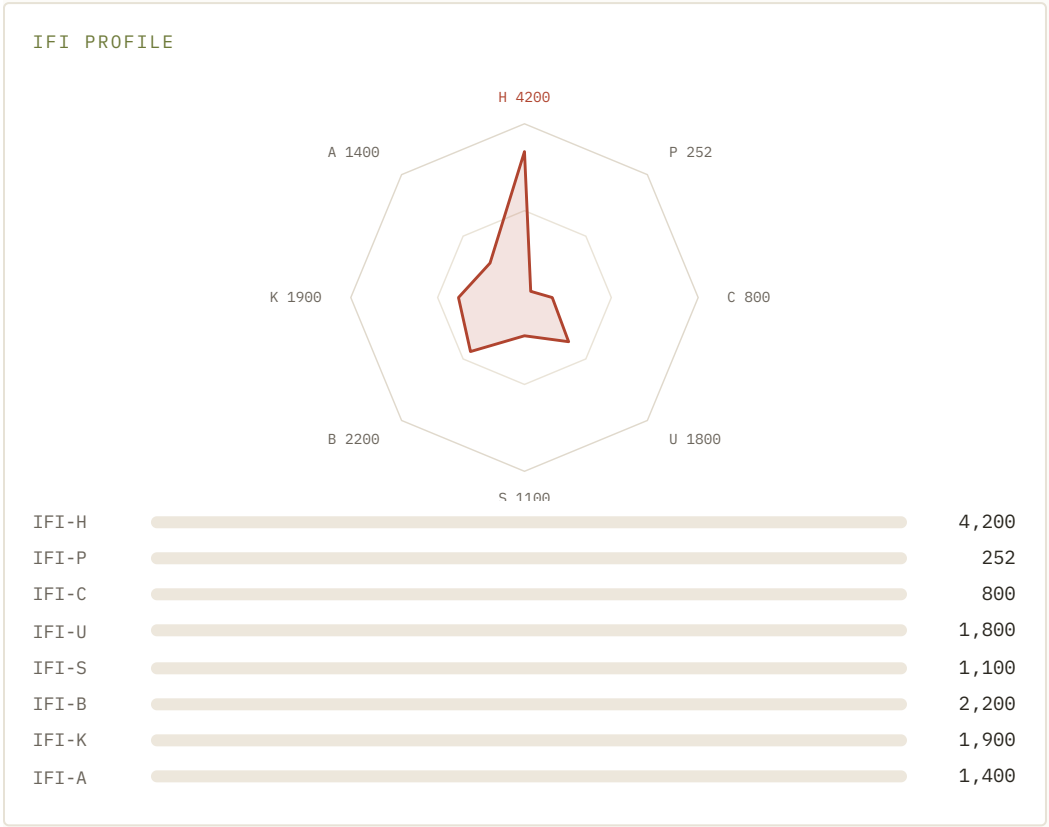

AXIS	RECEPTOR	IFI (PHOENIX)	POTENCY RATIO	EC_{50} AS × REFERENCE
H · Heat	TRPV1	4,200	4.20	0.238
B · Zero-Bitter	TAS2R38	2,200	2.20	0.455
K · Kokumi	CaSR	1,900	1.90	0.526
U · Umami	mGluR1	1,800	1.80	0.556
A · Astringency	Salivary PRPs	1,400	1.40	0.714
S · Sweetness	TAS1R2·R3	1,100	1.10	0.909
C · Cool-Burn	TRPM8	800	0.80	1.250
P · Pungency	TRPA1	252	0.252	3.97

Read the bottom row, because it is the most revealing. Pungency at **252** lands far below the anchor — a potency ratio of 0.252, an EC_{50} roughly four times the AITC convention reference of 2.1 µM. A score under 1,000 is not a shortcoming; it is the formula recording that this design is gentler on this axis than the reference. Each row scales against its own defined reference constant — the heat reference, the pungency reference, the kokumi reference are distinct conventions — so IFI values across axes are not directly comparable. What the scorecard shows, axis by axis, is the design decision made molecule by molecule, and that is exactly what it is meant to show.

Cooperativity, and the gift of steepness

Because CaSR works as a dimer, with two flytraps closing in concert, its dose-response is steeper than a simple one-site receptor's — and steepness, it turns out, is a gift to whoever is trying

to engineer the axis. The steepness is captured by the **Hill coefficient** *n*, and for CaSR it runs above one, typically between two and three. Fractional activation follows the Hill equation:

$$\theta = \frac{[A]^n}{\mathrm{EC}_{50}^n + [A]^n}$$

Set *n* = 2.5, a fair value inside the receptor's range, and watch what it buys. A non-cooperative receptor (n = 1) climbs from 10% to 90% activation only over an **81-fold** change in concentration — a long, lazy ramp. A cooperative dimer at n = 2.5 covers that same 10%-to-90% span over just **81^(1/2.5) ≈ 5.8-fold**. The receptor switches from quiet to loud over a concentration range a dozen times narrower. It behaves less like a dimmer and more like a threshold.

For the designer, the decisive number is the *local slope* at the working point. The sensitivity of activation to a fractional change in co-agonist is dθ/d(ln[A]) = n·θ·(1−θ), which at the midpoint (θ = 0.5) equals 0.25·n. For a cooperative dimer at n = 2.5 that is **0.625**; for a one-site receptor it is **0.25** — a 2.5-fold larger payoff for the same effort. Concretely: near the midpoint, nudging the co-agonist concentration up by a modest 10% (Δln[A] ≈ 0.095) raises activation by about 0.625 × 0.095 ≈ **6 percentage points**, where a non-cooperative receptor would yield barely 2.4. So a cultivar engineered to carry a little more kokumi peptide delivers rather more perceived richness than the raw increase in peptide would suggest. The receptor amplifies our effort. We put a small, clean change into the plant, and the dimer's cooperativity hands back a disproportionate gain in the mouth. Nothing in the kitchen tells you this is happening; the mathematics does.

Confidence: what a 7% CV buys

A number is only as trustworthy as its scatter, so the last piece of arithmetic is the error. IFI-K is a ratio of two predicted EC_{50} values, and for a ratio the relative errors add in quadrature:

$$\mathrm{CV}(\mathrm{IFI}) = \sqrt{\mathrm{CV}(\mathrm{EC}_{50}^{\mathrm{ref}})^2 + \mathrm{CV}(\mathrm{EC}_{50}^{\mathrm{cand}})^2}$$

Our engineered lines are designed to a batch-to-batch coefficient of variation of **≤7%**, the same design-precision floor we hold across every axis. Put 7% on both EC_{50} predictions and the ratio carries CV(IFI) = √(0.07² + 0.07²) ≈ **9.9%**. On Phoenix's kokumi score that is 1,900 × 0.099 ≈ **±190** at one standard deviation — a reported **IFI-K 1,900 ± 190**, or a 95% interval of roughly ±370. That is a number you can rank a cultivar by.

Now run the same arithmetic on an unengineered line, which scatters at the wild-type's **42% CV**: CV(IFI) = √(0.42² + 0.07²) ≈ **43%**, which puts the very same score at 1,900 ± 810. A kokumi reading that loose cannot tell Phoenix apart from a cultivar a third weaker — it is folklore with a decimal point. The ≤7% CV is not housekeeping; it is the thing that makes the Index legible at all. And where we need a score tighter still, the standard error falls as the square root of the replicate count: to pin IFI-K to ±2% we would average √m such that 0.099/√m ≤ 0.02, i.e. m ≈ 25 independent batches. Precision is not luck. It is replication, counted and paid for.

The kokumi larder

The peptides are not exotic. They are the reason certain foods have always tasted of care, and in our hosts they cluster, unsurprisingly, in the alliums, the pulses, and the ferments — the backbone of nearly every vegetarian kitchen worth sitting at.

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

Twelve hundred and forty constructs across five organisms, every one of them a plant or a microbe that makes its γ-glutamyl peptides for its own cellular reasons and that we have learned to coax into making more. Garlic and onion carry kokumi beside pungency, which is why an allium read raw and an allium read slow-cooked are almost different foods — the pungency mellows, the kokumi stays. Soybean and yeast pair kokumi with umami, the cousin axis, which is the molecular confession behind every great miso, every long-brewed shoyu, every spoonful of yeast-rich broth: two related receptors, lit together, reinforcing each other on the same taste cell. Vanilla is the quiet surprise on the list — its complexity is not only vanillin but a kokumi body that rounds the whole, which is part of why real vanilla and its imitation are not close. That these are hard and frontier organisms to engineer is the point: the richest, most generous savory depth is also the most demanding to build cleanly, and we would rather spend the difficulty than fake the result.

Proving the lock reads the key

We can prove the receptor is doing the work by silencing it, and the experiment is the one that moved kokumi from folklore to mechanism. In the sensory lab a trained panel evaluates a kokumi-enhanced broth, rating its mouthfulness and its continuity; then the same broth is re-evaluated after a specific **CaSR antagonist** has been introduced into the assay. The kokumi sensation collapses. The mouthfulness and the continuity fall away, and the panelists rate the two samples — identical in every ingredient — as suddenly, flatly, the same. Block the lock and the key does nothing.

That is a clean knockout, and it does two things at once. It shows that the sensation is not some diffuse property of a rich liquid but the specific output of a single, nameable receptor; and it grounds the assay our predicted EC_{50} values are referenced against, because the antagonist-sensitive fraction of the response *is* the kokumi signal, cleanly separable from everything else happening in the mouth. Every score in this chapter is anchored to that separability. Without the ability to turn the receptor off, we would have richness; with it, we have a defined target — something precise enough to design against.

Two sites, and the line we keep

One structural fact remains, because it matters when we come to drugs and to labels. The 7TM bundle — the part below the flytrap — has its *own* small binding pocket, entirely distinct from the flytrap's cleft. This is where the pharmaceutical **calcimimetics** bind: cinacalcet and its relatives, drugs used to turn CaSR up or down in diseases of calcium regulation. They are potent positive allosteric modulators, and they are absolutely not food. We name them only to be exact about the architecture, because the architecture draws a bright line we hold to: there are two ways to make CaSR more sensitive — the flytrap site and the 7TM site — and the kokumi axis lives *entirely* in the first.

Our molecules are peptides a cell assembles from glutamate, cysteine, and glycine, through the γ-bond, at the cost of two ATP. They bind in the flytrap, where food has always bound. They have been in the human diet since the first allium was dug out of the ground, and the receptor they speak to just happens to be one of the most eloquent in the body. We do not go near the drug pocket, not because we could not, but because the whole ethos of this work is to feed the body what it already knows how to read. Clean living is not a slogan here; it is a binding site we decline to touch.

Honest limits

Rigor means naming where the number stops, and kokumi has more edges than most axes. The first is structural to the sensation itself: because kokumi contributes no taste of its own, it can never be measured in isolation the way a sweetness or a heat can — it is always read as the *modulation* of other tastes, which is why the antagonist knockout, not a solo tasting, is the honest assay. The second is biological variability. CaSR is a polymorphic gene; people carry subtly different receptors, and a co-agonist that half-fills one person's flytrap at a given concentration may need a little more or less for another's. Our EC_{50} predictions are computed relative to a reference receptor, and the Index is a population instrument, not a promise about any single mouth.

The third is the gap between the dish and the dish-in-the-assay. An EC_{50} predicted for cells in a well is the cleanest possible version of a process that, in a real mouthful, competes with temperature, texture, dilution by saliva, and a dozen other molecules crowding the same taste cell — so the cell-assay score is a faithful *rank*, not a literal prediction of perceived intensity bite for bite. And the fourth is saturation: a flytrap can only close so far. Past the top of the dose-response curve, more peptide buys no more signal, and the generous steepness that rewards a small increase near the midpoint flattens to nothing near the ceiling. There is a loudest kokumi a receptor can report, and piling on peptide beyond it is waste, not depth. We predict that ceiling and design beneath it, because the most reverent thing you can do with a signal this beautiful is to know exactly how much of it the body can actually hear. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 3 · Scoring a modifier — *the mathematics of continuity.*

Every other axis in the Index measures a flavor you can taste alone. You can put capsaicin on the tongue and feel heat; you can put a sweetener on the tongue and feel sweet; the EC_{50} measurement is clean because the stimulus and the sensation line up, one molecule speaking and one receptor answering. Kokumi breaks that convenience. It has no sensation of its own to measure. Taste a kokumi compound on its own — a few grains of a γ-glutamyl peptide on a wet fingertip — and almost nothing happens; there is no sweetness, no salt, no heat, barely a whisper. Yet drop that same peptide into a simmering dal, a bowl of kombu dashi, a spoon of miso loosened in warm water, and the whole thing deepens, lengthens, gathers itself into something that tastes cared for. How do you put a number on a thing that is only visible in the presence of other things? This chapter is the answer, and the answer matters more here than on any other axis, because the measurement is the reason the axis can be designed at all. You cannot engineer toward a target you cannot read.

The receptor, and why it cannot be measured alone

To score kokumi honestly, you have to begin with the machine that senses it. Kokumi is the voice of a single receptor — the **calcium-sensing receptor, CaSR** — and that receptor is unlike the ones that read the other axes. TRPV1, the heat channel, is a pore that opens and lets ions through; it is essentially a gate. CaSR is a Class C G-protein-coupled receptor, and it is built more like a hinged hand than a gate. Two identical protein chains fold together into a dimer, and each chain carries, reaching out into the space around the cell, a large bilobed domain shaped like a Venus flytrap. Between the two lobes of that flytrap sits a cleft. When the cleft fills and the flytrap closes, the motion travels down through a cysteine-rich stalk into the seven-helix bundle buried in the membrane, and the receptor switches on.

What closes the flytrap, in the body, is calcium — the receptor's whole first job is to read the concentration of calcium ions in the fluid around the cell, which is why it is called the calcium-sensing receptor. But the cleft has room for more than calcium, and here is the molecular fact the entire kokumi axis rests on: **γ-glutamyl peptides bind in that cleft and hold the flytrap further closed.** They do not open it from nothing. They are not the thing the receptor is listening for. They are a thumb pressed gently on the hinge, making it easier for whatever calcium is already present to close the hand the rest of the way. In the language of pharmacology they are *positive allosteric modulators* — they tune the receptor's sensitivity rather than driving it directly. Glutathione does this. So does γ-glutamyl-valyl-glycine, far more strongly. So, more faintly, does a long tail of related peptides that collect in aged, slow-cooked, and fermented foods.

When the flytrap closes, the transmembrane bundle couples to its G-protein, Gq, which switches on the enzyme phospholipase C; that enzyme cleaves a membrane lipid to release a messenger, IP_3; and IP_3 opens stores inside the cell that flood the cytoplasm with a pulse of calcium. That internal pulse — not the external calcium the receptor was reading, but the cell's own answering surge — is the signal. It is what the taste cell reports to the nerve, and it is what, in the lab, we can make visible and measure. This is also, exactly, why kokumi feels the way it feels: it is a *modifier* at the receptor level before it is ever a modifier at the table. It changes how loudly the mouth hears a savoriness that is already there. The sensation of mouthfulness, of a broth that fills the whole mouth and keeps filling it, is the perceptual shadow of a receptor being held a little more closed than it would otherwise be.

The first instrument — the receptor, read in calcium

We score kokumi on **two instruments at once, and require them to agree.** The first is the receptor itself.

We express human CaSR in a cultured cell line and load those cells with a fluorescent dye whose brightness tracks the calcium concentration inside the cell. Then we apply a kokumi compound across a graded series of concentrations — a dilution ladder, each step a fixed fraction of the one before — and watch the dye. At low concentration, nothing; the cells sit dark. As the concentration climbs, the internal calcium pulse grows, the dye brightens, and plotted against the logarithm of concentration the brightness traces the classic sigmoid: a flat floor, a steep rising middle, a flat saturated ceiling. That curve is the molecule's whole conversation with the receptor, written out.



Reading the calcium flux on the microscope stage — the first of the two instruments, where a kokumi compound's dose–response curve is traced grain by grain.

The shape of the sigmoid is described by the Hill equation, which is the quiet workhorse of every axis in the Index:

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

Here *R* is the response — the height of the calcium signal — at ligand concentration *[L]*; *R_max* is the saturated ceiling; *n* is the Hill coefficient, which sets how steeply the curve rises; and **EC_{50} is the concentration that produces exactly half of the maximal response.** Set *[L]* equal to EC_{50} and the arithmetic collapses to R_max/2 no

matter what *n* is — that is the definition made visible. The EC_{50} is the single number that says how loudly this molecule speaks to this receptor. A low EC_{50} means the molecule reaches half-activation at low concentration: a loud molecule. A high EC_{50} means you need a great deal of it to be heard: a quiet one.

A worked reading makes the curve concrete. Take glutathione, whose measured CaSR EC_{50} is 5,000 µM, and read its response at three points, using the simplest case $n = 1$ to show the arithmetic plainly (the real CaSR curve is steeper, because a dimer with two cooperating flytraps gives *n* greater than 1):

- at 5,000 µM, exactly the EC_{50}: R = 5,000 / (5,000 + 5,000) = **0.50 R_max**, half-activation, by definition;
- at 1,250 µM, a quarter of the EC_{50}: R = 1,250 / (5,000 + 1,250) = **0.20 R_max**;
- at 15,000 µM, three times the EC_{50}: R = 15,000 / (5,000 + 15,000) = **0.75 R_max**.

Four-fold below the midpoint the receptor is barely answering; three-fold above it the signal is already three-quarters of the way to its ceiling. That compression — a whole world of sensation packed into a narrow band of concentration around the EC_{50} — is why the midpoint, and not the floor or the ceiling, is the honest place to anchor a score.

There is one subtlety we owe the reader, because kokumi compounds are allosteric and that changes what an EC_{50} even means. A positive allosteric modulator has no fixed potency in a vacuum; its effect depends on how much calcium is already present for it to amplify. Read against a low calcium background it looks faint; read against a high one it looks loud. So before we can report a single number, we fix the stage: every kokumi compound is measured against the same physiological extracellular calcium tone — roughly the calcium a taste cell actually bathes in — and its half-maximal concentration is read against that fixed background. We could instead report the *fold-shift* the peptide produces in the calcium curve, and in our internal records we track both. But the published number is the fixed-background EC_{50}, for one disciplined reason: every axis in the Index is an inverse-EC_{50}, and kokumi must resolve to the same kind of number as heat and sweetness and umami, or a tomato and a clove of garlic could not be read on one page.

The second instrument — the mouth

The receptor number is necessary and it is not sufficient. The second instrument is the mouth, and it has the final word.

Because kokumi is a modifier, we measure it as modification. A trained panel rates a standardized base — a defined umami-and-salt solution, chosen precisely because kokumi needs something to act *upon* — once plain and once with the compound added, scoring three things: **mouthfulness** (how fully the taste occupies the mouth), **continuity** (how it holds together from front to back), and **complexity** (how many layers it seems to carry). And they trace the

full **time–intensity curve**, pressing a lever second by second to record not just how strong the sensation is but how it rises, holds, and falls across the whole time it lives in the mouth.

The two instruments are then required to track each other. When they agree, we have a number. When they diverge, something is wrong, and we have learned to trust the disagreement as a signal rather than smooth it over. A compound that lights up the cell assay but goes silent in the mouth is almost always failing somewhere between the two: it is being cleaved by the enzymes in saliva before it reaches the taste cells, or it is bound up by the food matrix and never delivered, or it simply cannot cross the mucous layer at the concentration a recipe allows. These are the real failure modes of kokumi design, and the cell assay alone would never catch them — the receptor, in its clean dish, is a more forgiving listener than the tongue. So every IFI-K score for our designs in this book is a predicted EC_{50}, mapped onto the scorecard the two-instrument framework defines. A score no cook could taste would be no score at all.

The formula, worked

With both instruments agreeing, the score is an inversion. Every axis in the Index reads the same way — loudness is the inverse of the concentration it takes to be heard, scaled against a fixed reference:

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

The reference is glutathione, the first kokumi substance ever identified, and we define its half-maximal CaSR activation as **IFI-K 1,000**. Because its EC_{50} is 5,000 µM, the formula for kokumi reduces to a single clean constant — IFI-K(c) = 5,000,000 / EC_{50}(c), with EC_{50} in micromolar — and the whole axis becomes legible at a glance. Divide five million by the concentration a molecule needs, and you have its score.

Take γ-glutamyl-valyl-glycine, the loudest of the γ-glutamyl peptides we work with, the one that collects in long-aged and deeply fermented foods and does more with less than anything else on the axis. Its measured CaSR EC_{50} is 390 µM. Its pure-compound score is then:

$$\text{IFI-K} = 1000 \times 5{,}000 \,/\, 390 = 5{,}000{,}000 \,/\, 390 \approx \mathbf{12{,}800}$$

Nearly thirteen times as loud as the reference — half-activating the receptor at a concentration thirteen times lower. The inversion also runs backward, and running it backward is how you read the scale in your head: a molecule's required EC_{50} is simply 5,000,000 divided by its score. The table that follows reads the axis both ways, with the two measured anchors — glutathione and γ-Glu-Val-Gly — set among the round scale-points so the arithmetic has something true to stand on.

IFI-K SCORE	REQUIRED CASR EC_{50} (MM)	WHAT SITS HERE
1,000	5,000	glutathione — the reference, *measured*
2,500	2,000	scale point
5,000	1,000	scale point — and the scorecard ceiling
10,000	500	scale point
12,800	390	γ-glutamyl-valyl-glycine — *measured*

The units matter and are worth saying plainly: EC_{50} is a molar concentration, micromolar here, millionths of a mole per litre. The score it produces is dimensionless — a ratio of two concentrations — which is exactly what lets it live on the same page as a Scoville number or a sweetness index. We strip the chemistry of its units deliberately, so that loudness, and only loudness, is what crosses from one axis to the next.

Two numbers that must never be confused

Here is the distinction that took us the longest to make clean, and the one we ask the reader to carry out of this chapter most of all. There are **two different things an IFI-K number can describe,** and confusing them is the single most common way to misread the Index.

- The IFI-K of a **pure compound** is its intrinsic loudness — the inverse-EC_{50} ratio, a property of the molecule and nothing else. γ-Glutamyl-valyl-glycine is a loud molecule whatever it is dissolved in; its 12,800 travels with it.
- The IFI-K of a **cultivar or a food** is the *effective* loudness that food delivers on the kokumi axis at the concentrations its compounds actually occur — a composite of which γ-glutamyl peptides the organism makes and how much of each. This is the number on a scorecard, and it is mapped onto the Index's 0–5,000 scale.

The composite is a weighted sum. A food carries many γ-glutamyl peptides at once, each at its own concentration c_i and its own potency, and potency is just the inverse of that peptide's EC_{50}. The food's effective loudness scales with the sum of concentration times potency across all of them —

> $K_{effective} \propto \Sigma_i \left(c_i / EC_{50,i} \right)$

— and that sum is then passed through a saturating map onto the scorecard, so that no food can report past the ceiling. This is why a food can be loud without being pure. Glutathione alone, at its reference concentration, reads 1,000. A food made of nothing but γ-Glu-Val-Gly at that same molar concentration would read near its full compound loudness. Every real food lands somewhere between the two, pulled upward in proportion to how much of the louder peptides it happens to carry and how densely it carries them.



Engineered garlic in cross-section — dense in kokumi peptides, and the reason a food's score and a molecule's score are not the same number.

Our engineered garlic is the cleanest illustration we have. Our model predicts **IFI-K 3,200** — a magnificent number for a food, more than three times the reference, the savory foundation of nearly every cuisine made louder and longer. And yet 3,200 is a quarter of 12,800. The garlic is not failing; it is simply a food and not a flask. It is dense in kokumi peptides — that is what 560 constructs of patient work bought — but it is not saturated with the single loudest one, and the scorecard reports what the clove delivers at the table, not what its best molecule could deliver in a dish. The cultivar score is the one that matters when you cook, and it is the one our engineering moves. This next table is the kokumi axis as a cook meets it — the effective loudness of real foods, every value as our models compute it:

FOOD / CULTIVAR	EFFECTIVE IFI-K	WHAT IT IS
A natural stock, simmered well	700	glutathione and its kin, as a good cook reaches by instinct
Glutathione reference base	1,000	the anchor — the axis's natural zero-point
Engineered broth base	2,400	our savory foundation, kokumi lifted and held
Engineered garlic	3,200	560 constructs; dense in γ-glutamyl peptides
Scorecard ceiling	5,000	where perception caps, and the scale with it

Our full-spectrum flagship, designed to register on all eight axes at once, carries a predicted kokumi reading of **K 1,900** sitting among its other seven, computed the same way the garlic's 3,200 was: an EC_{50} estimate mapped onto the scorecard.

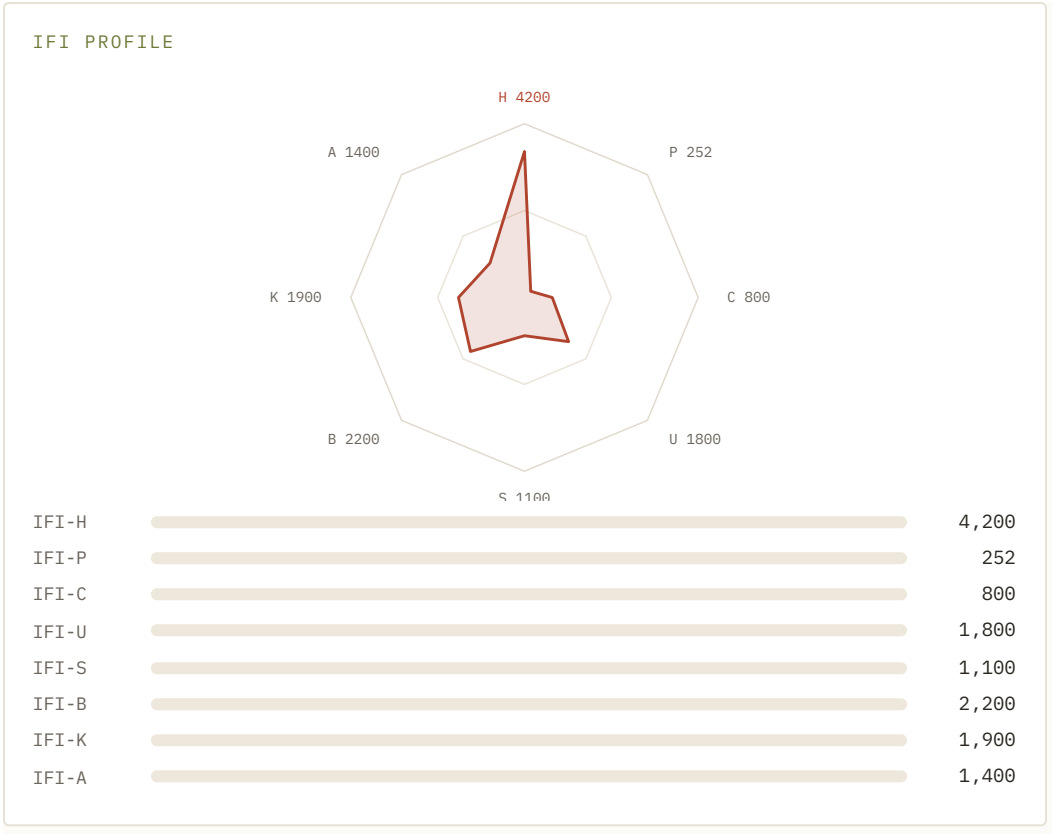

The ceiling, and why perception caps

The scorecard ceiling is 5,000 — the per-axis normalization maximum for the whole Index — and it is a deliberate choice with a perceptual, not a bureaucratic, basis. The sigmoid is the reason. Above roughly the concentration that maps to 5,000, the receptor is approaching the flat top of its dose–response curve: the flytrap can close no further, the internal calcium pulse can ring no faster or brighter. Return to the Hill arithmetic and the point is already visible — three-fold above the EC_{50} the signal is at three-quarters of its ceiling, and each further doubling of concentration buys a smaller and smaller sliver of additional response. Past a point, more loudness in the chemistry simply stops translating into more sensation in the mouth, because the receptor has nothing left to give.

So we cap the scorecard where perception caps. A scale that kept climbing into that flat region would be reporting chemistry the mouth can no longer hear — a vanity number, true on paper and false at the table. This is the clean resolution of the apparent paradox that a pure compound can score 12,800 while no cultivar scorecard ever exceeds 5,000. The first is a property of the molecule, measured in the unsaturated middle of its curve where every increment still counts. The second is a report on an experience, and the experience saturates. One is what the chemistry *could* say; the other is what the tongue can *still hear.* We hold the scorecard to the second, always.

The mathematics of continuity

Continuity is the sensation the Index is proudest to have captured, because it is the one cooks have always chased and never been able to name — the way a great broth does not so much taste strong as *stay*. It can be put in numbers, and the numbers are gentler and more surprising than you would guess.

Model the felt intensity of the sensation over time as a decay:

```
I(t) = I₀ · e^(−t/τ)
```

where I_0 is the peak height the moment the food is in the mouth and τ (tau) is the time constant of clearance — how slowly the sensation falls away. A kokumi substance does two things to this curve, and they are unequal. It raises I_0 a little. And — far more importantly — it **lengthens τ**, slowing the decay, because at the receptor the modulator is holding CaSR in its active, closed-flytrap conformation longer, letting the calcium signal persist after the first pulse instead of falling silent at once.

The whole felt experience is not the peak but the *area under the curve* — the sensation summed across all the time it lives in the mouth. For a decaying exponential that area has a beautifully simple form:

$$\int_0^\infty I_0 \cdot e^{(-t/\tau)}\, dt = I_0 \cdot \tau$$

The total is just the peak times the time constant. Now watch what a modest kokumi lift does to it. A reference stock might carry $I_0 = 1.0$ (arbitrary units) and $\tau = 4$ seconds, for a felt total of $1.0 \times 4 =$ **4.0**. A kokumi-enhanced stock raises the peak only a little, to $I_0 = 1.2$, while lengthening the tail to $\tau = 9$ seconds — and the felt total becomes $1.2 \times 9 =$ **10.8.** A twenty-percent taller peak and a longer ring together have not added twenty percent; they have **nearly tripled the whole experience** (10.8 / 4.0 = 2.7×). Said another way, the sensation's half-life — the time $\tau \cdot \ln 2$ for it to fall to half strength — stretches from $4 \times 0.693 \approx 2.8$ seconds to $9 \times 0.693 \approx 6.2$ seconds. The taste lasts more than twice as long.

STOCK	PEAK I_0	T (S)	FELT TOTAL ($I_0 \cdot T$)	HALF-LIFE (S)
Reference	1.0	4	4.0	2.8
Enhanced (peak only, same τ)	1.2	4	4.8	2.8
Enhanced (longer tail)	1.2	9	10.8	6.2
Enhanced (longest tail)	1.2	12	14.4	8.3

Read down the last two columns and the whole philosophy of the axis is there. In the integral, peak and tail have equal mathematical leverage — both multiply. But they are not equally available to the designer. The peak is bounded; it runs into the ceiling, the saturation of the receptor, and there is only so much taller it can go. The tail is not bounded in the same way — you can almost always slow the clearance further, hold the receptor a little longer, and every second of τ pays out in full. This is the arithmetic behind the word *generous,* and it is no accident that the word is a kindness: on the kokumi axis the richness lives in the tail, and a small investment in persistence returns large in the sum. It is also why two foods with the same peak can feel so different — the one with the longer τ is the one that stays with you, the one you are still tasting when you reach for the next spoonful. When an engineered broth base is predicted at IFI-K 2,400 against a natural 700, a good part of that climb is not a taller peak at all. It is a longer ring: the receptor, held in its oscillating state, taking its time to fall quiet.

Error, and the discipline of a number you can trust

A number is a form of respect for the reader, and respect means saying how well you know it. The IFI-K is a ratio of two measured concentrations, and the uncertainty of a ratio follows a rule as old as measurement: relative errors add in quadrature. If the reference EC_{50} is known to a fractional error of CV_ref and the candidate's to CV_c, then the fractional error of the score is their root-sum-square:

```
CV(IFI-K) = √( CV_ref² + CV_c² )
```



The bench where a score is earned and re-earned — EC_{50} measured across replicate runs until the number stops moving.

The lesson is in the worked numbers. Suppose both EC_{50}s are held to our engineered batch consistency of 7%. Then CV(IFI-K) = $\sqrt{(0.07^2 + 0.07^2)} = \sqrt{0.0098} \approx 0.099$, very nearly **10%.** γ-Glutamyl-valyl-glycine's 12,800 is really 12,800 ± about 1,300 — a number you can design against. Now suppose instead the candidate were as variable as a wild-type plant, which swings at a CV of 42% batch to batch. Then CV(IFI-K) = $\sqrt{(0.07^2 + 0.42^2)} = \sqrt{0.1813} \approx 0.426$, almost **43%.** The same 12,800 becomes 12,800 ± 5,400 — a number that means almost nothing, a score so loose it cannot tell a good batch from a poor one. This is why the ≤7% batch CV is not a vanity metric and never was. It is the thing that makes the published number *mean* something. Precision at the plant is what lets the decimal point survive the trip to the table.

There are two kinds of error braided together here, and honesty requires keeping them apart. The first is assay precision: any single dose–response fit carries fit error, and we beat it down by replication — the standard error of a mean EC_{50} shrinks as the square root of the number of runs, SE = SD / $\sqrt{n}$, so we run the ladder again and again until the number stops moving. The second is production precision: the cultivar's actual chemical output varies from harvest to harvest, and that is the error that actually reaches the eater. The assay can be made arbitrarily tight; the biology is where the real discipline lives. A food whose chemistry swings 42% from batch to batch cannot carry a stable score no matter how exquisite the assay, and a food designed to ≤7% can — which is the whole reason we chase consistency as hard as we chase loudness.

And then the honest limits, named plainly, because a measurement you cannot break is a measurement you do not understand. The panel is human: continuity and mouthfulness are rated by trained people, and training narrows their spread but never closes it to zero.

Saliva varies from person to person and hour to hour, and saliva is where fragile peptides go to be cleaved. The food matrix can bind a compound and withhold it, so that the same peptide scores differently in a clear broth than in a fatty sauce. We do not pretend these away. We bound them — fixed calcium background, standardized base, replicate assays, a panel held to agreement — and then we publish only the receptor number the mouth confirms. Where the two instruments will not agree, we report no score at all, and we go back to the bench. That refusal is not a weakness of the method. It is the method.

The anchor

Every axis needs a reference, a fixed point where the scale reads exactly 1,000, and the choice of that point is a statement of values. For kokumi we anchor to **glutathione** — the first kokumi substance ever identified, a simple tripeptide of glutamate, cysteine, and glycine, present in nearly every food that carries the sensation at all. It is the natural zero-point of the axis. We define IFI-K 1,000 as the half-maximal CaSR activation produced by our glutathione reference standard, and every other molecule and every cultivar is read against it.

It is a humble anchor on purpose. Glutathione is the kokumi of a decent homemade stock, the quiet depth a good cook reaches by instinct after an afternoon at the stove — nothing exotic, nothing you could name on the tongue, just the difference between water with vegetables in it and a broth that tastes like being fed. We set the scale's thousand-mark there, at the reach of patience and attention, so that every number above it measures something real that patience alone could not buy. Read the scorecard of the reference standard itself — glutathione in a clean umami base, the natural zero-point of the whole axis:

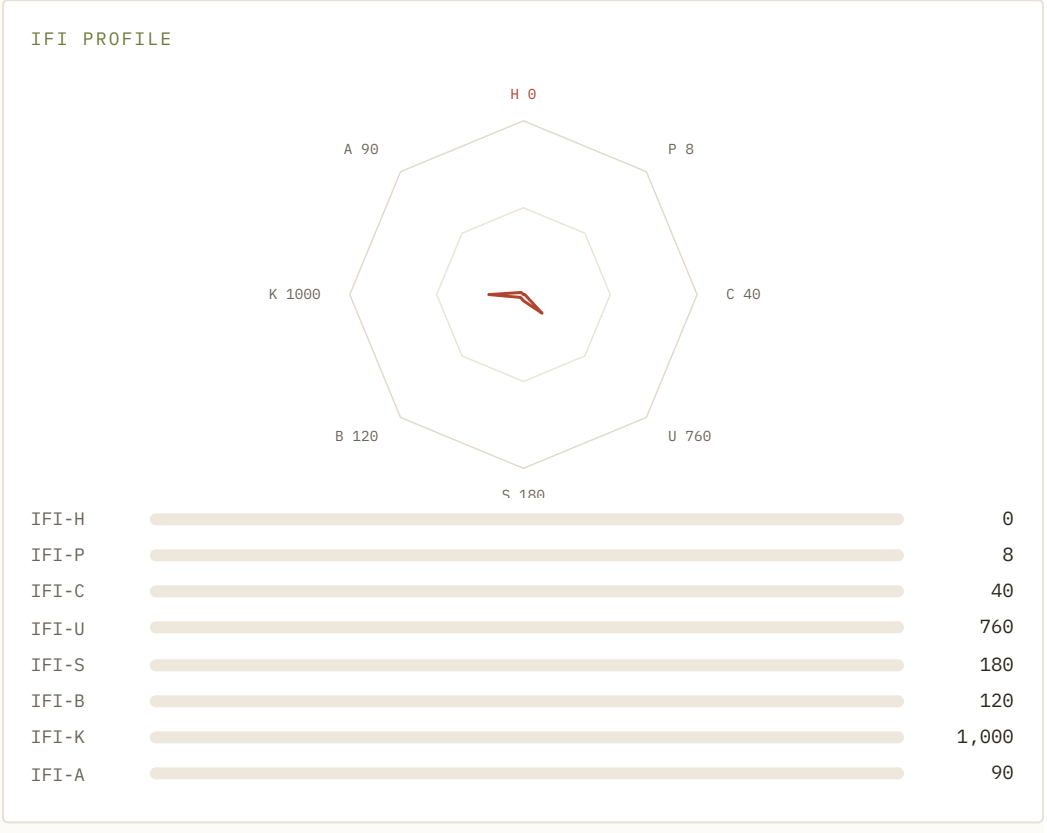

Almost silent on heat, on pungency, on cool and astringency — a trace, a rounding error. A modest savory warmth on umami, because kokumi needs a savoriness to deepen and a clean base supplies a little. And then K at 1,000, the needle resting exactly on its anchor. Everything that follows in this book is the work of making that single

number climb — gene by patient gene, peptide by peptide, second of τ by second of τ — without ever once making the food taste of anything you could name. That is the whole quiet art of the axis: to be felt everywhere and found nowhere, to make a broth more generous without making it taste of more. We learned to model it so that we could, at last, learn to design it. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.



PART II OF V

The Molecules of Richness

the γ-glutamyl family · the sulfur cycle · designing the pathway

Chapter 4 · The γ-glutamyl family — *the peptides that speak to calcium.*

Open the back of any food that tastes as though it was made slowly — a Gouda aged past its first year, a soy sauce drawn after a long dark ferment, roasted garlic gone soft and sweet, a spoon of yeast extract scraped from the jar, a fish sauce finishing a broth — and you are reading, without knowing it, a library of very small peptides built on a single structural trick. They are the **γ-glutamyl peptides**: chains of two or three amino acids whose entire character turns on the bond that opens them. There is no flavor molecule we love more and understand better, and none that rewards a close look the way this one does, because in this family the structure and the sensation are the same fact told twice.

The bond that moves one receptor

Glutamate is unusual among the common amino acids: it carries two carboxyl groups, not one. There is the **α-carboxyl** on the backbone — the carboxyl every amino acid has, the one that forms an ordinary peptide bond — and a second, the **γ-carboxyl**, at the far tip of its three-carbon side chain. An ordinary peptide links head to tail through the α-carboxyl. A γ-glutamyl peptide is joined through that side-chain γ-carboxyl instead. One bond, moved from the backbone to the end of the side chain, and the consequence is total: the glutamate's own α-head — its free α-amino group and its free α-carboxyl — is left sticking out of the molecule unbonded, in exactly the arrangement of a lone amino acid floating in solution. That protruding, free α-amino acid is the part the receptor reads. Everything in this chapter follows from that one displacement.

The γ-linkage does two things at once, and they are worth separating. First, it **presents** the molecule to the receptor in precisely the geometry that reads as activation — the free α-amino and α-carboxyl of the glutamyl head seated where the receptor expects an amino acid, the rest of the peptide trailing out of the binding cleft. The α-linked version of the very same sequence, with the glutamate tied in through its backbone so that head is consumed in the bond, is nearly silent: we have measured α-glutamyl-valyl-glycine side by side with its γ-isomer in our own assay and found the α-form essentially inactive. The receptor is not reading the atoms; it is reading the *bond*. Second, the γ-linkage is **invisible to the ordinary peptidases** of the mouth and gut, enzymes that evolved to cut α-bonds and cannot touch a γ-bond. Only one class of enzyme, γ-glutamyl transpeptidase, opens it. That is why kokumi *lasts* where an ordinary flavor peptide would be dismantled in seconds — the molecule survives on the tongue long enough to keep the receptor ringing. The bond that makes it loud is the same bond that makes it slow to leave. Loud, and lingering, for one structural reason.

A flytrap tuned for small, charged things

The receptor is the **calcium-sensing receptor, CaSR**, and to design for it you have to picture it honestly. It belongs to class C of the G-protein-coupled receptors — the same structural family as the metabotropic glutamate receptors and the sweet and umami taste receptors — and like its cousins it carries, hung off the cell surface on a stalk, an enormous bilobed extracellular domain shaped like a clam, the **Venus flytrap**. Two lobes, a hinge, and a cleft between them. A ligand that settles into the cleft lets the two lobes close around it; that closing is transmitted down through a cysteine-rich neck into the seven-helix bundle buried in the membrane, which rolls against its partner and releases the G-protein inside the cell. CaSR does not work alone — it is an obligate **homodimer**, two full receptors clasped together, so there are two flytraps and the cell counts votes between them. That cooperativity will matter when we read a dose-response curve.

What closes the trap is the receptor's name made literal: **calcium**. CaSR is first and last a calcium sensor, its native agonist the bare Ca^{2+} ion. But the trap has a second, gentler input. **L-amino acids** — especially the aromatic ones, phenylalanine and tryptophan — bind in the flytrap cleft not as the primary trigger but as *co-agonists*, tuning the receptor so that it answers a given calcium level more loudly. The γ-glutamyl peptides work exactly here. Their free glutamyl head docks in the amino-acid site; the rest of the peptide reaches into the cleft; and the receptor, sensitized, reports more signal for the calcium already around it. This is why kokumi is not a taste of its own. CaSR activation does not send the brain a flavor; it sends an *enhancement* — a deepening of continuity, thickness, and mouthfulness, and an amplification of the sweet, salty, and savory signals arriving alongside it. The peptides make the meal louder without adding a note. Downstream, the activated receptor couples through Gq/11 to phospholipase C, which splits a membrane lipid into the second messengers that release a pulse of calcium *inside* the cell — and that internal pulse is what we actually measure. Our receptor panel is a calcium-reporter assay: CaSR-expressing cells loaded with a fluorescent calcium indicator, read for the brightness of their flash as we titrate each peptide across them. The number that falls out is an EC_{50}, the concentration at which a peptide drives the receptor to half of its full response.

The grammar of a loud peptide

Measure enough of these peptides and a grammar emerges — a structure-activity rule precise enough to design against. The receptor reads three things.

It reads the **head**, which must be an intact, free γ-glutamyl group. Cap the α-amino group of that glutamate, block it, dress it up, and activity collapses, because the receptor needs a bare amino acid presented in the cleft and a blocked head is not one. It reads the **length**: dipeptides and tripeptides fit; add a fourth residue and the molecule no longer seats cleanly, and potency falls off a cliff. And it reads the **second residue** — the one just past the glutamate — whose side chain is the tuning knob. The amino-acid pocket of the flytrap opens onto a small hydrophobic subpocket, and a compact,

branched, greasy side chain dropped into that subpocket is what makes a peptide loud. **Valine** fits it best of all. That single fact is the whole reason γ-Glu-**Val**-Gly outruns the rest of the family: the valine lands in the receptor's favorite pocket and the terminal glycine caps the chain at the ideal length. This is not a curiosity; it is a blueprint. When we engineer an organism to build kokumi, we are not asking it to make more peptide at random — we are asking it to load valine into the second position, to make the one shape the receptor most wants to hold.



The sulfur-rich clove, specimen-isolated — garlic holds the densest natural store of γ-glutamyl peptides in the plant kingdom, and it is where we first learned to read the family.

The receptor panel — turning an EC50 into a number

Potency across the family varies by more than fortyfold, and that variation is the designer's palette. The panel below is our in-house measurement: CaSR EC_{50} from the calcium-reporter assay, with the potency index anchored against **glutathione** — the body's own γ-glutamyl tripeptide and the family's natural reference — set to 1.0. The lower the EC_{50}, the louder the molecule.

PEPTIDE	Γ-BOND	CASR EC_{50} (MM)	MG/L AT EC_{50}	POTENCY VS GSH	IFI-K (RAW)	FOUND NATURALLY IN
γ-Glu-Val-Gly	yes	390	118	12.8	12,821 †	fish sauce, soy sauce, aged Gouda, beer
γ-Glu-Cys-Gly (glutathione)	yes	5,000	1,537	1.00 *(ref)*	1,000	yeast, garlic, wheat germ
γ-Glu-Cys	yes	4,500	1,126	1.1	1,111	garlic, onion, germinating seeds
γ-Glu-Abu-Gly (ophthalmic acid)	yes	6,250	1,808	0.80	800	ferments, plant tissue
γ-Glu-Val	yes	8,300	2,044	0.60	602	aged cheese, fermented legumes
γ-Glu-Ala	yes	16,700	3,644	0.30	299	legumes, sprouts
α-Glu-Val-Gly	no	>100,000	—	<0.05	<50	— *(synthetic control)*
Ca^{2+} (reference ion)	n/a	~3,500	—	—	1,429	dairy, leafy greens, hard water

† exceeds the 5,000 per-axis normalization ceiling; clamped to 5,000 on a printed scorecard.

The index that turns an EC_{50} into a score is a single ratio — **IFI = 1,000 × EC_{50}(reference) ÷ EC_{50}(candidate)** — and it is worth running by hand, because the arithmetic is where the design decisions actually live. Take the loudest member first. γ-Glu-Val-Gly against the glutathione reference: 1,000 × 5,000 ÷ 390 = **12,821**. Glutathione against itself is the sanity check every scale needs — 1,000 × 5,000 ÷ 5,000 = **1,000**, the anchor sitting exactly where an anchor should. γ-Glu-Val lands at 1,000 × 5,000 ÷ 8,300 = **602**; γ-Glu-Ala at 1,000 × 5,000 ÷ 16,700 = **299**. And the α-isomer, the control that proves the bond is everything: 1,000 × 5,000 ÷ 100,000 = **50**, and because its EC_{50} is only a lower bound, even 50 is an overestimate. Moving one bond from the side chain back to the backbone costs the molecule a factor of 12,821 ÷ 50 ≈ **256** in loudness. The receptor is reading a single covalent choice and scoring it a quarter of a thousandfold.

Two more readings earn their place. Calcium itself, the receptor's native ligand, computes to 1,000 × 5,000 ÷ 3,500 ≈ **1,429** — reassuringly near the reference, which is the right answer: the ion the receptor was built to sense should sit in the same neighborhood as the peptide we anchor to. And the headline number, γ-Glu-Val-Gly at 12,821, overflows the Index. Our scorecards normalize each axis to a ceiling of 5,000, so on a printed profile the raw score is clamped: the single loudest kokumi peptide food has ever been found to contain is loud enough to saturate its own axis. That is not a flaw in the molecule; it is the Index telling you where the edge of the useful range sits.

A micromolar figure is hard to hold in the hand, so the panel carries a mass column too, and that conversion is pure chemistry you can check. γ-Glu-Val-Gly is three residues — glutamate (147.13 g/mol), valine (117.15), glycine (75.07) — joined by two peptide bonds, each shedding a water (18.02): 147.13 + 117.15 + 75.07 − 2(18.02) = **303.3 g/mol**. Its EC_{50} of 390 µM is therefore 390 × 10^{-6} mol/L × 303.3 g/mol ≈ 0.118 g/L, about **118 mg/L**. Glutathione, heavier at 307.3 g/mol and far quieter at 5,000 µM, needs **1,537 mg/L** to reach the same half-activation — thirteen times the mass for the same effect. The practical meaning of a potency index is exactly this ratio of masses: 1 ÷ 12.8 = 0.078, so **7.8%** as much γ-Glu-Val-Gly, by weight, matches the richness of a full dose of glutathione. Teach an organism to make the loud peptide where it would otherwise make only the quiet one, and you buy the same kokumi for a thirteenth of the material, with nothing else along for the ride.



The calcium-reporter stage — CaSR-expressing cells, a fluorescent calcium indicator, and a titration read for the brightness of the flash. This is where an EC_{50} becomes a measured fact.

The shape of a dose

An EC_{50} is one point on a curve, and the curve is where a formulator actually works. Receptor activation follows a sigmoid: near-flat at low concentration, steep through the middle, flattening again as every receptor fills. The simplest model — a Hill slope of one — gives the fraction of maximal response as $f = [L] \div (EC_{50} + [L])$, which is why the half-maximal point falls exactly at $[L] = EC_{50}$. Walk γ-Glu-Val-Gly up that curve from its 390 µM midpoint.

[Γ-GLU-VAL-GLY]	MULTIPLE OF EC_{50}	CASR RESPONSE *
43 µM	0.11×	10%
130 µM	0.33×	25%
390 µM	1.0×	50%
1,170 µM	3.0×	75%
3,510 µM	9.0×	90%

** fraction of maximal activation, modeled at a Hill slope of 1.*

Read the two ends. To climb from a tenth of full response to nine-tenths, the concentration must rise from 43 µM to 3,510 µM — a span of 3,510 ÷ 43 ≈ **81-fold**, the characteristic width of any Hill-slope-of-one curve. That width is the designer's friend and warning at once. It means kokumi is forgiving in the middle — a doubling of peptide around the EC_{50} moves you only from 50% to about 67%, so the sensation is stable against the ordinary drift of a real food. And it means the bottom of the curve is shallow: below a tenth of the EC_{50} the peptide is doing almost nothing you can taste, which is precisely why γ-Glu-Val-Gly stayed a laboratory curiosity for so long. Food holds it in traces far down the shallow foot of the curve, delicious in principle and inaudible in practice. The whole point of engineering is to move an organism up this curve, not to invent a new molecule but to make enough of the right old one to clear the knee. One honest caveat rides along: a Hill slope of one is the simple case, and CaSR is a dimer that counts votes between its two flytraps, so real curves often run steeper than this. A steeper slope narrows the 10-to-90

window — toward ninefold at a slope of two — which only sharpens the lesson, making the threshold crossing more abrupt and the dosing window tighter.

A worked plate

Put the curve to work on an actual bowl. Suppose we want a vegetable stock to sit at three-quarters of CaSR's full response — a deep, rounded richness that stops short of the flat top of the curve. Three-quarter activation falls at three times the EC_{50}, so for γ-Glu-Val-Gly that is 3 × 390 = 1,170 µM, which in weighable terms is 1,170 × 10^{-6} mol/L × 303.3 g/mol ≈ 0.355 g/L, or **355 mg/L**. A 250 mL bowl of that stock therefore carries about 355 × 0.25 ≈ **89 mg** of peptide — a quantity you could weigh on a jeweler's scale, present on the tongue as nothing but depth. Now cost the same richness in glutathione. Its three-quarter point is 3 × 5,000 = 15,000 µM; at 307.3 g/mol that is 4.61 g/L, and the same bowl needs 4.61 × 0.25 ≈ **1.15 g** — 1,150 mg against 89, thirteen times the material to reach the identical place on the curve. The ratio is exactly the potency index wearing kitchen clothes: 1.15 g ÷ 0.089 g ≈ 13. That single comparison is the whole case for the loud peptide. The same richness, the same receptor, the same curve — reached with a rounding error of peptide instead of a pinch you could see, which is the difference between a flavor that vanishes into the food and one that announces, faintly and wrongly, that something was added.

What a number owes the reader

A score with no error bar is a decoration, so it is worth being plain about where the uncertainty in IFI-K 602 actually lives. Two kinds stack up. The first is in the assay. The Index is a ratio of two measured EC_{50}s, and the fractional errors of a ratio add in quadrature: CV(IFI) = $\sqrt{(\text{CV(ref)}^2 + \text{CV(candidate)}^2)}$. Two EC_{50}s each determined to within 7% therefore yield an index good to $\sqrt{(0.07^2 + 0.07^2)} \approx$ **9.9%**, call it ten. A reported IFI-K of 602 is honestly a band of roughly 540 to 660 — which is why we never split hairs between, say, a 600 and a 620, and why we re-anchor against the glutathione reference on every plate so the reference error is shared, not compounded, across the panel.

The second kind of error lives in the plant, not the plate, and it is the one that decides whether a number means anything a year later. An engineered construct is only as good as its reproducibility, and the discipline we hold every construct to is a batch-to-batch **coefficient of variation of ≤7%** — against the 42% that raw biology hands you when you simply select a phenotype and hope. That gap is the real product. Anyone can hit a target peptide concentration once; building an organism that lands within 7% of it harvest after harvest is the difference between a measurement and a manufacturable flavor. Reverence for the plant and respect for the reader turn out to be the same discipline: a number is a promise, and a promise with a 42% spread is not one.

The quiet partner — calcium, carried in

A second mechanism rides alongside the receptor binding and deserves naming, because it explains why the kokumi of a food is never quite separable from its minerals. The γ-glutamyl peptides, glutathione most of all, are mild **calcium chelators**: their carboxylate oxygens and, in glutathione, the thiol sulfur, coordinate a calcium ion loosely in solution. Because CaSR is above all a calcium sensor, a peptide that arrives at the flytrap already carrying a Ca^{2+} ion may present both of the receptor's agonists at once — the co-agonist peptide and the primary ion, delivered together. The effect is subtle and we do not overstate it, but it is part of why calcium salts themselves register as weak kokumi (the ion is the receptor's own ligand), and part of why the same recipe reads differently in hard and soft water. That ~3,500 µM calcium EC_{50} is, in weighable terms, $3{,}500 \times 10^{-6} \times 40.08 \approx$ **140 mg/L** of calcium — the mineral load of a moderately hard water, close enough to the kitchen that a stock built on one tap and a stock built on another do not land in quite the same place on the axis. It is a small thing and a real one, and it is the sort of variable a precisely characterized flavor lets you finally see.

How the bond gets built

If the receptor reads the γ-bond, something has to write it, and in the living cell that something is **γ-glutamyl transpeptidase** — the same enzyme that, in the mouth, is the only thing able to take a γ-glutamyl peptide apart. It builds and breaks by one chemistry run in two directions. The enzyme is an N-terminal-nucleophile hydrolase, carrying a catalytic threonine whose side-chain oxygen attacks the γ-carbonyl of a donor — usually glutathione, the cell's most abundant γ-glutamyl molecule — snipping away the back end of the donor and parking the glutamyl group on the enzyme itself as a covalent **γ-glutamyl-enzyme intermediate**. What happens next is a competition for that loaded intermediate. If a water molecule reaches it first, the glutamyl group is hydrolyzed off as free glutamate and nothing is gained. But if an amino acid or a dipeptide reaches it first, the glutamyl group is handed to *that* acceptor, and a new γ-glutamyl peptide is born. The whole craft of building kokumi enzymatically lives in that fork: tilt the competition toward transpeptidation over hydrolysis, and stock the acceptor pool with the residue the receptor wants. Feed the enzyme valine and a glycine-capped acceptor and it writes γ-Glu-Val-Gly; the enzyme's own acceptor preference becomes the design lever. This is why our engineering works on the *flux* into the second position rather than on the receptor at all — we cannot make the flytrap love valine more than it already does, so instead we teach the cell to present more of the one peptide the flytrap was always going to answer.

Where the family grows, and where we grow it

The family's natural reservoirs are a map of slow food. Garlic and onion carry the densest γ-glutamyl stores in the plant kingdom, laid down as γ-Glu-Cys and γ-Glu-Val in the clove and the bulb. Yeast is thick with glutathione. Legumes sprout theirs on germination. And the loudest member, γ-Glu-Val-Gly, accumulates in the long ferments — the soy sauce drawn after months, the fish sauce aged in the barrel, the Gouda past a year, the beer left to condition — built slowly by microbial γ-glutamyl transpeptidase transferring glutamyl

groups onto valine and glycine over the patient weeks of a ferment. Our hands stay on the plant and microbial side of that map. The peptides occur in animal tissue too, and an honest panel says so, but what we build with is what plants and microbes already make by the billion on a pathway we can read and tune.



The soybean — reservoir of the long ferment, where microbial transpeptidase builds γ-Glu-Val-Gly molecule by molecule over the weeks that soy sauce and miso take to become themselves.

These are the hosts we build kokumi into, and the Index's own roadmap names them with their construct counts and the difficulty of the engineering:

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY	THE KOKUMI ROUTE
Garlic	560	K / P	Moderate	native γ-glutamyl pool — γ-Glu-Cys, γ-Glu-Val
Soybean	360	U / K	Frontier	ferment-borne γ-Glu-Val-Gly
Yeast	240	U / K	Hard	glutathione-rich; transpeptidation to the loud tripeptide
Onion	40	K	Moderate	allium γ-glutamyl precursors
Vanilla	40	K	Hard	kokumi mouthfeel beneath the vanillin

The engineering goal is the grammar of the receptor, read backward into a cell. Because the receptor's favorite shape is a valine in the second position, the enzyme work we build on aims the γ-glutamyl transpeptidase pathway specifically at loading valine there — nudging a glutathione-rich organism to make γ-Glu-Val-Gly where it would otherwise stop at glutathione, converting a quiet peptide into one thirteen times as loud on the same carbon. The payoff is a whole-food kokumi reading you can put on a profile. A full-spectrum cultivar carries it as one axis among eight:

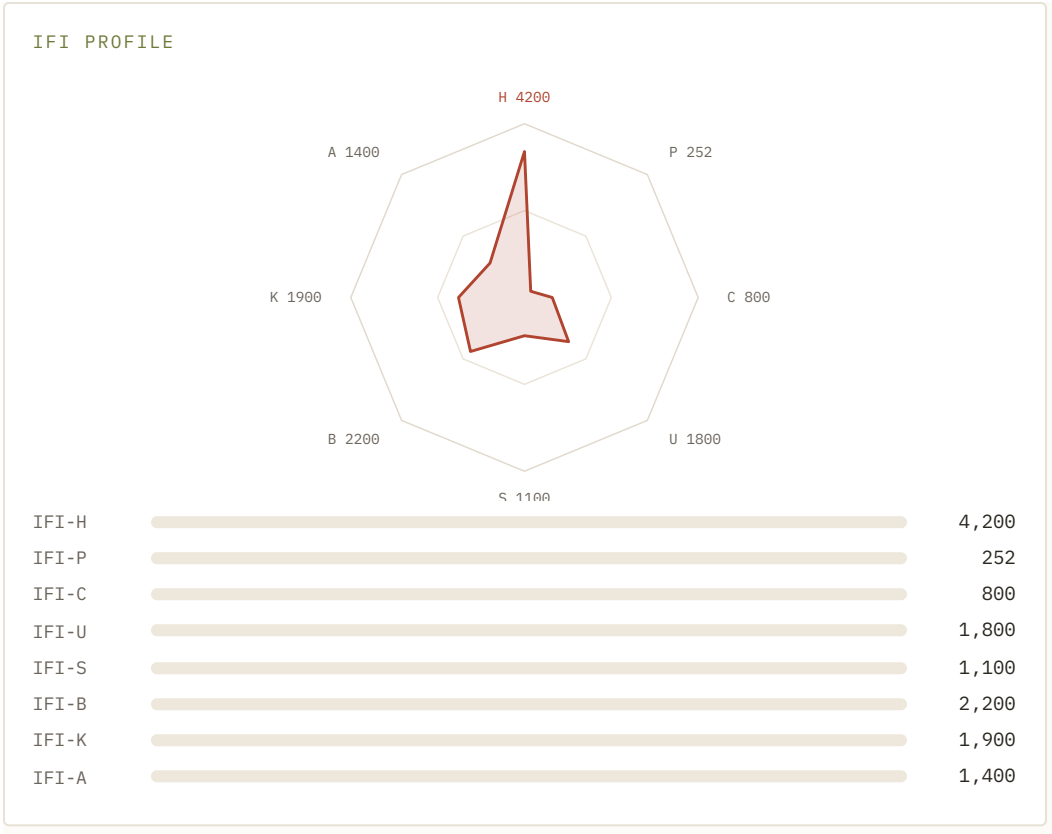

That **K of 1,900** is our model's prediction for the fruit speaking at nearly twice the glutathione reference — not a single peptide's pure loudness but the aggregate voice of the whole γ-glutamyl pool the plant has been designed to carry, computed from the same calcium-reporter framework as every molecule in the panel. The individual peptides saturate the axis; the fruit, holding them at designed and reproducible concentrations, is predicted to land at a score a formulator can actually design around.

The honest edges

Every family has limits, and naming them is how you stay trustworthy with a number. The loudest peptide overflows the scale, which means the Index stops discriminating at the very top — above an EC_{50} of about 250 µM the printed axis simply reads *maximum*, and finer ranking there is a laboratory matter, not a label one. The chelation link to calcium makes the axis genuinely sensitive to mineral content, so a kokumi score is a score *for a given water and a given salt load*, and we report it that way rather than pretend the peptide acts in a vacuum. The length rule is a hard wall, not a gentle slope: there is no reward for building bigger, because a tetrapeptide falls out of the cleft, so the entire design space is two and three residues wide. And the non-peptide agents that also move CaSR — calcium salts, the polyamine spermine, the polycation protamine — are real confirmations of the receptor's rule (it is moved by small, charged, amino-acid-like things) but minor players in food, and we do not build with them. The receptor's grammar is narrow by design, and the narrowness is what makes it designable.

The failure modes are chemical, and respecting them is part of designing honestly rather than hopefully. The thiol-bearing members — glutathione and γ-Glu-Cys — are **redox-sensitive**: left to the air, two of them oxidize into a disulfide-linked dimer that no longer presents the same clean monomer to the flytrap, so a glutathione reading taken fresh is not the reading a month of oxygen later. The free α-amino head the receptor grips is also the family's one liability in a hot pan. Protonated and positively charged at the pH of the mouth, that same amino group is a willing partner in **Maillard**

browning, and high-heat processing against a reducing sugar will quietly consume it, stripping the very handle the receptor holds — kokumi lost not to any flaw in the molecule but to the cook. And because the head must stay free and charged, strongly alkaline conditions, which pull the proton off it, dull the whole family in one stroke. None of this ends the usefulness of the peptides; these are the ordinary cautions of working with a real molecule, and naming them is the price of trusting the number.

Then there is the quiet miracle, which is also the reason this is the cleanest target in the whole Index. Look again at the panel and note what is *not* there: no bitterness, no sweetness, no flavor of any kind. At the concentrations that matter, γ-Glu-Val, γ-Glu-Val-Gly, glutathione — none of them tastes of anything you could pick out of a lineup. Most flavor molecules are compromised, each gift arriving with a tax of three things you did not order; push the compound you want and you spend the rest of the formulation hiding its side effects. The γ-glutamyl peptides give richness, continuity, and harmony, and ask nothing back. They are silent on every axis but their own. The hardest part of designing with them, then, is not managing what they do wrong — it is the purely quantitative problem of building enough of the right one, high enough up the shallow foot of that dose-response curve, reproducibly enough to promise, that a plant grown in the ground can speak to a receptor on the tongue in a language the body was always listening for. That is the work: one bond, one pocket, one valine, made loud and made honest. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 5 · The γ-glutamyl cycle — *glutamate, cysteine, glycine, ATP.*

To build more of the right peptide you have to know exactly how the plant builds any of it — down to the enzyme, the intermediate, and the single phosphate bond that each step spends. Every lever we pull to make a food richer is a lever on this one pathway, so the whole of our kokumi work rests on reading it correctly first. And the news that makes the program possible — the news we return to each time a new organism arrives on the bench — is that the core of this machinery is **universal**. Every vascular plant, every yeast, every mold, every green thing we have grown toward fullness already runs the γ-glutamyl cycle, because it runs it to make **glutathione**, and glutathione is not a flavor molecule the cell can take or leave. It is the cell's master redox buffer, the small tripeptide that stands between a living cell and oxidative ruin, among the most essential molecules life makes. We are never installing this pathway. We are turning up a pathway the organism cannot live without, and that distinction is the quiet mercy at the center of the work: to make a food richer, we lean on the same chemistry that keeps the plant well. Nothing here is grafted on. We are asking the cell to do more of what it already does to protect itself, and tasting the surplus.

Glutathione is assembled in two steps, from three amino acids, at the cost of two molecules of ATP, and those two steps are the whole beginning of the story. A third enzyme — the most interesting of the three — takes the finished tripeptide apart and redistributes it, and it is that third enzyme that turns a redox molecule into a flavor library. Before we open any of them, though, we have to say what the tongue is actually doing when it meets these peptides, because the receptor is where the number comes from, and the number is how we design.

The receptor that reads fullness

Kokumi is the quietest of the eight axes because, strictly, it is not a taste. Put a kokumi peptide on the tongue alone and it is nearly silent — no sweetness, no savor, no edge. Its gift is **modulation**: it takes a food that already has flavor and gives it body, continuity, and a mouth-filling persistence, the sense of a broth that has been tended rather than assembled, of a sauce that keeps speaking after you have swallowed. The word, borrowed from the Japanese kitchen, points at exactly the thing our family had no number for across a decade of cooking — the roundness of something made with care.

The receptor underneath it is the **calcium-sensing receptor, CaSR**, and it is worth lingering on, because its architecture explains why kokumi behaves the way it does on the palate. CaSR is a **class C G-protein-coupled receptor**, which means it does not look like the compact seven-helix receptors of sweet or bitter. It is a large, obligate **dimer**, and each half carries an enormous bilobed extracellular domain shaped like a **Venus flytrap** — two lobes hinged on a cleft — perched above the seven transmembrane helices that actually talk to the cell. In its day job CaSR reads extracellular **Ca^{2+}**:

when calcium is plentiful the flytrap closes, the dimer twists, and the cell is told the mineral is abundant. Kokumi peptides do not do this job. They bind at or near the hinge of the flytrap and act as **positive allosteric modulators** — they do not open the door themselves so much as make the existing calcium signal easier to send, lowering the threshold at which the receptor turns on. That is why kokumi reads as *enhancement* and not as a flavor of its own: it is the tongue's richness sensor being made more willing, against a background the food itself supplies.

Two consequences fall straight out of that mechanism, and we design around both. First, because the ligand is an allosteric modulator rather than a lone agonist, its effect is **context-dependent** — it needs the background it is modulating, so kokumi is always measured against a fixed, defined base, never in isolation. Second, because CaSR is a cooperative dimer, its response to a stimulus is **steep**: the dose–response curve rises faster than a simple one-site receptor's, which means the line between "flat" and "full" is a narrow band of concentration. Cross that band and the food turns round; fall short of it and nothing happens. Designing kokumi is, in large part, designing to land a food inside that band and keep it there batch after batch.



Tomato in cross-section — our highest-construct savory organism carries both the umami and the kokumi machinery, the two depth axes running on receptors from the same class of large, hinged, dimeric proteins.

To turn that band into a number we use the same instrument we use for every axis — the half-maximal activation of the receptor, predicted and inverted against a defined reference convention:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — here, half-maximal potentiation of CaSR against the defined background. A smaller EC_{50} means the molecule speaks more loudly: less of it is needed to reach the half-way mark. Invert the ratio, scale by a thousand, and loudness becomes a single

score. The reference molecule of the kokumi axis is **glutathione** itself — γ-Glu-Cys-Gly, the tripeptide at the heart of this whole pathway — so by construction the reference scores exactly:

$$\text{IFI-K(glutathione)} = 1{,}000 \times EC_{50}(\text{ref}) \div EC_{50}(\text{ref}) = \mathbf{1{,}000}$$

Everything louder than plain glutathione scores above 1,000; everything softer, below. With that anchor set, the two enzymes that make glutathione, and the one that redistributes it, become the three places we go to move the number.

Glutathione, the molecule the cell cannot do without

We treat glutathione with the reverence it earns, and not only because it is the reference. At millimolar concentrations inside nearly every living cell, it is the thiol that quenches reactive oxygen, recycles vitamins C and E, and keeps proteins folded with their cysteines in the right redox state. A plant under the ordinary stresses of light and drought leans on its glutathione pool the way a body leans on rest. When we speak of growing food that holds the most nourishment it can, this molecule is a literal part of what we mean: a cell kept well-buffered is a cell that stays whole, and a plant that stays whole carries more of everything we want it to carry. So the first rule of the kokumi program is a rule of restraint — we never drain the pool to feed the fingerprint. We ask for surplus, not sacrifice.

The assembly is spare and elegant. Two bonds, two ATP, three amino acids:

Step one — the γ-bond is forged. The enzyme is **glutamate–cysteine ligase**, written GCL, known by its older name γ-glutamylcysteine synthetase (γ-GCS); in yeast and many plants the gene is *GSH1*. It takes L-glutamate and L-cysteine and, spending one ATP, joins them through glutamate's **γ-carboxyl** — not the usual α-carboxyl of a normal peptide bond, but the side-chain acid — forging the signature **γ-linkage** into **γ-glutamyl-cysteine**. This unusual bond is the whole reason the molecule survives the gut's ordinary peptidases long enough to reach a receptor: the standard machinery that chews α-peptides cannot read it. Step one is the committed step and the **rate-limiting** step of the pathway, held on a short leash for two reasons that become our two main design targets.

Step two — the tripeptide is closed. The enzyme is **glutathione synthetase** (GSS; *GSH2*), which takes the γ-glutamyl-cysteine, adds **glycine** at the far end, spends the second ATP, and closes the tripeptide into **glutathione** — γ-Glu-Cys-Gly. This step is fast and rarely limiting; it waits patiently on whatever step one hands it.

That makes the reference. But the kokumi *family* — γ-Glu-Val-Gly, γ-Glu-Val, γ-Glu-Cys, the whole louder library — needs the third enzyme, and the third enzyme is where flavor is actually distributed. We will take the three in order, because the mechanism of each is exactly where our edits bite.

Step one — forging the γ-bond

GCL is an **ATP-grasp** enzyme, and the choreography is worth drawing slowly. The enzyme clamps a molecule of ATP and the glutamate together in a closing cleft; it transfers the γ-phosphate of ATP onto glutamate's γ-carboxyl, producing a high-energy **γ-glutamyl phosphate** intermediate and leaving ADP behind. That phosphorylation is the trick: a carboxylic acid will not react with an amine on its own, so the cell spends ATP to turn the sluggish carboxyl into a reactive acyl-phosphate. Now cysteine's amino group attacks that activated carbon, the phosphate leaves, and the γ-bond snaps shut. The energy of one ATP has been spent buying a thermodynamically awkward linkage that ordinary peptide chemistry would never make.

Two brakes hold this step, and both are where kokumi is lost. The first is **product feedback inhibition**: glutathione, the pathway's own end product, binds back onto GCL at a regulatory site that **overlaps the glutamate pocket** — the product sitting in the doorway of its own synthesis. The cell wires this deliberately, a thermostat that keeps the glutathione pool from running away. The second brake is the **cysteine ceiling**: cysteine is the only sulfur-bearing amino acid in the recipe, sulfur is metabolically expensive, and the cell guards its cysteine jealously, so GCL often runs starved of one of its two substrates. Between the feedback thermostat and the cysteine ceiling, the committed step is kept modest on purpose. That modesty is exactly what a cook tastes as "thin."

We can put arithmetic on the first brake, because competitive inhibition at the glutamate pocket follows a clean law. Write the uninhibited step as Michaelis–Menten, with flux v through GCL at glutamate concentration S:

```
v = V_max · S ÷ (K_m + S)
```

Product inhibition that competes at the glutamate site does not lower V_max; it **raises the apparent K_m**, by a factor set by how much product is present relative to how tightly it binds:

```
K_m(apparent) = K_m · (1 + ρ), where ρ = [GSH] ÷ K_i
```

Here K_i is the product's inhibition constant and ρ is a dimensionless measure of brake strength — how many K_i's worth of glutathione the cell is holding. Work an example at the substrate level S = K_m, where an unbraked enzyme runs at exactly half its maximum:

- **No brake (ρ = 0):** v = V_max · K_m ÷ (K_m + K_m) = **V_max⁄2**.
- **Resting cell (ρ = 9):** K_m(app) = 10·K_m, so v = V_max · K_m ÷ (10 K_m + K_m) = **V_max⁄11** — about 18% of the unbraked rate.

A cell holding nine K_i's worth of its own product is running the committed step at under a fifth of its capacity. That is the thin broth in molecular form. The **H370L edit we write into GCL** reshapes the regulatory doorway so the product no longer lodges there — it *raises K_i*, which *lowers ρ* at the same glutathione concentration, without

touching the catalytic residues that do the actual chemistry. Suppose the edit raises K_i tenfold; then at the same resting glutathione the brake parameter falls from ρ = 9 to ρ = 0.9, and:

- **Edited (ρ = 0.9):** K_m(app) = 1.9·K_m, so v = V_max · K_m ÷ (1.9 K_m + K_m) = **V_max⁄2.9** — about 34% of V_max.

Flux through the committed step rises from ~18% to ~34% of maximum — roughly a **1.9-fold** increase — while V_max, the enzyme's intrinsic speed, is never altered. We have not made a faster enzyme; we have taken a working enzyme off a leash the cell put there for a stress it is not currently under. And because we leave the cysteine ceiling in place and never over-relieve the brake, the glutathione pool rises toward richness without ever falling below the floor the cell needs to protect itself.

Step two — closing the tripeptide

Step two is the easy one, and we rarely touch it. Glutathione synthetase accepts the γ-glutamyl-cysteine, positions a glycine, and — through its own ATP-grasp chemistry, phosphorylating the cysteine's α-carboxyl this time to activate it — closes the α-peptide bond to glycine and releases finished glutathione. It is fast; its K_m for γ-glutamyl-cysteine is low; it almost never backs up. In practice it means that whatever surplus γ-glutamyl-cysteine step one produces is promptly converted to glutathione, which is both the reference molecule and, now, the raw stock for the enzyme that makes the fingerprint. Step two is the courier between the rate-limiter and the redistributor, and a good courier is one you do not have to think about.

Step three — the transpeptidase sets the fingerprint

γ-Glutamyl transpeptidase (GGT) is the only enzyme that can cut the γ-bond, and it cuts it not to destroy the molecule but to **move** the γ-glutamyl group somewhere new. This is the enzyme that turns a redox buffer into a flavor catalog.

Its mechanism is a two-step **ping-pong** (double-displacement) run through a catalytic **threonine**. In the first half-reaction, the threonine's hydroxyl attacks the γ-glutamyl group of glutathione, forming a covalent **γ-glutamyl-enzyme intermediate** and releasing the rest of the donor — the cysteinyl-glycine tail — off the back. The enzyme now sits loaded, a γ-glutamyl group parked on its threonine, waiting. In the second half-reaction an **acceptor** arrives and attacks the intermediate, walking off with the γ-glutamyl group. If the acceptor is water, the product is free glutamate — plain **hydrolysis**, the γ-glutamyl group simply let go. If the acceptor is an amino acid — **valine**, say — the product is **γ-glutamyl-valine**, and with a glycine already positioned, the prized **γ-Glu-Val-Gly**, the loudest peptide in the kokumi catalog. GGT builds nothing from scratch; it **redistributes** the γ-glutamyl groups the first two enzymes made, spreading them across whatever acceptors it finds and favors. The particular mix of γ-glutamyl peptides in a food — its kokumi **fingerprint** — is set almost entirely by which acceptors win at that second step.

Which acceptor wins is, again, arithmetic. When two acceptors compete for the same loaded intermediate, the ratio of their products is the ratio of their **specificity constants** (k_cat⁄K_m) times their concentrations:

```
products to A ÷ products to B = (k_cat⁄K_m)_A · [A] ÷ (k_cat⁄K_m)_B · [B]
```

Frame the real competition as **transfer to valine** versus **hydrolysis to water**, and the fraction of flux that becomes useful kokumi peptide rather than wasted free glutamate is:

```
f(transfer) = (k_cat⁄K_m)_Val · [Val] ÷ [ (k_cat⁄K_m)_Val · [Val] + (k_cat⁄K_m)_water · [water] ]
```

The **T391S acceptor-pocket edit** reshapes where the acceptor binds so that valine is favored — it *raises (k_cat⁄K_m)_Val*. Suppose a food starts with transfer and hydrolysis evenly matched, f = 0.50, and the edit **doubles** valine's specificity constant. Then the ratio goes from 1 : 1 to 2 : 1, and:

```
f(transfer) = 2 ÷ (2 + 1) = 0.67
```

A third more of every γ-glutamyl group now lands on valine instead of leaking to glutamate — the fingerprint pushed, in one pass, toward γ-Glu-Val-Gly. Push the acceptor preference and you push the whole pool; that is the entire design handle of step three. The three enzymes, laid out with their intermediates named:

STEP	ENZYME	GENE	SUBSTRATES	PRODUCT	ATP	ROLE
1	GCL / γ-GCS	*GSH1*	L-Glu + L-Cys	γ-Glu-Cys	1	**rate-limiting**; feedback-inhibited by GSH
2	GSS	*GSH2*	γ-Glu-Cys + Gly	glutathione (GSH)	1	fast; closes the tripeptide
3	GGT	—	GSH + acceptor a.a. (e.g. Val)	γ-Glu-Val(-Gly)	0	**sets the fingerprint**; transpeptidation

Our highest-construct organism carries this whole machinery tuned at once. Phoenix (SS-0096) is the cultivar we designed to register on all eight axes, and our model predicts **IFI-K 1,900** for its kokumi axis — nearly twice the loudness of plain glutathione — alongside the umami, heat, and the rest:

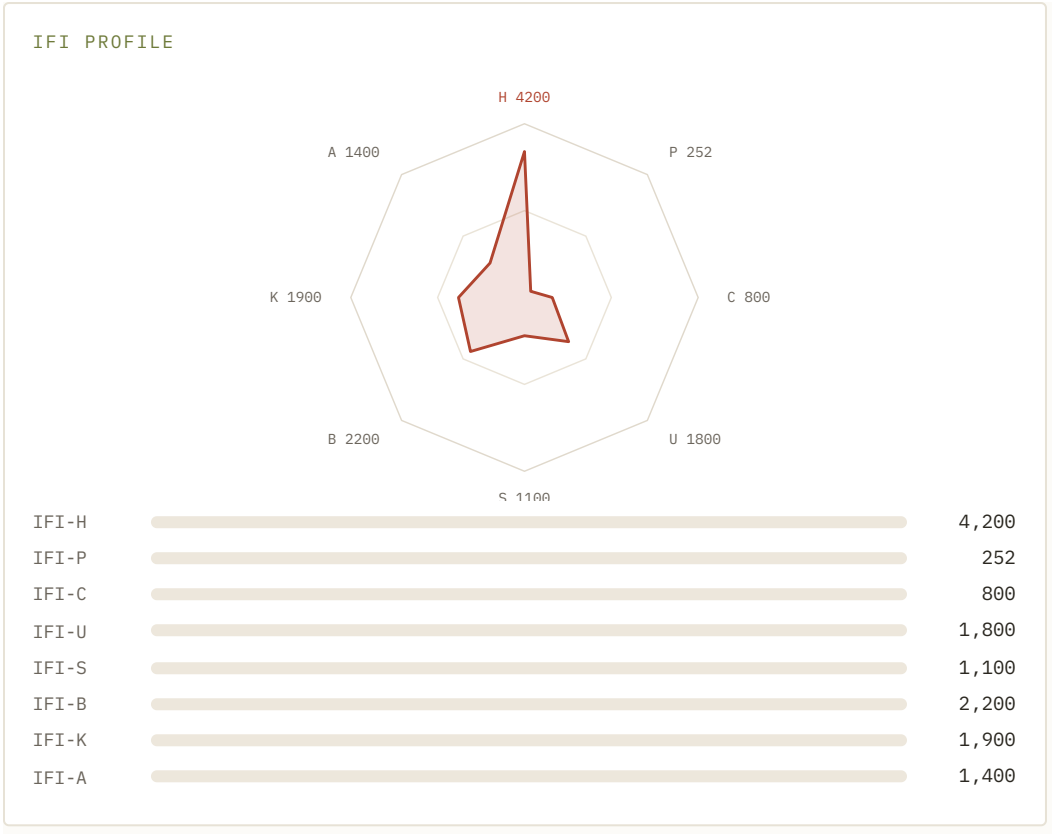

The cycle, not the dead end

The pathway is not a straight line to a finished peptide; it is a **cycle**, and the cycling is why aged foods accumulate specific γ-glutamyl peptides while fresh ones do not. In the full γ-glutamyl cycle — the Meister cycle, named for the biochemist who first drew it — GGT transfers the γ-glutamyl group to an acceptor; the resulting γ-glutamyl-amino-acid is closed by **γ-glutamyl cyclotransferase** into **5-oxoproline** (pyroglutamate), releasing the free acceptor amino acid; 5-oxoproline is then reopened to glutamate by **5-oxoprolinase**, spending one ATP; and that glutamate re-enters at GCL to begin again. The cycle's original purpose is to ferry amino acids and to recover glutamyl groups rather than waste them.

The ATP accounting is worth tallying, because it tells you what richness costs the cell. One full turn that moves an amino acid and regenerates glutathione spends **three ATP**: one at the oxoprolinase step to recover glutamate from 5-oxoproline, and two more to rebuild glutathione through GCL and GSS. The transfer and cyclization steps themselves are free — no ATP — which is why a slowly aging food can run the loop again and again on cheap thermodynamics, each pass nudging the peptide population toward the stable, accumulating members. In the lucky cases, those stable members are the loud ones, which is the whole reason aged things taste rich: they have had the months for the transpeptidase to do its patient redistribution. A long-fermented sauce, a cured cheese, a vanilla pod cured in the sun for weeks — all of them are kokumi that time built, one free transfer at a time.

Our engineering short-circuits the patience. Rather than let a year of partial cycling slowly enrich γ-Glu-Val-Gly, we bias GGT's acceptor preference so the enrichment happens in a single growing season, on purpose, inside the living plant. We do what aging does by accident, and we do it where the food grows, so that the richness arrives with the harvest rather than after it.

The allium's sulfur coat

In the allium this same pathway wears a sulfur coat, and the branch is the clearest molecular picture of why one plant can read loudly on two different axes at once. Garlic alkylates cysteine — it attaches an allyl or propenyl group to the sulfur — and feeds the resulting **S-alk(en)yl-cysteine** into the γ-glutamyl machinery. So the stored peptides are **γ-glutamyl-S-allyl-cysteines**: kokumi-active through their γ-glutamyl head, and at the very same time the plant's stable storage form of its pungency precursors. One molecule, two jobs.

Read the branch left to right and it is kokumi becoming pungency. GGT strips the γ-glutamyl group to free **S-allyl-cysteine**; that oxidizes to **alliin**; and when the clove is crushed, the enzyme **alliinase** — held apart from its substrate in the intact tissue — meets alliin and cleaves it in seconds to **allicin**, the sharp, volatile compound the tongue reads as garlic's bite. This is why garlic registers on both the richness axis and the sharp axis: the γ-glutamyl peptide is the quiet, stored, round form, and allicin is the loud, released, sharp form, and they are two readings of one pathway caught at two moments. The engineering choice in an allium is therefore a choice of *when* to stop the branch — hold the γ-glutamyl peptides for depth, or let them run to allicin for pungency — and we can tune the balance by where in that chain we place our hand.



Garlic in cross-section — the allium stores its flavor as γ-glutamyl-S-allyl-cysteine, a single molecule that reads as kokumi at the head and becomes pungency when the clove is cut.

The organisms where we do this work are the ones whose native chemistry already runs the γ-glutamyl cycle hard, and the construct counts track how much room each one gives us:

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

Garlic leads because its sulfur branch gives the most handles; soybean and yeast pair kokumi with umami, the two depth axes that so often arrive together in a broth or a ferment; onion and vanilla are the smaller, harder frontier, where the pool is subtle and the margins are thin. In yeast especially the cycle can be pushed until glutathione becomes a measurable fraction of the cell's dry mass — the organism is, in a sense, built to overproduce exactly the molecule we want.

The math of loudness

With the enzymes understood, the IFI math becomes concrete, and it is worth working in full because the number is what makes kokumi designable rather than merely describable. Take the kokumi axis of Phoenix, our model's predicted **IFI-K 1,900**. Invert the index formula to recover what the receptor actually saw. Since

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

solving for the ratio of concentrations gives

> $\mathrm{EC}_{50}(\mathrm{candidate}) \div \mathrm{EC}_{50}(\mathrm{ref}) = 1{,}000 \div \mathrm{IFI} = 1{,}000 \div 1{,}900 =$ **0.526**.

The designed peptide pool is predicted to reach half-maximal CaSR potentiation at **0.526 times** the concentration the reference needs — which is to say it is **1 / 0.526 = 1.9 times more potent** than plain glutathione in our model. The score and the potency are the same fact read two ways. Lay the axis out as a short ladder from the reference to the scorecard ceiling:

IFI-K	EC_{50}(CANDIDATE) ÷ EC_{50}(REF)	POTENCY VS. GLUTATHIONE
1,000 — glutathione (reference)	1.00	1.0×
1,900 — Phoenix · SS-0096	0.526	1.9×
5,000 — scorecard normalization ceiling	0.200	5.0×

The per-axis scorecard saturates at 5,000, which the table shows is a candidate five times louder than the reference — a deliberate ceiling, because beyond roughly five-fold potentiation the cooperative receptor is already near the top of its curve and extra loudness buys little mouth-perceptible richness.

That cooperativity is the second piece of math, and it governs how forgiving the axis is to formulate against. Model the receptor's fractional response with a Hill function, writing concentration in units of its own EC_{50} so no absolute anchor is needed — let $x = C / \mathrm{EC}_{50}$:

> $\theta = x^n \div (1 + x^n)$

At $x = 1$ the response is exactly half, $\theta = 0.5$, for **any** Hill coefficient n — that is what makes EC_{50} a clean anchor. But n sets the **steepness**, and CaSR, as a cooperative dimer, reads its ligands with n commonly

near 2–3 rather than the n = 1 of a simple one-site receptor. The span of concentration needed to climb from 10% to 90% of full response is:

- **n = 1:** x must run from ⅑ to 9 — an **81-fold** concentration window.
- **n = 2:** x runs only from ⅓ to 3 — a **9-fold** window.

Cooperativity collapses the window nine-fold. That is a gift and a hazard at once: a gift, because a food tipped into the band turns round quickly and convincingly; a hazard, because the margin between "flat" and "full" is narrow, so the concentration of the loud peptides has to be held tightly from batch to batch. Which is precisely why the whole program lives or dies on consistency.



Phoenix on the microscope stage — the kokumi axis is read as half-maximal CaSR potentiation against a fixed background, and held to a batch-to-batch CV at or below 7%.

What it costs to be sure — error, CV, and honest limits

A number is only as good as its reproducibility, and here the arithmetic of error is as important as the arithmetic of potency. The IFI-K score is a **ratio** of two predicted EC_{50} values, so its relative uncertainty is the two underlying uncertainties added in quadrature:

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

Our engineered cultivars are designed to hold a batch-to-batch coefficient of variation at or below **7%**. If both predicted EC_{50} values carried that 7%, the score would inherit

```
CV(IFI) = √(0.07² + 0.07²) = 0.07 × √2 ≈ 9.9%,
```

and when the reference is treated as a fixed, repeatedly calibrated standard — its variation driven toward zero — the score inherits essentially the candidate's **7%** alone. On an IFI-K of 1,900 that is a band of about ±133 to ±188. Now run the same arithmetic on an unengineered plant, whose capsaicinoid and peptide output we have measured at a wild-type CV near **42%**:

```
CV(IFI, wild-type) = √(0.42² + 0.42²) = 0.42 × √2 ≈ 59%,
```

which on 1,900 is a band of roughly ±1,120 — an uncertainty so wide the number stops meaning anything. This is the whole case for engineering stated in one line of algebra: a 7% process turns a flavor into a specification; a 42% process leaves it a rumor. Against a receptor whose 10-to-90 window is only nine-fold wide, a 59%

swing in loudness routinely throws a food out of the band entirely, which is why wild material tastes round one season and thin the next.

We are careful to name what the number does not capture, because a measurement you oversell is a measurement you will be caught by. Four honest limits bound the kokumi axis:

- **Kokumi is a modulation, not a flavor, so it is only ever measured in context.** IFI-K is defined against a fixed background of calcium and base taste; change the background and the same peptide reads differently. The score is a property of the peptide-against-a-standard, not a property the peptide carries alone into every dish. We publish the background with the number, always.
- **The fingerprint is a distribution, not a single peptide.** GGT's acceptor pocket is reshaped toward valine, but it is never made perfectly selective; a real food carries a spread of γ-glutamyl peptides, and the IFI-K we report is the integrated loudness of the mix. Push the pocket too hard toward one acceptor and the others thin out in ways the single number hides.
- **Transpeptidation is an equilibrium, and it can run backward.** GGT does not only build γ-glutamyl peptides; under the wrong balance of substrate and acceptor it will hydrolyze them back toward free glutamate. A fingerprint engineered in the plant has to be stable to the conditions of cooking and storage, or the richness drains back out of the pool it was concentrated into.
- **The committed step must never be over-relieved.** Lift GCL's product brake too far and the enzyme runs down the cell's cysteine and skews its redox balance — the plant pays in the exact currency glutathione exists to protect. So the H370L edit is tuned to raise K_i modestly, enough to lift flux toward richness and no further, and we watch the glutathione floor the way we watch anything that keeps the plant well.

That last limit is where the chemistry and the ethos meet. Three enzymes, two ATP, three amino acids and a single sulfur atom — that is the entire molecular cost of the richest sensation in food, and all of it is already running, quietly, in every organism on the bench. In the allium it does double duty, the γ-glutamyl peptides serving as both the kokumi pool and the stored precursors of pungency, which is why garlic reads loudly on two axes from one pathway. In yeast it can be pushed until glutathione is a weighable fraction of the cell. In fermenting molds and the long age of cheese it works slowly over months, redistributing a little at a time, which is why aged things are rich: they have had the time to let the transpeptidase do its patient work. We do in a growing season, and on purpose, what aging does by accident over a year — and we do it without ever asking the cell to spend below the floor that keeps it, and the food it becomes, whole. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 6 · Designing the pathway — *flux, feedback, and the AI that reads proteins.*

Knowing a pathway is not the same as commanding it. For almost the whole history of rich food the γ-glutamyl pathway ran where it wished — a little more glutathione in this clove, a different spread of γ-glutamyl peptides in that one, set by the plant's own quiet priorities and discoverable only by tasting the result and keeping the seed of whatever moved us. It made richness the way weather makes rain: reliably on the whole, never twice the same. Our method is to make that pathway answer to a number. We decide the IFI-K we want — the loudness of kokumi, read off the calcium-sensing receptor — and the fingerprint of peptides we want underneath it, and then we design the organism to deliver both with a batch-to-batch steadiness no field of garlic ever had.

The method does not change when the pathway does. The same two models that taught a chili to burn on command teach a clove to deepen on command, carried across without a line rewritten, and they run in two layers. One model decides *where* to intervene — which single reaction, out of thousands, actually holds the richness back. The other decides *how* — which one amino acid, out of a protein's hundreds, to change, and to what, so the enzyme runs harder without losing its fold. One is arithmetic over an entire metabolism; the other is a kind of reading, a machine that has read the protein universe and can tell a plausible enzyme from a broken one. Between them they turn a design problem that used to take a career into a ranked shortlist a few dozen edits long — which we then specify down to the last residue. These are computational predictions; the model narrows the search, ranks the candidates, and delivers the design.

Before any of that, though, it is worth being exact about the thing we are writing toward, because the receptor sets every number that follows.

The receptor we are writing toward

Kokumi is the axis you have tasted ten thousand times and never once named. It is not a basic taste the way sweet and salt are; on its own, a kokumi compound is nearly flavorless. What it does is *modulate* — it reaches under a broth or a long-cooked sauce and makes everything already there feel fuller, rounder, more continuous, arriving a beat sooner and lingering a beat longer, with the body and the mouthfulness a Japanese cook would call *koku* centuries before anyone found the receptor. Richness, continuity, mouthfeel: three words for one sensation, the sensation of a food being more wholly *present* on the tongue. Because kokumi works by lifting what is already there rather than adding a taste of its own, it reads most vividly alongside the savory and the salty — it makes an umami

broth taste brothier and a salted dish taste rounder at less salt — which is exactly why our loudest kokumi work lives in the organisms that also carry umami, the soybean and the yeast, where the two axes reinforce each other in the same mouthful.

The receptor is CaSR, the calcium-sensing receptor — a class C G-protein-coupled receptor the body first evolved to watch the calcium level in the blood, and that the tongue quietly repurposed to read richness. Its architecture is the key to everything we do to it. A large, bilobed extracellular domain hangs off the cell like a Venus flytrap and closes around its ligand the way a flytrap closes on a fly; a cysteine-rich collar carries that motion down through the membrane to a bundle of seven helices, which passes the signal inward. CaSR works as a homodimer — two flytraps hinged together — and that pairing is why its response is *cooperative*: one lobe closing primes the other, so activation climbs more steeply than a single binding site ever could. In the language of dose–response, the Hill coefficient sits above one, and the curve is a staircase rather than a ramp. We will use that steepness; it is what lets a modest shift in concentration produce a large shift in richness.

The agonists that matter for food are the **γ-glutamyl peptides** — short peptides built on a glutamate joined through an unusual bond. An ordinary peptide links the α-carboxyl of one residue to the amine of the next; a γ-glutamyl bond links glutamate's *side-chain* carboxyl instead, the γ-carboxyl, a linkage most of the mouth's proteases cannot cut. That single structural quirk is why these peptides survive long enough to be tasted, and why only two specialized enzymes make and move them — one that writes the γ-bond and one that transfers it. The quietest of the family is **glutathione**, γ-Glu-Cys-Gly, the cell's own antioxidant tripeptide, present in every leaf and every ferment; we were glad to find the whole axis already vegetarian, already the plant's own chemistry. Louder, far louder gram for gram, is **γ-Glu-Val-Gly**, among the most potent kokumi peptides known. Louder still in an allium sense is **γ-Glu-S-allyl-Cys**, the sulfur-scented γ-glutamyl peptide that gives a clove of garlic its depth. The loudness varies enormously from peptide to peptide, which is the whole reason a kokumi design is two numbers and not one: not only *how much* richness, but *which peptide* carries it. We tune the total with one hand and the fingerprint with the other.

What a number buys us

The Index turns that loudness into a score. Every axis is anchored to the half-maximal activation of its receptor — not how much of a compound a food holds, but how *loudly* that compound speaks to the nerve — and kokumi is no exception:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response; a smaller EC_{50} means a louder compound, because less of it is needed to reach the same half-way point. The reference is fixed so that its IFI-K reads **1,000 by construction** — set a candidate equal to the reference and the ratio is one, times a thousand. Everything else is read against that anchor. Because the

score only ever contains EC_{50} as a *ratio* to the reference, it is portable: the absolute concentration can be called c_ref and left as a symbol, and no number below depends on its value.

Work an example in full. Wild garlic registers a natural IFI-K of **520**; our model predicts an IFI-K of **3,200** for the garlic we design. Invert the formula to recover each EC_{50} as a multiple of the reference:

- natural garlic: EC_{50} = 1,000 × c_ref ÷ 520 = **1.923 · c_ref** — quieter than the reference, needing nearly twice its concentration to reach the half-point;
- engineered garlic: EC_{50} = 1,000 × c_ref ÷ 3,200 = **0.313 · c_ref** — louder, reaching the half-point at under a third of it.

The loudness gain is the ratio of the scores, 3,200 ÷ 520 = **6.15×**, and the EC_{50} falls by exactly the same factor, 1.923 ÷ 0.313 = 6.15. To make the units concrete, carry an *illustrative* reference concentration of 1.0 mM — a round anchor, not a measured constant, chosen only so the arithmetic lands somewhere. Engineered garlic's lead peptide, γ-Glu-S-allyl-Cys, has a molecular weight of 290.3 g/mol (glutamate 147.13 + S-allyl-cysteine 161.22 − one water), so its EC_{50} of 0.313 mM converts to 0.313 mmol/L × 290.3 mg/mmol ≈ **91 mg/L**, against the wild clove's 1.923 mM ≈ 558 mg/L. Six-fold louder is six-fold less peptide to say the same thing.

The steepness of CaSR turns that EC_{50} shift into a perceptual one. Model the receptor's fractional activation θ with a Hill curve — $\theta = [c]^n \div (EC_{50}^n + [c]^n)$ — and read both cloves at the same serving concentration, a dose matched to the reference, [c] = c_ref. With the conservative, non-cooperative case n = 1 the expression collapses to a small, memorable identity:

> `θ = IFI-K ÷ (IFI-K + 1,000)` *(fractional CaSR activation at a reference-matched dose, simple Hill)*

At a reference-matched dose, wild garlic drives the receptor to 520 ÷ 1,520 = **0.342**, barely a third; the engineered clove drives it to 3,200 ÷ 4,200 = **0.762**, better than three-quarters. Let the dimer show its true cooperativity at n = 2 and the gap widens to a chasm — 0.213 against 0.911, a fifth of the receptor against nearly all of it — because a steep curve rewards loudness twice over. The same bite of food, the same amount in the mouth, and the difference between a flicker of richness and a flood of it.



Garlic in cross-section — 560 constructs ride on this one species, more than any other organism on the kokumi axis. The γ-glutamyl peptides that make a clove taste like depth itself are the cell's own glutathione chemistry, read and turned up, never invented.

Where to cut — flux-balance analysis

Before we change a single base, we build a genome-scale metabolic model of the organism: a complete accounting of its reactions, each written as a column of what goes in and what comes out, stacked into a stoichiometric matrix **S** with one row per metabolite and one column per reaction. At steady state nothing accumulates — every metabolite is produced exactly as fast as it is consumed — which is the single linear constraint $\mathbf{S} \cdot \mathbf{v} = \mathbf{0}$, where **v** is the vector of fluxes we are solving for. Bound each flux by its thermodynamics and its enzyme's capacity, then, within the whole space of flux patterns that satisfy the balance, maximize one objective by linear programming: here, the flux reaching the γ-glutamyl peptide pool. The model does not simulate the cell through time or track a single concentration; it finds the flux distribution most consistent with the stoichiometry and reports which reactions, if their bounds were relaxed, would most raise the objective. Those shadow prices are the bottlenecks, and they are rarely where intuition points.

For kokumi the answer is robust across organism after organism, and it comes in two parts. The **first bottleneck is the feedback brake on glutamate-cysteine ligase (GCL)** — the enzyme that writes the γ-glutamyl bond, product-inhibited by the very glutathione it makes, so that the pathway throttles itself long before the enzyme runs out of capacity. It is a thermostat the cell installed to keep its antioxidant pool from overflowing, and it holds flux far below the ceiling. The **second is cysteine supply** — the sulfur ceiling — because cysteine is the pathway's one sulfur-bearing input, and even with the brake off the enzyme starves for it unless the upstream sulfate-assimilation route is fed in parallel: sulfate reduced to sulfide, serine acetylated to O-acetylserine, the two condensed to cysteine. The two are bound by a hard stoichiometry the model makes explicit: every γ-glutamyl peptide demands one cysteine, each cysteine one atom of reduced sulfur, so the sulfate-assimilation flux must climb step for step with the peptide flux — a one-to-one coupling, not a loose dependency. Relieve only the feedback and the organism pushes for a moment and then stalls, cysteine-limited; relieve only the cysteine ceiling and the feedback brake simply tightens its grip.

The flux model's gift is to tell us, before any bench work, that **kokumi is a two-handed problem**, and to spend our edits accordingly.

It also keeps us honest about cost. Each molecule of glutathione is paid for with two ATP — one at the ligase, one at the synthetase that adds the glycine — and an organism asked to flood its tissue with γ-glutamyl peptides is an organism spending real energy it would otherwise put into growth. We spend the plant's energy the way we'd spend our own: carefully, to a limit. The model sets the ceiling past which vigor pays for richness, and we design to the edge of it, never past. That is the honest boundary of the method, too — flux balance sees stoichiometry and thermodynamics, not kinetics, not the clock. It tells us where to cut and how much we can afford; it cannot tell us how the edited cell behaves minute to minute. For that, we build kinetic models on top of the stoichiometric foundation — another computational layer before a prediction becomes a design.

How to edit — the model that reads proteins

Once the flux model names GCL and its partner transpeptidase as the targets, a harder question opens: which exact substitutions relieve the feedback, sharpen turnover, and reshape the acceptor pocket toward valine, all without collapsing the fold. The search space is merciless — a few hundred residues, nineteen alternatives each — so we read it with a 650-million-parameter protein language model, the same ESM2 that scored our capsaicin-synthase edits. Trained on hundreds of millions of natural protein sequences, it has learned the deep grammar of which residues belong where, and it scores a proposed substitution the way a fluent reader scores a word swapped into a sentence: by how much the change disturbs what the protein universe says is plausible. Concretely, it is the difference of log-likelihoods — how much more or less probable the mutant residue is than the wild-type one, given every other residue as context — and that number tracks both the change in catalytic efficiency and the change in fold stability. We generate every single-residue variant in the regulatory and active-site regions in silico, score them all, rank by predicted effect × predicted stability, and keep only the handful that promise more function while holding or improving the fold.

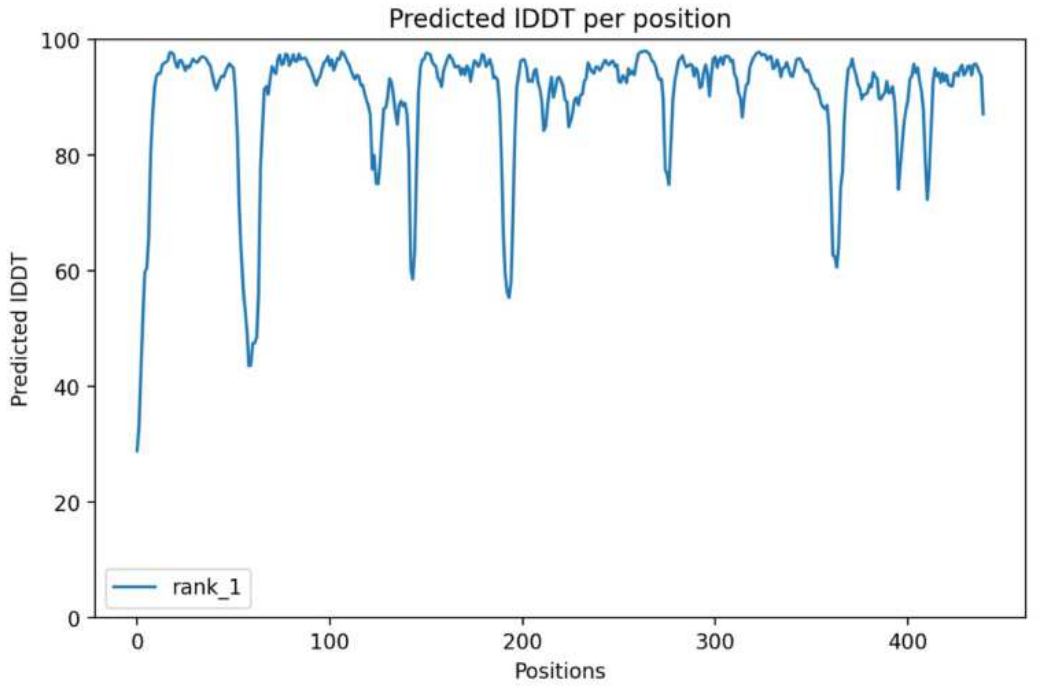

FIG. 11 — a pathway enzyme coloured by model confidence (pLDDT): the active-site cleft and the regulatory pocket resolve sharply, the surface loops less so. The same per-residue confidence scoring guides every edit we make to the γ-glutamyl enzymes; we write where the colour runs hot and leave the cool, uncertain loops alone.

A plausibility score alone is not enough, because the model is only as trustworthy as our picture of the protein. So we cross-check every survivor against a structural model whose per-residue confidence — pLDDT, a 0-to-100 scale — tells us which parts of the fold we understand well enough to touch. Where pLDDT runs hot, in the buried active-site cleft and the regulatory pocket, the structure is sharp and an edit's consequences are legible; where it runs cold, in the flexible surface loops, we leave the protein alone no matter how tempting the score. A second map carries the same caution one level up. The predicted aligned-error plot asks a different question — not how sure we are of each residue, but how sure we are of the *relative position* of any two — and its answer is what licenses an edit in one pocket to be judged safe for another.

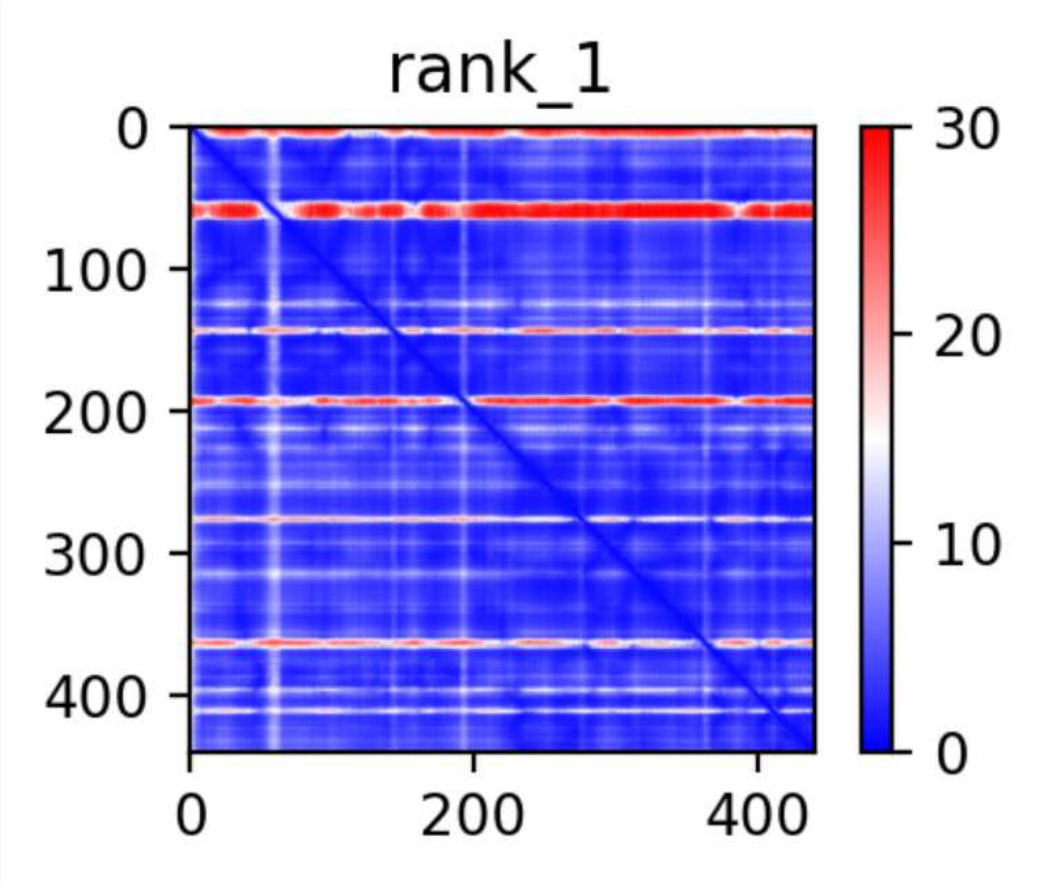

FIG. 12 — the predicted aligned-error map for the same model: dark blocks mark residue pairs whose relative position the model is confident of. Low error between the regulatory pocket and the active site is our permission to loosen the feedback residue without fear of swinging the catalytic one out of place.

Every surviving edit carries a predicted **ΔΔG**, the change in folding free energy in kilocalories per mole, and the sign convention is the same throughout this Library: a *more negative* ΔΔG is a *more stabilizing* edit, a protein that holds its fold more firmly under the heat and stress of a real growing season. That number is not a mood; it is a Boltzmann factor. The folded and unfolded states sit in equilibrium, and a mutation shifts the ratio between them by exp(−ΔΔG ÷ RT), where RT at 25 °C is 0.593 kcal/mol. A −1.6 kcal/mol edit therefore tilts the folded-to-unfolded balance by exp(1.6 ÷ 0.593) ≈ **15-fold** toward folded — a protein fifteen times more likely to be found doing its job than coming apart. We require our edits to be net-stabilizing *and* functional; we will not trade richness for a fragile enzyme that frays in the field.

The four edits

The flux model's two hands and the language model's shortlist resolve, across the two target enzymes, into four substitutions no breeding program could ever have found — because no breeding program selects on a regulatory pocket it cannot see. We compute the stabilization each one buys directly from its ΔΔG:

EDIT	ENZYME	EFFECT	ΔΔG (KCAL/MOL)	FOLDED : UNFOLDED SHIFT (25 °C)
H370L	GCL (*GSH1*)	glutathione feedback-pocket relief; brake released	−1.6	14.9×
S495A	GCL (*GSH1*)	active-site turnover ↑; k_cat improved	−1.2	7.6×
T391S	GGT	acceptor pocket reshaped toward L-Val	−0.8	3.9×
R107K	GGT	γ-glutamyl donor affinity ↑; transfer efficiency ↑	−1.1	6.4×

Read them as the mechanism made real. The two GCL edits are the two hands of the flux problem: **H370L** reaches into the regulatory pocket and loosens the grip of glutathione on its own maker — the wild-type histidine offers a polar contact the inhibitor leans on to dock, and swapping it for a nonpolar leucine withdraws that handhold, so glutathione can no longer hold the brake while the distant catalytic site is left entirely untouched; **S495A** sharpens the active site so that, once freed, it runs faster. The two edits on the γ-glutamyltransferase redirect where the enlarged pool goes — **T391S** reshapes the acceptor pocket so the transferred γ-glutamyl group prefers valine, and **R107K** tightens the enzyme's hold on its glutathione donor so the transfer happens cleanly — steering the chemistry toward γ-Glu-Val-Gly, the loud peptide, rather than scattering it across quieter acceptors. Feedback released, turnover raised, pool redirected: four edits, two enzymes, and an organism designed to make the kokumi we specified. Taken together — and treating the four as first-order additive, which they roughly are

across two independent folds — the combined −4.7 kcal/mol tilts the enzymes' folding equilibria by about **2,800-fold** toward folded, which is why our designs are computed to maintain their richness through a real growing environment rather than only under idealized conditions.

We write the four into the design for CRISPR delivery, and the predicted payoff is not only loudness but *constancy*. A wild crop's richness drifts with weather, soil, and the plant's private moods; our designed lines are specified to hold a batch-to-batch coefficient of variation at or below **7%**. That number propagates cleanly through the Index. Because IFI-K is the reference concentration — an exact anchor — divided by a predicted EC_{50}, its relative error is just the relative error of that prediction, so a 7% CV design target translates directly to a 7% CV on the score. At IFI-K 3,200 that is a standard deviation of ±224, a one-sigma band of 2,976–3,424 and a two-sigma band of roughly 2,752–3,648 — a richness you can formulate against and trust to arrive, batch after batch, against the open-ended scatter of a crop that was never designed.

One pathway, every organism

The reach of all this is the quiet thesis of the whole kokumi program. The γ-glutamyl pathway is not a specialty of garlic; glutathione is universal, running in every vascular plant and every yeast, which means the four-edit logic — relieve the feedback, raise the turnover, redirect the transpeptidase, feed the sulfur — **transfers unchanged to any organism with a sequenced genome**. The work distributes across the kokumi organisms roughly in proportion to how much of the axis each is asked to carry, and the loudness gain is simply the ratio of the two scores:

ORGANISM	CONSTRUCTS	PRIMARY AXES	NATURAL IFI-K	DESIGNED IFI-K	GAIN	LEAD PEPTIDE
Garlic	560	K / P	520	3,200	6.15×	γ-Glu-S-allyl-Cys / GSH
Soybean	360	U / K	640	2,400	3.75×	γ-Glu-Val-Gly / GSH
Yeast	240	U / K	700	2,800	4.00×	glutathione
Koji mold	80	U	400	2,100	5.25×	γ-Glu-Val-Gly (built)
Onion	40	K	180	1,900	10.6×	γ-Glu peptides
Vanilla	40	K	150	1,600	10.7×	GSH / γ-Glu (mouthfeel)
Total	**1,320**	—	—	—	—	–

The same arithmetic that read engineered garlic reads every row. Convert each designed score to an EC_{50} as a multiple of the reference (1,000 ÷ IFI-K) and to a fractional CaSR activation at a reference-matched dose, under both the flat and the cooperative Hill cases:

ORGANISM	DESIGNED IFI-K	EC_{50} (× REFERENCE)	θ AT N = 1	θ AT N = 2
Garlic	3,200	0.313	0.762	0.911
Yeast	2,800	0.357	0.737	0.887
Soybean	2,400	0.417	0.706	0.852
Koji mold	2,100	0.476	0.677	0.815
Onion	1,900	0.526	0.655	0.783
Vanilla	1,600	0.625	0.615	0.719

The ladder is not confined to the traditional carriers of richness, either. We have written kokumi into organisms no cook would think of as savory — and, going the other way, into our full-spectrum heat flagship, which carries a designed kokumi signal of K 1,900 (θ = 0.655 at n = 1) folded in among its other seven axes, a chili that is also, on purpose, round:

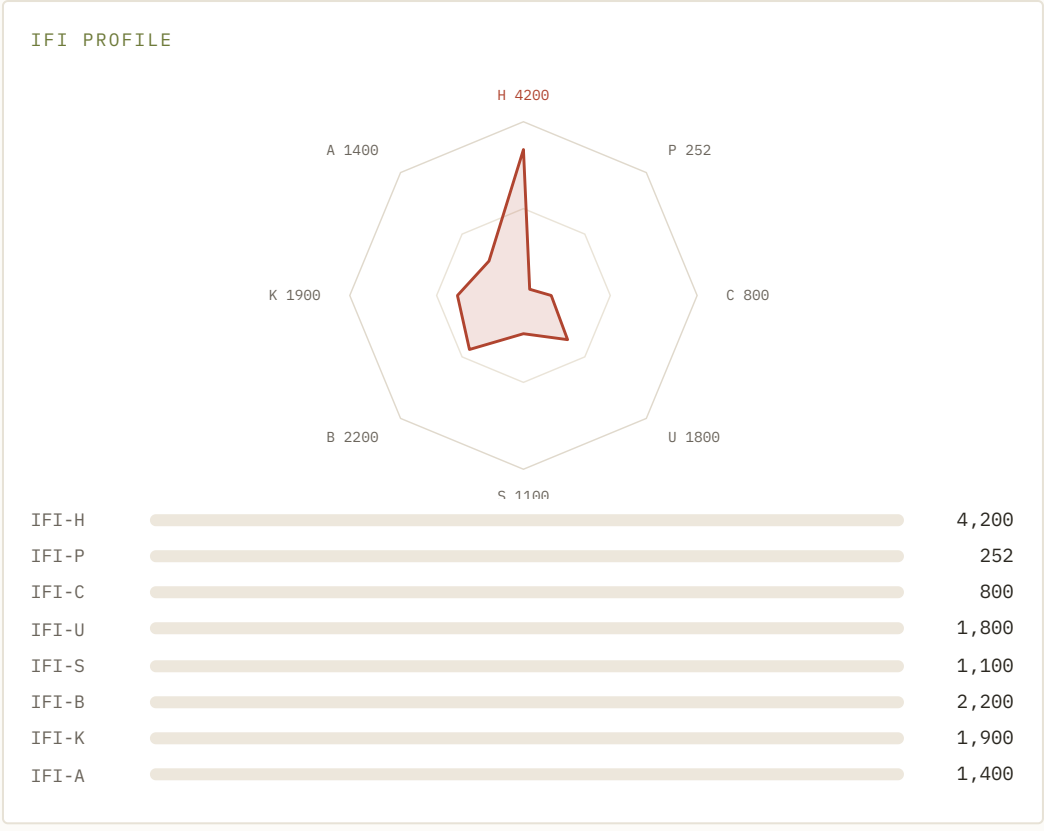

On a scorecard, where each axis is normalized to a per-axis maximum of 5,000, engineered garlic's 3,200 reads as 64% of full scale and onion's 1,900 as 38% — room still above every one of them, by design, because we build to the edge of the vigor ceiling and not past it.

The honest limits sit in plain view on these tables. Koji mold's lead peptide is marked *built*, not merely amplified: it lacks the native route to γ-Glu-Val-Gly, so there the work is to install a pathway rather than tune one, and the edit count is small because the hard part is

upstream. Every row depends on feeding the sulfur in parallel — relieve the feedback alone and the brightest of these organisms stalls, cysteine-starved, within a generation. Every engineered score carries the vigor cost of two ATP per peptide, which is why none of these numbers runs to the 5,000 ceiling. And the Hill columns are a model: θ is the receptor's fractional activation under a simple curve, a computed estimate from the first principles of cooperative binding. That is the discipline the whole program runs on — the model narrows the search, ranks the candidates, and delivers the design. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The twenty-one papers and three patents in this library document the computational method in full.

What emerges, across more than thirteen hundred constructs and six organisms, is a kind of reverence turned into arithmetic. Not one of these peptides was invented. The γ-glutamyl chemistry was already there, in every clove and every ferment, written by the cell for its own quiet reasons — the richness it makes a byproduct of keeping itself whole. All we did was learn to read the pathway carefully enough to ask it, politely and exactly, for more.



PART III OF V

The Kokumi Organisms

garlic · the alliums · yeast & koji · soybean · vanilla · the aged

Chapter 7 · Garlic — *the savory foundation.*

Begin where kokumi was discovered, which is also where it is densest in the living world: the garlic clove. Cut one and the smell comes up at once, sharp and total, and that smell is a decoy — it is not kokumi at all, it is pungency, the broken-cell alarm of allicin firing on the pungency axis (IFI-P). The kokumi is the other thing, the quiet thing underneath, the reason a head of garlic roasted slow and whole turns sweet and deep and *round* in a way that has nothing to do with the raw clove's bite. Garlic is the organism that taught the world the word, and in our program it is the flagship of the axis: **560 constructs**, more than any other kokumi organism we grow, and the one where the biology of richness and the biology of pungency are the same biology, read in two directions.

We came to garlic the way our family comes to everything — at the table first, the receptor second. For years we cooked with it by instinct, the single clove that changes a whole pot, the roasted head that turns to butter; only later did we learn to read what our hands already knew. What follows is that reading: the molecule, the receptor it speaks to, the exact arithmetic of how loudly it speaks, and the clove we built to speak louder without ever raising its voice into a shout.



The whole head, backlit — a warehouse of γ-glutamyl peptides, kokumi held stable until the knife turns it toward pungency.

The duality, exactly — one molecule, both the gift and the weapon

Here is the duality, because it is the key to garlic and found nowhere else so cleanly. The allium stores its flavor as **γ-glutamyl peptides** — chief among them γ-glutamyl-*S*-allyl-cysteine and its near relatives, a cysteine carrying a sulfur side-group, capped by a glutamate joined through an unusual bond. That bond is the whole story. Glutamate has three handles: an α-amino group, an α-carboxyl, and a second carboxyl out at the end of its side chain, the γ-carboxyl. Ordinary peptides are built off the α-carboxyl; these storage

peptides are built off the **γ-carboxyl** instead — a γ-glutamyl linkage. These peptides are stable, odorless, and — this is the point — **kokumi-active**: the γ-glutamyl group is precisely the shape the calcium-sensing receptor reads. So intact, uncut, a clove is a warehouse of kokumi held in a form that will keep for months.

Then the knife comes. Tissue damage brings two enzymes to bear in sequence. First, **γ-glutamyl transpeptidase** (GGT) strips the γ-glutamyl cap from the storage peptide, releasing the free cysteine sulfoxide — **alliin**, *S*-allyl-L-cysteine sulfoxide, chief among them. Second, **alliinase**, a pyridoxal-5′-phosphate–dependent carbon–sulfur lyase held in separate cellular compartments until the cell wall breaks, cleaves alliin in milliseconds into a reactive allyl sulfenic acid, pyruvate, and ammonia; two sulfenic acids condense into the thiosulfinate **allicin** — the pungency, the sharpness, the alarm on TRPA1.

Read the pathway in the direction the plant runs it and the irony is complete: garlic's own chemistry spends the kokumi store *down* to build the pungency *up*. Richness is the precursor of sharpness. The same γ-glutamyl molecules that speak softly to the calcium receptor are the raw material for the ones that will, moments later, shout at TRPA1. GGT is the hinge between the two — the branch-point enzyme that decides whether a given molecule stays in the stored, rich, kokumi-active pool or is handed down the line toward the weapon.

It is worth pausing on *why* the allium, of all plants, became the kokumi champion, because the answer is sulfur and the answer is defensive. Garlic and its kin invest heavily in sulfur chemistry as a defense against being eaten — the pungency is a weapon — and to build that weapon safely they must hold its precursors in a form that will not poison their own living tissue, which is the γ-glutamyl peptide: stable, inert, locked. The allium accumulates kokumi peptides not to be delicious but to hold a loaded gun safely, and the richness we prize is the safety catch on the plant's chemical defense. This is a recurring grace of the living world, that so much of what we find delicious is a plant's attempt to be *un*-eaten — and nowhere is it cleaner than here, where the same molecule is both the stored threat and the gift of depth. Our engineering, in deepening the γ-glutamyl store and slowing its conversion, is in a real sense asking the plant to keep more of its defense in the safe, rich, stored form and convert less of it to the weapon. We are breeding, you might say, for a garlic that holds its fire as flavor.

The receptor it speaks to — CaSR, step by step

To know what kokumi *is*, you have to meet the receptor, because kokumi is defined by it the way heat is defined by TRPV1. The calcium-sensing receptor, **CaSR**, is a class C G-protein-coupled receptor — the same structural family as the metabotropic glutamate and sweet-and-umami taste receptors, and nothing like the single-pass ion channels that carry heat and pungency. A class C receptor

is a large, two-lobed thing shaped like a Venus flytrap, working as a homodimer: two identical halves clasped together, each with a cleft that opens and closes.

Its primary agonist is the calcium ion itself, Ca^{2+}, which settles into orthosteric sites in the cleft and at the seam between the two halves. The kokumi peptides are not agonists in that primary sense — they are **positive allosteric modulators**. A γ-glutamyl peptide binds a site in the extracellular flytrap and, by its presence, stabilizes the closed, active clasp, so that the receptor fires more readily at any given calcium background. This is the molecular root of everything a cook knows about kokumi: it does not taste of anything itself, it *amplifies and prolongs* whatever else is on the tongue — the thickness, the continuity, the sense of a flavor that arrives, stays, and finishes rather than flashing and vanishing. Kokumi is a modulator, and this is not a limitation to hide but the fact to design around.

The activation runs in four steps, and they matter for the math that follows:

1. **Occupancy.** Ca^{2+} fills the orthosteric sites in the flytrap cleft and at the dimer interface, setting a baseline tension in the clasp.
2. **Modulation.** A γ-glutamyl peptide docks at its allosteric site and stabilizes the closed conformation — it lowers the receptor's effective threshold, left-shifting the whole response.
3. **Propagation.** The flytrap closure travels down through the cysteine-rich linker domain into the seven-pass transmembrane bundle, twisting it into its active geometry.
4. **Transduction.** The active bundle engages Gq/11, which fires phospholipase C-β; PLCβ splits membrane PIP_2 into IP_3 and diacylglycerol; IP_3 releases calcium from internal stores; the taste cell depolarizes, and the signal climbs toward a brain that will, a quarter-second later, call it *savory depth*.

Two molecular facts from this cascade become the whole engineering problem. The first is the **γ-glutamyl recognition motif**: because the receptor reads the γ-glutamyl group specifically, the very enzyme that builds and dismantles that group — GGT — is also the lever on the kokumi signal, which is why deepening the store and slowing GGT are the two hands of one design. The second is **cooperativity**: because CaSR is a dimer with several occupancy sites acting in concert, its dose-response is steep — not the gentle slope of a single lock and key, but something closer to a switch. That steepness is a number, and the number is where we go next.

Why kokumi needs a number — anchoring IFI-K to glutathione

For ten thousand years kokumi had no number because it had no name you could measure against. You cannot design what you cannot measure, so the first thing we did with the calcium receptor was anchor it. The Index scores every axis the same way — as loudness, not quantity — with one formula:

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

EC_{50} is the concentration that carries a receptor to half of its full response; invert it against a fixed reference and scale by a thousand, and a soft molecule and a loud one can be read on the same page. For kokumi we anchor the reference to **glutathione** — γ-glutamyl-cysteinyl-glycine, the archetypal γ-glutamyl tripeptide, the one every cell already carries. Against a fixed 1.0 mM extracellular calcium background, we set the convention reference at glutathione's EC_{50} of **2.0 mM (2,000 µM)** — a stated reference constant, not a literature-derived value — and define that as **IFI-K 1,000** by convention. Everything else on the kokumi axis is read against it.

The formula then becomes a simple inversion. With the reference fixed at 2,000 µM:

> $\text{IFI-K} = 1{,}000 \times 2{,}000\ \mu\text{M} \div EC_{50}(\mu\text{M}) = 2{,}000{,}000 \div EC_{50}(\mu\text{M}),$

so a candidate's EC_{50} falls straight out of its score, $EC_{50}(\mu\text{M})$ = 2,000,000 ÷ IFI-K. Run the measured garlic numbers through it and the biology turns into arithmetic. A conventional culinary clove reads **IFI-K 520**, which is an effective half-activation at 2,000,000 ÷ 520 = **3,846 µM** — weaker, note, than pure glutathione, because a whole clove's kokumi is diluted across everything else in the cell. Our engineered kokumi clove reads **IFI-K 3,200**, an EC_{50} of 2,000,000 ÷ 3,200 = **625 µM** — more than three times as loud as the glutathione reference, and six times as loud as the conventional clove. The individual γ-glutamyl peptides line up the same way, each one's score a direct readout of how tightly it holds the receptor's attention.

KOKUMI PEPTIDE	Γ-LINKAGE CARRIES	EC_{50} AT CASR (MM)	IFI-K
Glutathione (γ-Glu-Cys-Gly) — *reference*	Cys-Gly	2,000	1,000
γ-Glu-Cys (store precursor)	Cys	2,500	800
γ-Glu-*S*-methyl-cysteine	methyl-Cys	1,250	1,600
γ-Glu-*S*-allyl-cysteine (garlic's main store)	allyl-Cys	625	3,200
γ-Glu-valine	Val	435	4,600
γ-Glu-valyl-glycine	Val-Gly	95	5,000*

*γ-Glu-Val-Gly's true EC_{50} of 95 µM computes to a raw IFI-K of ~21,000; the scorecard clamps every axis at the normalization maximum of **5,000**, so this peptide pins the ceiling. It is the loudest kokumi molecule known, and a useful reminder that the Index measures loudness, not virtue — a trace of it does the work of a spoonful of weaker peptide.

Notice what the table makes plain: garlic's dominant storage species, γ-glutamyl-*S*-allyl-cysteine, is intrinsically **3.2 times** as potent at the receptor as the glutathione reference (2,000 ÷ 625 = 3.2). The sulfur side-group that the plant grew as a defensive precursor is, by lucky geometry, exactly the decoration that makes the γ-glutamyl cap read louder at CaSR. Garlic's weapon-chemistry and its kokumi-chemistry are, again, the same chemistry.

The ladder of the clove — conventional to engineered

The climb from a conventional clove to our engineered line is a factor of 3,200 ÷ 520 = **6.15×** in the quiet thing, and it is worth taking apart, because it is not one trick but two multiplied together. A kokumi score rises for only two reasons: the delivered molecules are individually louder (*quality*), or there are more of them in the delivered pool (*quantity*). We can read each contribution off the numbers.

The **quality** term is the shift in which peptide dominates. A conventional clove's kokumi is carried mostly by glutathione-type peptides near the reference potency; our line shifts the balance toward γ-glutamyl-*S*-allyl-cysteine, worth 2,000 ÷ 625 = **3.2×** on its own. The **quantity** term is everything left over: 6.15 ÷ 3.2 = **1.92×**, the deepening of the delivered pool itself. The design identity is exact and multiplicative:

> `ΔIFI-K(total) = ΔIFI-K(quality) × ΔIFI-K(quantity) → 6.15 = 3.2 × 1.92.`

Both hands pull the same lever from different ends. We shift the pool toward the *S*-allyl species by routing more flux through the γ-glutamylation step, and we deepen the pool by slowing the GGT branch-point so that less of it is spent toward pungency between the field and the pan. Neither edit touches a capsaicin gene; garlic carries no capsaicinoids at all, so its heat on the Scoville scale is a true **zero SHU**, and the warmth you sometimes feel is a trace, a floor, not Scoville fire — the clove's bite lives entirely on TRPA1, never on TRPV1.



A clove in cross-section — the γ-glutamyl warehouse is stored in the living tissue, stable and odorless until the knife divides the cell.

THE CLOVE, LADDERED	IFI-K	EC_{50}-K (MM)	DOMINANT KOKUMI SPECIES	IFI-P (BITE)
Conventional culinary clove	520	3,846	glutathione-type	high
Engineered kokumi clove, raw	3,200	625	γ-Glu-*S*-allyl-cysteine	high (raw)
Same clove, slow-roasted to paste	2,600	769	γ-Glu-*S*-allyl-cysteine	63

Here is the raw engineered clove as a full fingerprint — the loud bite of raw garlic sitting right on top of the deep, quiet kokumi floor that is the whole point of the line:

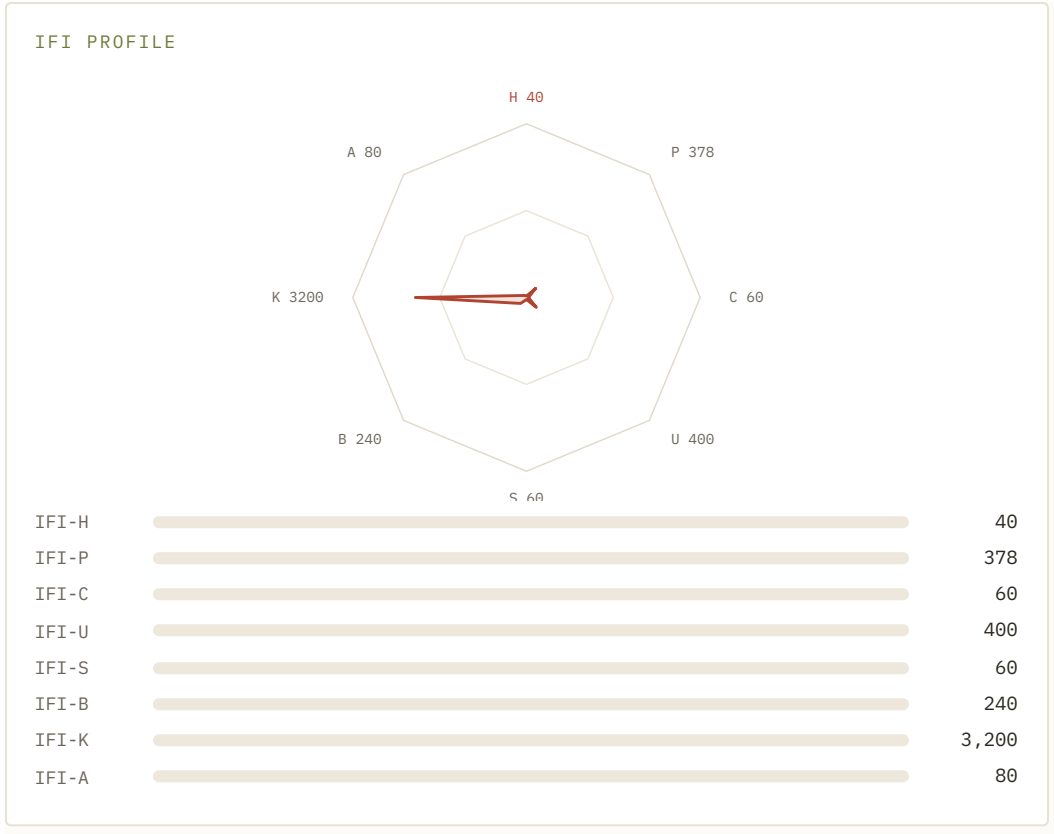

The raw clove is a paradox in eight numbers: pungency (P) up at 378, the allicin alarm firing hard on TRPA1, and kokumi (K) at 3,200 underneath it, the storage pool the knife has not yet spent. A clove bred only for bite — the industrial default — runs the storage pool down fast and hard toward allicin, spending the kokumi to maximize the P. Our line does the opposite: it carries both at once, so you can cut it raw into a sharpness that still has a floor of body under it, or you can turn the whole profile toward depth with heat. That turning is the next thing to understand.

Dose, response, and the kokumi threshold

Because CaSR is a cooperative dimer, kokumi does not fade in gently — it *arrives*. The dose-response follows a Hill relation,

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

with a measured Hill coefficient of about **n = 2.3** for the γ-glutamyl potentiation at a fixed calcium background. That exponent is the steepness, and it compresses the whole on-ramp. Solve the relation for the concentration that reaches any given response fraction,

$$[A] = EC_{50} \times (R \div (1 - R))^{(1/n)},$$

and run the engineered clove's EC_{50} of 625 µM through it: the receptor passes 10 % of full kokumi at **240 µM**, 20 % at **342 µM**, half at **625 µM** by definition, 80 % at **1,142 µM**, and 90 % at **1,625 µM**. The entire climb from barely-there to nearly-full spans only 342 → 1,142 µM — a factor of **3.34×** in concentration from 20 % to 80 %.

That 3.34 is itself a formula worth keeping: the 20-to-80 span of any Hill curve is 16^(1/n). For a non-cooperative receptor (n = 1) it would be a lazy 16-fold; the dimer's cooperativity folds that down to 3.3-fold. This is the molecular reason kokumi has the quality cooks describe as a thing that "clicks in" — a dish that was almost there and then, with one more spoonful, suddenly *coheres*. You are not imagining the switch. You are feeling a Hill coefficient of 2.3.

From clove to pot — the arithmetic of a single tablespoon

The density is why a little garlic goes so far, and the claim can be made exact. Garlic's dominant storage peptide, γ-glutamyl-*S*-allyl-cysteine, has a molar mass of **290.3 g/mol**. Its half-activating concentration of 625 µM, converted to mass, is 625 × 10^{-6} mol/L × 290.3 g/mol = 0.181 g/L, or **181 mg/L**. That is the number to hold: a little under two hundred milligrams of this one peptide, dissolved in a litre, brings the calcium receptor to half of its full kokumi response.

Now weigh a tablespoon of the roasted paste. It is a concentrated thing — the long roast has driven off water and the engineered store is deep — and our model predicts approximately **12 g/kg (1.2 % by weight)** γ-glutamyl-*S*-allyl-cysteine equivalent. A rounded tablespoon is about **16 g**, so it carries 0.012 × 16,000 mg = **192 mg** of the peptide. Divide by the 181 mg/L threshold: that single tablespoon can carry **1.06 litres** of otherwise-flat food all the way to the half-activation point of CaSR. One spoon, roughly half a roasted head, enough kokumi to lift a whole pot of beans or a litre of broth across the threshold where savor switches on.

This is what the cook has always done by feel. A sofrito, a masala base, the first hiss of garlic in oil — these are all, in receptor terms, the identical move: loading the dish with γ-glutamyl peptides before anything else goes in, so that everything after lands on a floor of richness. Kokumi is the integrator, the molecule that makes a plate of separate things taste like one cared-for thing. The cook does it by instinct. We can now do it by number, and the number is 181 mg/L.

The allium among its kin — why garlic wears the crown

Garlic is the flagship of this axis for a reason the construct count makes plain: at **560 constructs** it carries more engineering than any other kokumi organism in the Library, because it has the deepest native γ-glutamyl machinery to build upon. Its relatives share the chemistry but not the depth.

KOKUMI ORGANISM	CONSTRUCTS	PRIMARY AXES	KOKUMI ROUTE
Garlic	560	K / P	deep γ-glutamyl-*S*-allyl-cysteine store
Soybean	360	U / K	γ-glutamyl peptides + glutamate (ferment-ready)
Yeast	240	U / K	glutathione + peptide pool
Onion	40	K	shallower γ-glutamyl store, sweeter sulfur
Vanilla	40	K	kokumi mouthfeel beneath vanillin

Read down the column and the allium family tells its own story: garlic and onion share the pathway, but garlic holds fourteen times the construct depth and a storage pool to match, which is why a clove changes a pot where a slice of onion only sweetens it. Soybean and yeast reach kokumi by a different door — they carry glutathione and free glutamate together, which is why a ferment reads on both the umami and kokumi axes at once and why miso, koji, and aged things taste *built*. Vanilla's kokumi is the faint body beneath its aromatics, the reason real vanilla feels round where the synthetic single-note feels thin. Garlic sits at the head of all of them because, alone among the kin, it grew the deep store as a weapon-magazine and left us the richest safe to open.



Onion, the sister allium — the same γ-glutamyl chemistry at a shallower depth, which is why it sweetens a dish where garlic founds it.

Error, reproducibility, and the honest limits

Numbers are a form of respect for the reader, so they come with their error bars. The storage-pool chemistry of a conventional clove is volatile — field conditions, storage time, and the moment of cutting all move it — and a wild-type allium's kokumi output swings at a batch coefficient of variation near **42 %**. Our constructs hold to a batch **CV of ≤ 7 %**. That tightening is the quiet half of the engineering, and it propagates cleanly into the Index. Because IFI-K is a quotient of two measured EC_{50}s, its own variation combines in quadrature,

$$CV(\text{IFI-K}) = \sqrt{CV(\text{reference})^2 + CV(\text{candidate})^2}.$$

With a reference held to 3 % and a candidate line at 7 %, the score carries $\sqrt{(0.03^2 + 0.07^2)} = \mathbf{7.6\ \%}$. For the engineered clove at IFI-K 3,200 that is ±244 — a batch-to-batch range of **2,957 to 3,443**. Set that beside a wild-type clove at IFI-K 520 and 42 % CV, which swings ±218, a range of **302 to 738**. The engineered line's *worst* batch (2,957) still reads four times louder than the wild clove's *best* (738). Reproducibility is not a footnote to the richness; it is the richness, made dependable enough to cook and formulate against.

Honesty asks for the limits, too, and kokumi has three worth naming:

- **It needs a partner.** Kokumi is a positive allosteric modulator — it potentiates a receptor already bathed in calcium and a tongue already carrying other tastes. In plain water the engineered clove underdelivers; its 3,200 is a loudness with nothing to amplify. It shines only in a composed dish, under umami, salt, or sweetness it can thicken. This is why the paste is a foundation and not a dish.
- **The branch-point can steal the gains.** Deepen the γ-glutamyl store but leave GGT and alliinase running fast, and the knife simply makes *more allicin* — more pungency, not more kokumi. The deep pool and the slowed conversion must move together; one without the other is a louder weapon, not a richer gift.
- **Heat is a trade, and overheat is a loss.** Roasting frees glutamate and tames the bite, but it also degrades some γ-glutamyl peptides through the Maillard route — which is exactly why the slow-roasted paste reads K 2,600, a little below the raw clove's 3,200. Push the oven too hot and the trade turns against you: the store burns off, bitterness climbs, and the acrid note takes over. The

recipe's low, patient 150 °C is the whole mitigation. (A methods note for the bench, too: glutathione reference oxidizes to its disulfide and quietly loses activity, so the reference control is made fresh for every assay, or the whole axis drifts upward under you.)

None of this is a health claim, and we are careful never to make one. Kokumi is a flavor the body reads as depth; that it lets a cook reach for less salt, or coax savor from a pot of plain beans and greens, is a kindness of the chemistry, not a cure. We revere the clove for what it is: a plant that grew a defense so elegant it reads, on our tongues, as being cared for.

The method on the plate — the roasted paste

Garlic's kokumi is best understood by making it, so here is the foundation we build a hundred dishes on — a slow-roasted paste that is almost pure richness, the pungency cooked down to a whisper, the γ-glutamyl store turned toward body and held there by the long low heat. It is the single most efficient kokumi move a kitchen keeps in a jar.

```
ROASTED GARLIC KOKUMI PASTE
──────────────────────────────
4 whole heads     engineered kokumi garlic (our line)
3 tbsp            cold-pressed olive oil
1 tsp             flaky salt
2 sprigs          thyme
──────────────────────────────
yield  ~120 g paste   ·   oven 150°C   ·   time 90 min
```

1. Slice the top quarter off each head to open every clove. Sit the heads cut-face up in a small dish, pour the oil over so it pools into the cloves, scatter the salt, lay the thyme across.
2. Cover tight with foil and roast at **150 °C for 75 minutes** — low and slow is the whole method, because it is the temperature that lets the γ-glutamyl store caramelize and sweeten without driving the alliinase reaction toward sharpness. Rushed hot, you make a harsher, more bitter thing.
3. Uncover, raise to **190 °C for 15 minutes**, until the cut faces are amber and the cloves slump.
4. Cool until you can handle them, then squeeze every clove out of its skin into a bowl and mash smooth with the roasting oil. It keeps, covered, a week in the cold; it only gets rounder.

IFI PROFILE

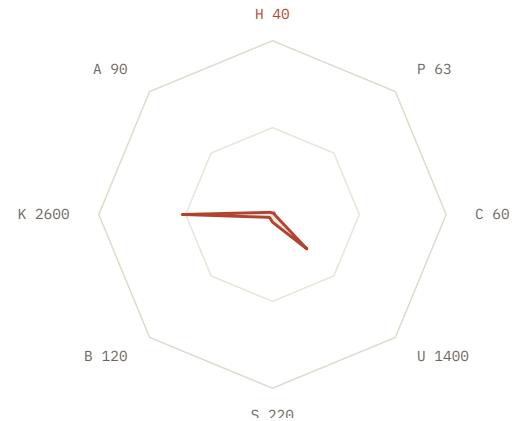

IFI-H	40
IFI-P	63
IFI-C	60
IFI-U	1,400
IFI-S	220
IFI-B	120
IFI-K	2,600
IFI-A	90

Set this fingerprint beside the raw clove's and you can watch the roast as a transformation in eight numbers. Pungency (P) falls 378 → **63**, the allicin alarm cooked down to a shadow — the heat GGT and alliinase would have made is simply never made, because the slow warmth outruns them. Umami (U) climbs 400 → **1,400** as the heat frees bound glutamate. Sweetness (S) rises 60 → **220** as sugars caramelize. And kokumi (K) settles at **2,600** — a little below the raw 3,200, the honest cost of the Maillard trade, but still the headline, still the richness the whole method was built to protect, now at an EC_{50} of 769 µM, loud enough that a spoonful carries most of a litre across the threshold.

A spoonful of this stirred into anything — a pan of greens, a bowl of beans, a sauce that was almost but not quite there — is the most efficient kokumi move in the kitchen. It is the floor every savory thing we cook is built upon: a plant's safeguarded defense, roasted patiently into depth, measured now to the milligram, and given back to the body as savor. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 8 · The wider alliums: onion, leek, scallion — *the gentler route.*

Garlic is the allium of intensity; the onion is the allium of patience, and it teaches a different lesson about kokumi — that richness can be *coaxed*, slowly, from a thing that holds very little of it raw. Cut a garlic clove and the intensity is already there, loud in the nose before the knife has left the board. Cut an onion and what you get is tears and a thin, green sharpness; the richness is not in the cut at all. It is folded into the bulb as a quiet chemical promise, and it is released not by the blade but by the pot — by heat, and time, and the kind of attention our family learned to give a dish long before we could measure what that attention was doing. On the axis, a raw onion reads modestly, around **180**, a fraction of garlic's store. The whole of this chapter is the arithmetic of turning that 180 into something that stays with you.

The receptor that reads richness

Kokumi is the strangest of the eight axes because it is the one with no taste of its own. You cannot point to it the way you can point to sweet or to the first warmth of heat. What it does instead is make everything else *more* — rounder, thicker, more continuous, lasting longer on the tongue and arriving with a fullness the French call *bouche* and the Japanese named *kokumi*, "rich taste." It is the difference between a broth that is correct and a broth that tastes like being cared for. And all of it traces to a single protein: the calcium-sensing receptor, **CaSR**.

CaSR is a class C G-protein-coupled receptor, and that family membership is the first and most important fact about it, because the sweet receptor (TAS1R2·TAS1R3) and the umami receptor (TAS1R1·TAS1R3) and the brain's metabotropic glutamate receptors are all class C GPCRs too. They are cousins. They share the same deep architecture — a large, hinged **Venus-flytrap domain** reaching out from the cell surface, a cysteine-rich stalk, and a seven-pass transmembrane bundle that does the signaling — and they work as paired dimers, two receptors clasped together, closing on their ligand like the two lobes of a trap shutting on a fly. This shared ancestry is why kokumi amplifies umami so faithfully: the receptor that reads richness sits in the same structural family as the receptor that reads savor, and when both are lit the tongue reads a single, deepened sensation rather than two separate ones.

The orthosteric trigger of CaSR — the thing it evolved to sense — is extracellular calcium, binding down in the cleft of the flytrap. Kokumi compounds do not compete for that site. They are **positive allosteric modulators**: they bind at the margins of the flytrap and in the transmembrane bundle, and by stabilizing the closed, active shape

of the receptor they lower the calcium threshold needed to fire it. The stimulus on the tongue does not change. The receptor simply becomes easier to satisfy, and the brain reads that greater ease as *more* — more body, more continuity, more of the round depth under everything. CaSR also signals with pronounced positive cooperativity, a Hill coefficient near 3 for calcium, which means its response is not a gentle ramp but a switch with a steep middle; a small nudge to the threshold, delivered by a kokumi peptide, moves the whole steep curve sideways and produces a large change in what you feel.

What are these peptides? They are **γ-glutamyl peptides** — short chains built on an unusual bond. In an ordinary peptide, glutamate links to the next amino acid through its α-carboxyl group. In a γ-glutamyl peptide, the link runs through the side-chain γ-carboxyl instead, and that one structural quirk is the pharmacophore, the part the receptor reads. The archetype is **glutathione**, γ-Glu-Cys-Gly, the tripeptide every living cell carries as its master antioxidant — the molecule the body itself reveres as the keeper of its own clean inner chemistry. There is a quiet grace in that: the compound we taste as richness is the same one the cell uses to stay well. Potency climbs with the second and third residues; the dipeptide γ-Glu-Cys and the tripeptide γ-Glu-Val-Gly are both more eloquent at CaSR than glutathione itself. The alliums are, in the most literal sense, a reservoir of these ligands.



An onion in cross-section — the concentric scales are the vault, each layer holding the γ-glutamyl store that the knife cannot reach and the pot slowly can.

A bulb that stores its own receptor ligands

An allium does not keep its flavor as flavor. It keeps it as a set of odorless, non-volatile precursors, locked away from the enzymes that would act on them, and it assembles the finished molecule only at the moment of injury — a defense evolved against the animal that would eat it, repurposed by us into the foundation of nearly every savory dish there is.

The precursors are **S-alk(en)yl-L-cysteine sulfoxides**, the ACSOs. Onion's dominant one is **isoalliin** (S-1-propenyl-L-cysteine sulfoxide); garlic's is **alliin** (S-allyl-L-cysteine sulfoxide); both carry a little **methiin** (the S-methyl form) as well. And crucially, the bulb stores and shuttles these as their **γ-glutamyl conjugates** — γ-

glutamyl-S-alkenyl-cysteines — which means the storage forms themselves are γ-glutamyl peptides, CaSR ligands in their own right. The enzyme **γ-glutamyl transpeptidase** (GGT) liberates and interconverts them; during sprouting and cold storage it drives the γ-glutamyl pool toward the free ACSOs, which is why an old, sprouting onion tastes different from a fresh one. This is the vault. The richness is in there the whole time, held as peptide.

Now the knife. Damage a cell and it releases **alliinase** — a pyridoxal-5′-phosphate-dependent C–S lyase (EC 4.4.1.4) held apart from its substrate until the tissue is broken. Alliinase cleaves the ACSOs into reactive sulfenic acids, pyruvate, and ammonia, and here the two alliums part ways. In **garlic**, allyl sulfenic acid condenses with itself almost instantly into **allicin**, the diallyl thiosulfinate that is garlic's whole nose-forward signature — which is why garlic's intensity is a thing of the cutting board, present before any heat. In **onion**, 1-propenyl sulfenic acid is seized by a second, dedicated enzyme, **lachrymatory-factor synthase** (LFS), and turned into **syn-propanethial-S-oxide**, the tear molecule. That is the onion's defense: not an aroma in the nose but a sting in the eyes. Its sharpness goes *up*, not out, and — this is the whole point — none of it is the richness. The γ-glutamyl store, the kokumi, is untouched by the knife. It waits for the pot.



Garlic, the intense sibling — its richness is released by the blade, where the onion's must be drawn out by slow heat. Same pathway, opposite temperament.

In our program this difference in temperament is written straight into the numbers of the work. Garlic carries **560 constructs** and the onion only **40**, and that is not a ranking — it is a division of labor. Garlic's job is density: to be the loudest small source of γ-glutamyl richness and the sharp pungency that rides alongside it. The onion's job is to be the *gentlest* — the kokumi you can lay down in volume, under everything, as a base that deepens a dish without ever announcing itself. The forty constructs on our onion line do not chase garlic's density. They deepen and clean the γ-glutamyl store, and tighten its variability, so that the slow method below has more to concentrate and lands on its number every time. The engineering fills the vault; the cooking opens it.

The number, derived

The Iyengar Flavor Index turns every one of these sensations into a single figure by the same construction, anchored to how *loudly* a stimulus speaks at its receptor rather than how much of it a food contains:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. A smaller EC_{50} means a louder voice — less material needed to be heard — so inverting it against a fixed reference, and scaling by a thousand, makes loudness a score. The IFI is a stated convention: each reference constant defines IFI = 1,000 on that axis by convention, in the same spirit as Scoville's dilution-to-threshold definition. For the kokumi axis we fix the reference at reduced **glutathione**, the canonical CaSR agonist, and we set its half-maximal activation at **EC_{50}(ref) = 1.80 mM** by convention. By construction, then, glutathione itself scores **IFI-K 1,000** — the anchor of the axis. Rearranged, the formula reads the other way just as easily:

$$EC_{50}(\text{candidate}) = 1{,}000 \times EC_{50}(\text{ref}) \div \text{IFI-K} = 1{,}800\ \mu M \div \text{IFI-K} \times 1{,}000$$

so every kokumi score we quote implies an exact concentration at which that material half-fills the receptor. The canonical onion figures fall straight out:

MATERIAL	IFI-K	IMPLIED EC_{50} (GSH-EQUIV.)	IN MG/L
Reduced glutathione (reference anchor)	1,000	1.80 mM	553
Raw onion, cut (its active fraction)	180	10.00 mM	3,073
Slow onion stock	1,500	1.20 mM	369
Nine-hour caramelized onion	1,900	0.947 mM	291

Read the first two rows together and the raw onion's modesty is exact: at IFI-K 180 its kokumi-active fraction reaches half-max only at **10.0 mM** glutathione-equivalent — 1,000 × 1.80 ÷ 180 = 10.0 — more than five times the concentration the reference peptide needs, which is to say roughly a fifth as loud. The caramelized onion inverts that story completely: 1,000 × 1.80 ÷ 1,900 = **0.947 mM**, a little louder than free glutathione itself. (The milligram-per-litre column is only a unit change, multiplying millimolar by glutathione's molar mass, 307.3 g/mol: 1.80 mM × 307.3 = 553 mg/L, and so on down the column — the same quantities an analyst actually weighs out.)

The EC_{50} is one point on a curve, and it is worth seeing the curve, because kokumi is not an on/off sensation. CaSR follows a Hill relationship,

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

where θ is the fraction of maximal receptor activation, [A] the concentration of the kokumi principle in glutathione-equivalents, and *n* the Hill coefficient — the steepness. For the γ-glutamyl potentiation of CaSR, published studies report an apparent slope of

about *n = 1.4*. Taking the caramelized onion (EC_{50} = 0.947 mM) and walking [A] in doubling steps around that half-max gives the dose–response the index compresses into a single number:

[KOKUMI PRINCIPLE], GSH-EQUIV.	RATIO [A]/EC_{50}	θ (FRACTION OF MAX)
0.24 mM	0.25	0.13
0.47 mM	0.50	0.27
0.95 mM	1.00	0.50
1.89 mM	2.00	0.73
3.79 mM	4.00	0.87

The middle row is the EC_{50}, θ = 0.50, by definition — that is the point the index reports. What the full table shows, and the single figure cannot, is the *shape*: doubling the dose from half-max carries you to 0.73, not to 1.0, and you spend four-fold more material to crawl from 0.73 to 0.87. Richness has steeply diminishing returns, which is the receptor-level reason a dish can be made rounder and rounder only up to a point, and then no further — the flytrap is nearly shut, and more peptide has almost nothing left to close.

The same Hill relation makes exact what it means for a kokumi peptide to be an allosteric modulator rather than a flavor of its own — the single most important and least intuitive fact about this axis. Hold the stimulus fixed and move the receptor instead. Suppose the calcium bathing the taste cell sits at 2.0 mM while CaSR's intrinsic half-max for calcium is 3.0 mM, at the receptor's cooperative *n ≈ 3*. Then baseline activation is $\theta = 2^3 \div (3^3 + 2^3) = 8 \div 35 =$ **0.23** — the tongue's quiet resting read. Now add enough γ-glutamyl peptide to left-shift that calcium EC_{50} down to 2.0 mM, which is all an allosteric agonist does: at the *same* 2.0 mM calcium the activation becomes $\theta = 2^3 \div (2^3 + 2^3) =$ **0.50**. The calcium never changed. Nothing new was tasted. And the receptor's output more than doubled — 0.23 to 0.50 — which the brain can only report as *more*: more body, more fullness, more of the sensation lasting. That doubling, delivered by a peptide the tongue cannot taste on its own, is kokumi, written as arithmetic.

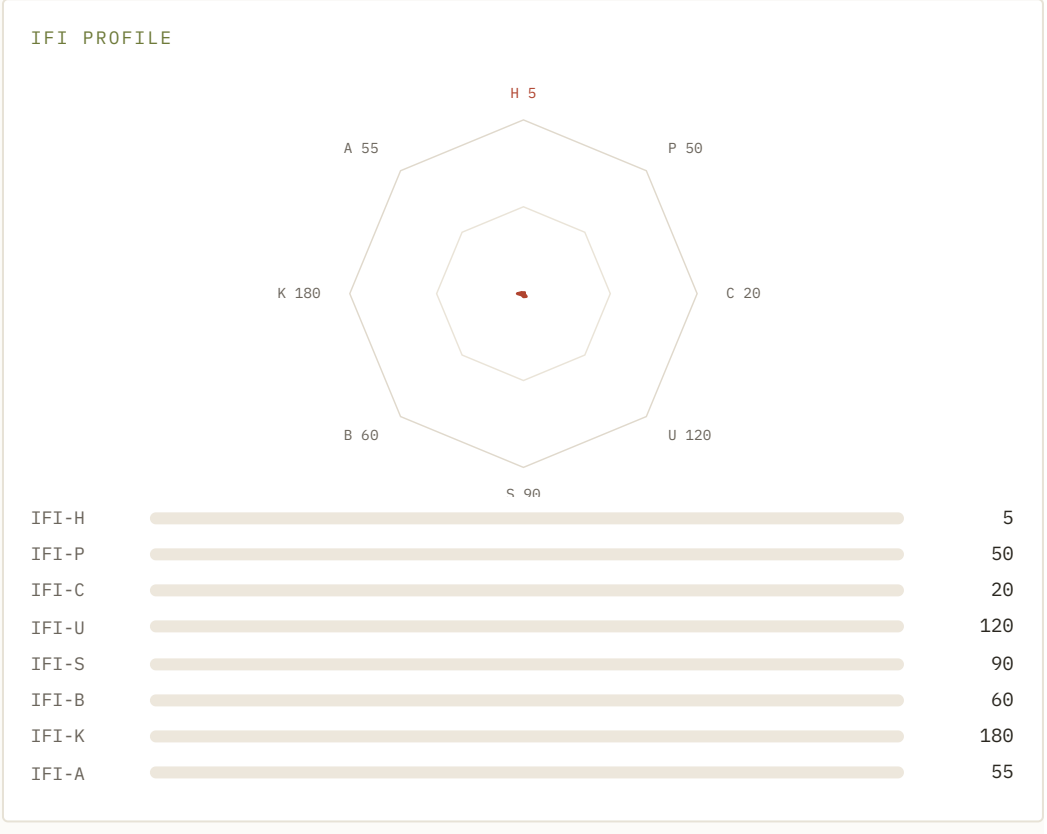

That is the raw bulb on all eight axes: a tall spike of pungency (**P 50** — the LFS lachrymator and the raw sulfur bite), almost nothing of heat or cool, and a kokumi reading of just **180**. Everything that matters is latent. The pot is what collects the debt.

What browning actually does, counted

The single most dramatic kokumi transformation available at any stove has a name every cook knows — **caramelization** — and it is worth taking apart into the exact chemistry, because two distinct reactions are running at once and only one of them is what the receptor reads.

The first is caramelization proper: the onion's own sugars, under long heat, fracturing into hundreds of brown, aromatic compounds. That is sweetness turning into complexity, and it is most of the color and much of the perfume. The second — the one that builds kokumi — is the **Maillard reaction**, between those same sugars and the onion's free amino acids, which generates savory, roasted, almost meaty compounds *and*, alongside them, drives GGT-style transpeptitation that concentrates and reshapes the γ-glutamyl pool, shifting it toward the more potent peptide forms. Underneath both, as water leaves, everything simply grows more concentrated. And all the while the harsh raw-sulfur compounds and the lachrymator — volatile, heat-labile — break down and boil off, taking the sharpness out of the eyes and out of the dish.

We can account for the whole climb from 180 to 1,900 as a product of two honest factors. A kilogram of raw onion is near 89% water; cooked low and slow for nine hours to a dark, jammy mass, it sheds most of that water and falls to roughly 200 g — a **five-fold concentration** by mass. Concentration alone, holding the chemistry fixed, lowers the effective EC_{50} five-fold and so lifts the score about five-fold, from 180 toward 900. The remaining factor to reach 1,900 — a further **2.1×** — is the chemistry proper: the new kokumi-active Maillard species, the transpeptidation shift of the γ-glutamyl pool toward its louder members, and the clearing-away of the sharp sulfur compounds that otherwise mask and compete.

STAGE OF THE SLOW COOK	WHAT HAPPENS, CHEMICALLY	IFI-K	STEP FACTOR
Raw bulb	the latent γ-glutamyl store, masked by sharp sulfur	180	–
Water driven off (~5×)	peptides concentrate as mass falls five-fold	~900	×5.0
Maillard + transpeptidation	new kokumi species; pool shifts to potent forms; harsh sulfur volatilizes	1,900	×2.1

Multiply the two factors — 5.0 × 2.1 ≈ **10.6** — and you have the whole gain: 180 × 10.6 ≈ 1,900, a more-than-tenfold lift, most of it unlocked by the cooking, the engineering having deepened the store the cooking then concentrates. This is why the first motion of so many savory dishes across so many cuisines is the same — onion, in fat, until brown — and why skipping it leaves a dish that is technically

correct and somehow hollow. The browning is not a flourish. It is the kokumi foundation, laid in the first ten minutes, and it is also, read honestly, where the number is *made* rather than merely given.



The onion whole — the allium of patience, whose richness is not in the cut but in the long, low heat that follows it.

IFI PROFILE

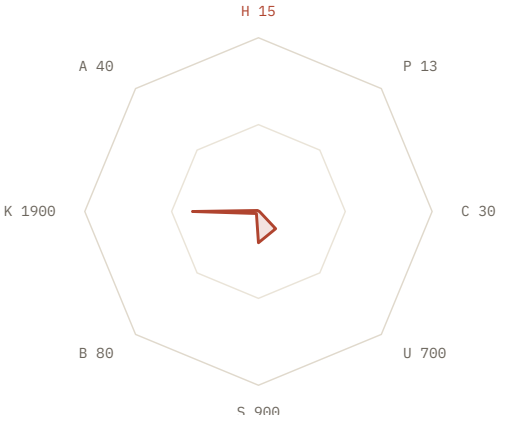

IFI-H	15
IFI-P	13
IFI-C	30
IFI-U	700
IFI-S	900
IFI-B	80
IFI-K	1,900
IFI-A	40

Set that caramelized profile against the raw one and the transformation is legible at a glance: pungency collapses from 50 to **13** as the lachrymator and harsh sulfur boil away; sweetness climbs from 90 to **900** with caramelization; umami rises to **700** on the Maillard savor; and kokumi arrives at **1,900**, the whole patient point of the exercise. Nothing was added. Water left, and chemistry happened, and a modest bulb became one of the richest vegetarian foundations in all of cooking.

The allium ladder — a palette of speeds

The gentler route is not one road but several, because the onion's near relatives run the identical γ-glutamyl pathway at different densities, different water contents, and — most usefully for a cook — different *speeds*. In our model, raw and slow-cooked values for our engineered lines lay out as a ladder, and it is the final column, the cook-gain, that carries the argument:

ALLIUM (ENGINEERED LINE)	RAW IFI-K	SLOW-COOKED IFI-K	COOK GAIN	CULINARY ROLE
Garlic — the intense sibling	820	2,400	2.9×	assertive; releases on the knife, not the pot
Shallot	240	2,050	8.5×	fast and sweet; the kokumi of a pan sauce
Scallion (white)	210	1,400	6.7×	sharp and immediate; a finishing note
Yellow onion	180	1,900	10.6×	the patient base; stocks, soffritto, foundations
Leek (white and pale green)	120	1,250	10.4×	gentle and green; will not dominate

Read down the cook-gain column and you read the whole thesis of the gentle route. Garlic begins loud — 820 raw — and the pot barely doubles it, because garlic's richness is a knife phenomenon, released and largely finished at the cutting board. The onion and the leek begin quietest of all, 180 and 120, and the slow method *multiplies them more than tenfold*. That is the gradient a serious kitchen is really keeping when it keeps all of them: not five similar ingredients but five different relationships between time at the stove and depth on the tongue. The shallot gives you most of the gain fast, browning sweetest and soonest — the right choice when a sauce has minutes, not hours. The scallion's white carries a sharp, high raw reading you deploy uncooked, at the end, for a bright edge. The leek offers a soft, green depth that will sit under a dish without ever stepping forward. And the yellow onion is the workhorse of the base, the one whose patience pays the largest dividend, the reason a stock built on it tastes finished when a stock built on water does not.

The ladder answers to the same formula as everything else on the axis, and it is worth doing the arithmetic to feel how far apart these raw bulbs really are. Leek at IFI-K 120 implies a half-max of 1,000 × 1.80 ÷ 120 = **15.0 mM** glutathione-equivalent — the quietest active fraction in the family, which needs the most material to be heard, exactly why it will not dominate. Shallot at IFI-K 240 implies 1,000 × 1.80 ÷ 240 = **7.5 mM**, precisely twice as loud as the leek and the reason it reads sweet and forward in a fast pan. Garlic's raw 820 works out to 2.20 mM, already near the reference peptide's own 1.80 mM before a drop of heat. The ladder is not a metaphor for the gradient; it is the gradient, in millimolar, and the cook simply chooses where on it to begin.

Error, honest limits, and the edges

A number is only a form of respect for the reader if it is also honest about what it is, so a few words on precision and on what the kokumi axis does not say.

Precision. Our engineered onion line is designed to hold its kokumi reading to a batch coefficient of variation of **≤7%** — against the 42% batch-to-batch swing of wild-type bulbs, whose γ-glutamyl content lurches with cultivar, soil, storage time, and whether the bulb has begun to sprout. Seven percent on a caramelized reading of 1,900 is **±133**, a window of roughly 1,767 to 2,033. That is the figure for a calibrated reference and a controlled onion. If one also propagates uncertainty in the reference standard itself, two independent 7% sources combine in quadrature — $\sqrt{(0.07^2 + 0.07^2)} \approx$ **9.9%** — widening the band to about ±188; in practice the reference is a tightly held calibrator, so the candidate's ≤7% dominates and ±133 is the honest claim. The point of the engineering, as much as the depth, is that the number *repeats* — that a formulator can build on an onion the way one builds on a weighed reagent, which is precisely what a wild onion never allowed.

What the axis does not capture. IFI-K is a loudness at one receptor under a fixed assay, and three real things fall outside it. First, **aroma**: the volatile sulfur and Maillard notes that make a caramelized onion smell the way it does are a separate sensory channel entirely, read by the nose, not by CaSR, and the kokumi number says nothing about them. Second, **time**: kokumi in the mouth is partly a phenomenon of *continuity* and persistence — the sensation that lasts — and a single EC_{50} compresses that whole temporal shape to one scalar; the Hill table above is as close as the index comes to admitting the curve underneath. Third, **matrix**: fat, salt, and co-present umami all change how much kokumi you perceive from the same peptide, and the assay fixes a clean background that a real pot never has. The index is a true construction of one dimension, not the whole of what a broth does to you.

Failure modes. The gentle route has edges you can fall off. Push browning past Maillard into pyrolysis — a scorched, bitter onion — and the bitterness axis climbs while the kokumi you were building is masked, not gained; the window closes abruptly, and there is no recovering a burnt base. A hard, rolling **boil** in place of a bare simmer does not destroy the peptides but shears and clouds the liquid, emulsifying fine solids into a coarse, muddy stock whose richness reads lower than its chemistry deserves. And a **sprouted** bulb has already let GGT shift its γ-glutamyl pool toward the free sulfoxides, trading some of its latent kokumi for sharpness before you ever light the burner. Each failure is the same lesson from a different side: the richness is real but conditional, and the condition is patience held at the right temperature.

The stock, built in water

Here is the clearest demonstration we know that kokumi is a thing you build slowly, out of patience, in nothing but water — no fat, no body of its own, only the allium's released store of peptides speaking to the receptor.

SLOW ONION STOCK

1 kg	engineered kokumi onion, halved
2	leeks, white and pale green, split
1 head	garlic, halved across the equator
2 tbsp	olive oil
2 L	water
1 sheet	kombu (optional, for an umami floor)

yield ~1.6 L · time 3 h total

1. Sit the onion halves cut-face down in a dry heavy pot over medium heat, no oil yet, and leave them untouched for **12 minutes** until the cut faces are deeply, evenly browned — almost too dark. That fond is pure concentrated kokumi, Maillard laid directly onto the pan; do not rush it and do not stir it, because stirring is how you trade browning for steaming.
2. Add the oil, leeks, and garlic, lower the heat, and cook **20 minutes**, stirring now, until everything is soft and sweet and the pot bottom is layered with brown. The garlic goes in here, late and whole-faced, so its knife-released sharpness never fully forms — you want its depth, not its bite.
3. Pour in the water, scraping up every bit of fond so none of that made kokumi is left on the metal. Add the kombu if using, for a glutamate floor the γ-glutamyl peptides will stand on. Bring it barely to a simmer — never a boil, which clouds and coarsens — and hold it there, uncovered, **2 hours**, until the liquid has reduced by a third and gone the color of weak tea.
4. Strain, pressing the solids hard to recover every drop of the γ-glutamyl liquor they are still holding. Season at the very end, when you can finally taste what the patience built.

IFI PROFILE

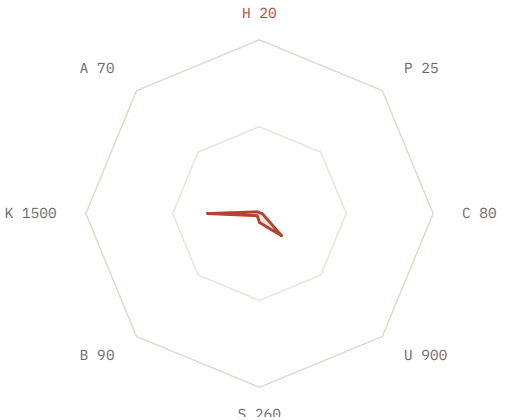

IFI-H	20
IFI-P	25
IFI-C	80
IFI-U	900
IFI-S	260
IFI-B	90
IFI-K	1,500
IFI-A	70

A kokumi of **1,500** in a clear liquid with no fat and almost no body of its own — in EC_{50} terms, a half-max at 1.20 mM glutathione-equivalent, which is to say the stock's active fraction is half again louder than free glutathione itself. That is the number our grandmother was chasing when she said the soup *stays with you*. There is no mystery left in the staying, only the arithmetic of it:

water, patience, and an allium's slow-released store of peptides left-shifting the calcium receptor until the tongue reads depth where the eye sees only pale liquid. Use it anywhere you would use water and the dish will feel, inexplicably, more finished — except that it is not inexplicable anymore. It is 1,500 on the axis, give or take seven percent, and it is the gentlest route to the richest thing a kitchen can make.

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

Chapter 9 · Yeast & koji — *the richness factories.*

The organism of bread is the organism of richness

A loaf going slack and domed under a towel; the cidery breath of a crock three days into its work; the soft collapse of foam as a beer finishes and settles — the organism behind all of it is *Saccharomyces cerevisiae*, and read a second way it is the most concentrated source of richness our family grows. The plants store kokumi, a little of it, folded into a clove of garlic or a vanilla pod. The microbes **manufacture** it, and at a density nothing green can match.

The reason is a molecule the cell cannot live without. Yeast builds **glutathione** — the tripeptide γ-glutamyl-cysteinyl-glycine, γ-Glu-Cys-Gly — not to flavor anything but to defend itself: it is the cell's own redox currency, the small reservoir of reducing power it spends to keep its interior from oxidizing and its proteins intact. A yeast cell under good conditions carries glutathione at a substantial fraction of its dry weight, because the need is constant and the stakes are survival. That survival chemistry is, from the kokumi engineer's bench, a gift. The same compound the cell hoards to stay alive is one of the purest agonists of the receptor that reads richness. Grow the cell, and you are growing kokumi whether you meant to or not.

This is why **yeast extract** — the autolyzed, broken-open contents of yeast cells — is the single most important industrial kokumi ingredient in the world: the brown savory paste and the tan powder standing behind an enormous share of processed savory food. Read plainly, it is a concentrated delivery of glutathione and free glutamate together, kokumi and umami in one spoonful. Beside yeast stands **koji**, *Aspergillus oryzae*, the noble mold of miso and soy sauce and sake — not a store of richness but a factory of the enzymes that *make* it. Between the two organisms sits the whole engine of fermented depth: one that contains richness, one that builds and redistributes it. This chapter is about teaching both to speak the richness axis louder, and about the mathematics that lets us say exactly how much louder.



The sequencer bench — where a yeast line's glutathione pathway is read, and the feedback-relief edits are characterized, before a single tank is pitched.

What the calcium receptor is listening for

Richness is not a taste. It has no flavor of its own, which is precisely why it took so long to name. What we call kokumi — mouthfulness, continuity, the long round backbone under a good broth — is the work of a single protein reading the mouth: **CaSR**, the calcium-sensing receptor, the IFI-K axis. It is a class C G-protein-coupled receptor, built as a homodimer, each half carrying a large bilobed extracellular domain shaped like a *Venus flytrap*. The flytrap's natural job is to clamp shut around extracellular calcium and the aromatic amino acids; clamping shut is the "on" signal, which travels down through a seven-helix membrane bundle to a Gq protein, to phospholipase C, to a rise of calcium inside the cell.

Kokumi compounds work on this receptor not by being tasted but by making it easier to close. Glutathione, γ-Glu-Val-Gly, the other γ-glutamyl peptides — they bind at or beside the flytrap's amino-acid site and stabilize the shut, active conformation. In pharmacological terms they are **positive allosteric modulators**: they do not carry the signal themselves so much as lower the price of the signal, sliding the receptor's whole dose-response for calcium and amino acids to the left. That leftward slide is the entire sensation. It is why a spoonful of broth with no new flavor in it makes everything already present taste deeper, rounder, and longer — the receptor, potentiated, now reports a fuller mouth from the same tastants. It is also why the kitchen test at the end of a good kokumi preparation is *no single flavor stepping forward*: the compound that did the work was, by design, flavorless.

The pharmacophore — the business end of every one of these peptides — is the **γ-glutamyl group**. The glutamate is joined to the next residue not through its α-carboxyl, the way an ordinary peptide bond runs, but through the γ-carboxyl on its side chain. That unusual bond is what CaSR's flytrap is shaped to recognize, and it is also what ordinary proteases cannot cleave, so the γ-glutamyl peptides survive a ferment that dismantles everything around them. Hold that detail. It is the hinge on which koji's whole contribution turns, because the only enzymes that build and move γ-glutamyl bonds are the γ-glutamyltransferases — and those are exactly what the mold brings.

Reading IFI-K as a number — from the index back to the receptor

Because kokumi is a left-shift of a dose-response, it is unusually clean to put a number on. The Iyengar Flavor Index scores every axis the same way:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. For the kokumi axis, read it as the half-maximal concentration of the potentiating effect — the dose at which a compound delivers half of its own maximum left-shift of CaSR. Invert it against a fixed reference, scale by a thousand, and loudness

becomes a score. The index is anchored so that **IFI-K 1,000 is the native wild-type reference**: at 1,000, a candidate's half-activation sits exactly at the reference's, EC_{50}(c) = EC_{50}(ref).

That anchoring hands us an identity worth keeping in the hand, because it makes every score legible on sight. Rearrange the formula and the EC_{50} ratio falls out as 1,000 divided by the score, and the fold-potency — how many times more potent than the reference — is simply the score over a thousand:

SOURCE	AXIS	IFI	FOLD-POTENCY (IFI ÷ 1000)	EC_{50} ÷ REFERENCE (1000 ÷ IFI)
Reference anchor	—	1,000	1.00×	1.000
Phoenix · SS-0096	K	1,900	1.90×	0.526
Kokumi broth base	K	2,400	2.40×	0.417
Phoenix · SS-0096	U	1,800	1.80×	0.556
Kokumi broth base	U	2,000	2.00×	0.500

So the broth base this chapter ends on, scoring **K 2,400**, reaches half-activation of the richness receptor at 1,000 ÷ 2,400 = 0.417 of the reference concentration — it is 2.40× more potent. To put that in units a formulator can weigh out, set the reference anchor at 1.00 mmol·L^{-1} of glutathione-equivalent (a scale chosen only to carry the arithmetic; the index itself is dimensionless by normalization). Then the broth's half-activation lands at 0.417 mmol·L^{-1}, and with glutathione's molar mass of 307.3 g·mol^{-1} that is 0.417 × 307.3 = **128 mg·L^{-1}**. Expressed instead as γ-Glu-Val-Gly, molar mass 303.3 g·mol^{-1}, it is 126 mg·L^{-1} — the two peptides are near-isobaric, so the mass barely moves even though their potencies on the receptor do not; the index has already absorbed that difference into the score.

The left-shift is where the number becomes a sensation. Model CaSR potentiation as a sigmoid with a Hill coefficient — the dimeric receptor reads cooperatively, so take n ≈ 2 — where the response R at concentration [A] is $R_{max}\cdot[A]^n \div (EC_{50}^n + [A]^n)$. Walk it out in multiples of EC_{50}: at half the EC_{50} the receptor is at 0.25 ÷ 1.25 = 20% of maximum; at the EC_{50}, 50%; at twice, 4 ÷ 5 = 80%; at three times, 90%. The steep middle of that curve is the whole game. Now fix a tasting concentration at the reference EC_{50} — one unit — and compare. The reference compound sits at its own half-point, 50%. The engineered broth, 2.40× more potent, meets that same dose at 2.40 EC_{50} of its own, and reads $2.40^2 \div (2.40^2 + 1)$ = 5.76 ÷ 6.76 = **85%**. A 2.4-fold gain in potency became a 35-percentage-point jump in receptor response, because the dose was sitting on the steep part of the curve. That jump, repeated a billion times a second across the tongue, is what the mouth finally calls *deep, round, and long*.

Two enzymes, two ATP, and the two brakes we reached

The glutathione route is short, which is its charm — two enzymes, two molecules of ATP, and the product falls out. It is also a route whose brakes sit where an engineer can reach them.

STEP	GENE	ENZYME	REACTION	ATP	NATIVE BRAKE	PROGRAM EDIT
1	**GSH1**	γ-glutamylcysteine ligase (GCL)	L-Glu + L-Cys → γ-Glu-Cys	1	feedback-inhibited by glutathione; short protein half-life	**H370L** (feedback relief) · **S495A** (turnover slowed)
2	**GSH2**	glutathione synthetase (GSS)	γ-Glu-Cys + Gly → γ-Glu-Cys-Gly	1	—	downstream flux amplified

Step one is the one that matters. GSH1 forms the γ-glutamyl bond itself — the unusual side-chain linkage the receptor later reads — by condensing glutamate and cysteine at the cost of one ATP. It is the rate-limiter, and the cell governs it with exactly the kind of elegance that makes it hard to push: glutathione, the end product, loops back and inhibits the very first enzyme, so the pathway throttles itself the moment the pool is full. A cell optimized for staying alive wants precisely this restraint. A tank optimized for richness does not.

Two edits reach the two brakes. **H370L** relieves the product feedback — a single substitution at the residue through which glutathione reports back to the enzyme, so that a full pool no longer tells synthesis to stop; the line keeps building past the wild cell's comfortable setpoint. **S495A** addresses turnover, removing a regulatory serine that marks the enzyme for scheduled disposal, so the ligase persists longer and the catalytic machinery stays on the floor. Relieve the feedback and slow the demolition, and a wild strain's ceiling becomes a floor. Yeast genetics make both edits express cleanly, which is the deeper reason the program leans on this organism — not because its richness is the most complex we can grow, but because its pathway is the most *tractable* one we hold.

We carry **240 constructs** in yeast, and the reach here is less about any one flavor than about supply. An engineered yeast line is a renewable, fermentable, specifiable source of exactly the γ-glutamyl peptides the index rewards — grown in a vessel in days rather than coaxed from a field across a season. Where a garden plant is a thing you tend, a yeast line is a thing you *run*: pitch it, feed it sugar and nitrogen, hold the temperature, and it builds richness because building richness is, for this organism, the same act as staying alive.

Koji — the enzyme factory, and the γ-glutamyl hand

Where yeast contains kokumi, koji generates it. *Aspergillus oryzae* arrives at a substrate of steamed soybean or rice as a comparatively empty mold and leaves it, weeks later, as miso or the first pressing of soy sauce — a transformation run almost entirely by the enzymes it secretes. Three of those enzymes matter to the richness axis, and they work in a relay.

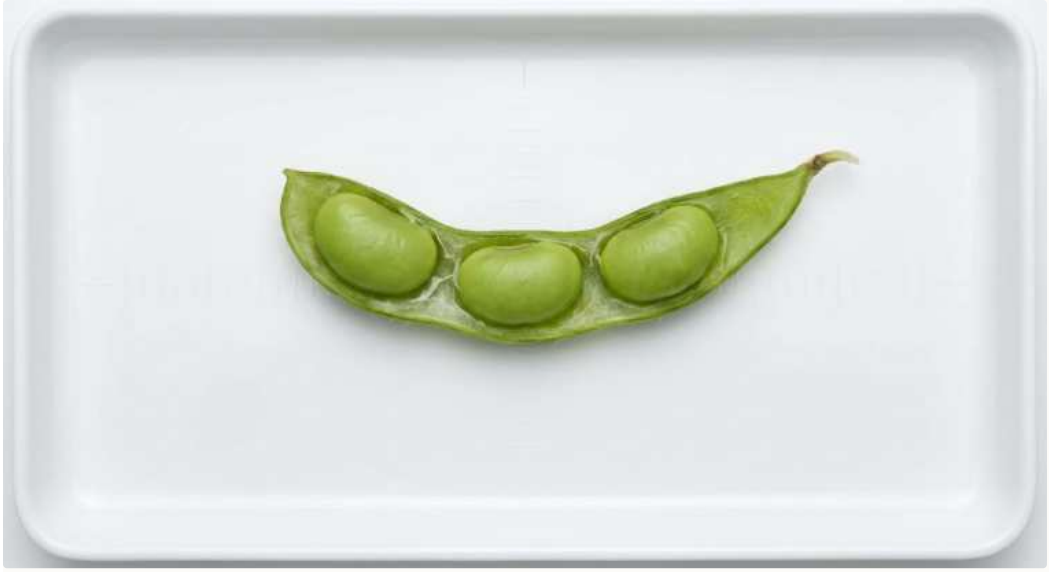

Soybean in cross-section — the protein-dense substrate koji takes apart, and the raw material of every miso the richness broth leans on.

First the **proteases**, acid and neutral, cleave the substrate's storage proteins into peptides and then into free amino acids — the slow dismantling that turns a bland bean into something savory. Among the amino acids they free is glutamate, the carrier of umami, but much of a soybean's glutamate is locked as glutamine, which tastes of almost nothing. So the second enzyme, **glutaminase**, hydrolyzes that glutamine to glutamate and converts silent nitrogen into the loud free glutamate that rings the umami receptor. The third is the one that sets the kokumi fingerprint. **γ-Glutamyltransferase** — GGT — is the only hand in the ferment that can build and move the γ-glutamyl bond, grafting γ-glutamyl groups onto small acceptors like valylglycine to assemble γ-Glu-Val-Gly, the loudest of the common richness peptides, far more potent on CaSR than glutathione itself.

Left to its defaults, a koji ferment builds a quiet, broad spectrum of γ-glutamyl peptides. Our **80 constructs** in koji are tuned mostly toward that GGT activity, so a ferment can be directed — pushed toward the loud γ-Glu-Val-Gly rather than left at its gentle average, the transpeptidase redistributing γ-glutamyl groups onto the acceptor we want instead of the ones chance hands it. That is the whole difference between a ferment that merely tastes full and one tuned to a target on the index. Yeast is the store; koji is the hand that builds and redistributes; and the γ-glutamyl bond, invisible on the tongue, is what both of them are quietly working on.

Store and hand — the roster of richness

Richness is carried across the program by a small group of hosts, and the construct counts tell you where the work has gone. The microbes anchor it, the alliums ground it in the kitchen, and a handful of fungi and a single frontier legume fill out the spectrum.

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Soybean	360	U / K	Frontier
Yeast	240	U / K	Hard
Garlic	560	K / P	Moderate
Koji mold	80	U	Hard
Shiitake	80	U	Hard
Onion	40	K	Moderate
Mushroom	40	U	Hard
Vanilla	40	K	Hard

Two things read straight off the table. First, kokumi and umami travel together — soybean and yeast both carry the U/K pairing because the same ferment that liberates glutathione liberates free glutamate; richness and savor are two outputs of one process, which is why the broth at the end of this chapter rings both receptors at once rather than choosing. Second, difficulty does not track with construct count. Garlic carries the most constructs of any richness host at 560 and is only *Moderate* to engineer, because it is a plant with a mapped genome and a cooperative pathway; yeast and koji carry fewer and are rated *Hard*, because a microbe's whole regulatory web has to be respected to push one output without starving the cell of another. The frontier legume, soybean, carries the most promise and the most risk at once — 360 constructs against a *Frontier* rating — because a seed crop that both fixes nitrogen and builds the richness peptides is a longer engineering horizon than a tank you can re-pitch next week.

Why the world's richness runs on a tank

The scale argument is where kokumi stops being a garden project and becomes an industrial reality, and it is worth stating plainly because it is the honest reason the world's savory food leans on yeast. A field of garlic is a season, an acre, a weather risk, and a harvest that varies clove to clove. A tank of engineered yeast is a few days, a defined medium, a held temperature, and a glutathione yield you can specify to the gram — and when it finishes you grow it again, identically, next week. Yeast does not make the richest kokumi we know; a long-aged miso is more complex, a roasted allium arguably deeper. Yeast makes the most *controllable* kokumi, in the most scalable form, on the shortest clock. The engineered line sharpens every one of those advantages: more glutathione per cell from the feedback-relieved ligase, a fingerprint shifted toward γ-Glu-Val-Gly, a strain stable across many generations of propagation.

Control is a number, not a feeling, and the program holds it to the same standard everywhere. The batch-to-batch coefficient of variation we require is **≤ 7%**, against the roughly **42%** a wild, unselected process drifts across. Put that on the broth's K 2,400 and the difference is stark. At CV 7%, the standard deviation is 0.07 × 2,400 = 168, and a 95% band runs 2,400 ± 1.96 × 168 — about **2,071 to 2,729**, a window 658 points wide that a formulator can build a product inside and trust. At the wild 42%, the standard deviation is 1,008 and the same band runs 2,400 ± 1,976 — **424 to 4,376**, from barely-there to overwhelming, a span no one can design against. Because both share a mean, the width ratio is just the CV ratio: 42 ÷ 7 = 6, so the uncontrolled process is literally six times vaguer. Consistency is not a bureaucratic virtue here. It is the thing that lets a number on the index mean the same thing in the thousandth tank as in the first.

The factory pays out by taking itself apart

There is an elegance in the fact that the richness is released by the organism's own undoing. **Autolysis** — self-digestion — is how yeast extract is made and how a long-aged ferment deepens: at the end of a fermentation, under heat or salt or simple exhaustion, the cell's vacuole ruptures and its own hydrolases spill into the cytoplasm,

turning the cell's enzymes against the cell's own contents. Proteases cut the proteins to peptides and amino acids; nucleases cut the RNA to 5′-nucleotides; the glutathione and the free glutamate the living cell had hoarded are released into a usable liquor. The store opens because the storekeeper is gone.



A cell held under the stage during an autolysis time-course — the moment the richness a cell built to stay alive is released by its own self-digestion.

The autolysate carries a second gift alongside the kokumi, and it is one the index records on the umami axis. Those 5′-nucleotides — chiefly 5′-GMP — are not savory on their own in any great way, but they are powerful potentiators of the glutamate receptor: free glutamate and a 5′-nucleotide together drive mGluR1 far past the sum of their separate effects, a synergy that is the quiet backbone of nearly every good broth. So a single self-digesting cell releases, in one liquor, the glutathione that rings CaSR and the glutamate-plus-nucleotide that rings the umami receptor — richness and savor, kokumi and umami, delivered together because they were built together. When we taste the deep savory note under a well-made lager, or the round backbone of a properly proofed bread, we are tasting yeast kokumi released by the yeast's own death. The richness factory pays out, in the end, by taking itself apart.

A broth base — richness you can pour

Here the two organisms come together, and here the processed-food cook recognizes their own problem: how to get the depth of the brown savory paste without the brown savory paste. This is our answer — a vegetarian, from-scratch liquid kokumi concentrate built on an engineered yeast line and a koji-aged miso. No animal stock, no added MSG, nothing on the label a cook could not pronounce.

```
KOKUMI BROTH BASE
──────────────────────────────
1.5 L          slow onion stock or water
30 g           engineered nutritional yeast (our line)
40 g           koji-aged white miso
1 tbsp         roasted garlic kokumi paste
1 sheet        kombu
2 dried        shiitake
──────────────────────────────
yield  ~1.4 L concentrate   ·   time 45 min + 30 min steep
```

1. Warm the stock to a bare steam — **not a boil**. The glutathione and the γ-glutamyl peptides keep better below a rolling boil, and the kombu turns bitter if you boil it. Drop in the kombu and shiitake and hold at 70–80 °C for **30 minutes** off the direct heat, steeping.

2. Lift out the kombu and shiitake (save the shiitake for another use). Whisk in the yeast and the garlic paste until dissolved.
3. Take the pot fully off the heat, let it settle a minute, then whisk in the miso last — miso's own kokumi and its live character are blunted by high heat, so it always goes in off the flame.
4. Taste. It should read deep, round, and long, with no single flavor stepping forward — if one does, it is not finished. Salt only at the very end, and notice how little it asks for.

IFI PROFILE

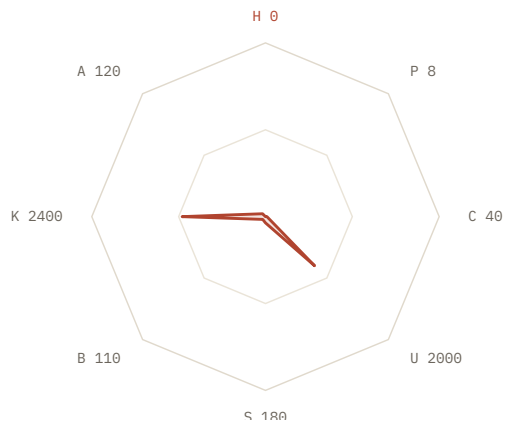

IFI-H	0
IFI-P	8
IFI-C	40
IFI-U	2,000
IFI-S	180
IFI-B	110
IFI-K	2,400
IFI-A	120

Kokumi **2,400**, umami **2,000**, and a sodium level set deliberately low because the richness is carrying the weight salt usually carries — a potentiated CaSR makes a given amount of salt read as more, so you reach for less. In receptor terms the broth is a precise object: a liquid tuned to ring the calcium receptor long and the glutamate receptor loud at the same time, which is the exact pairing a great savory base has always been reaching for by instinct.

It helps to see where 2,400 sits against a generalist. Our most complete cultivar, Phoenix (SS-0096), is built to register on all eight axes at once, and even its considerable richness lands at K 1,900:

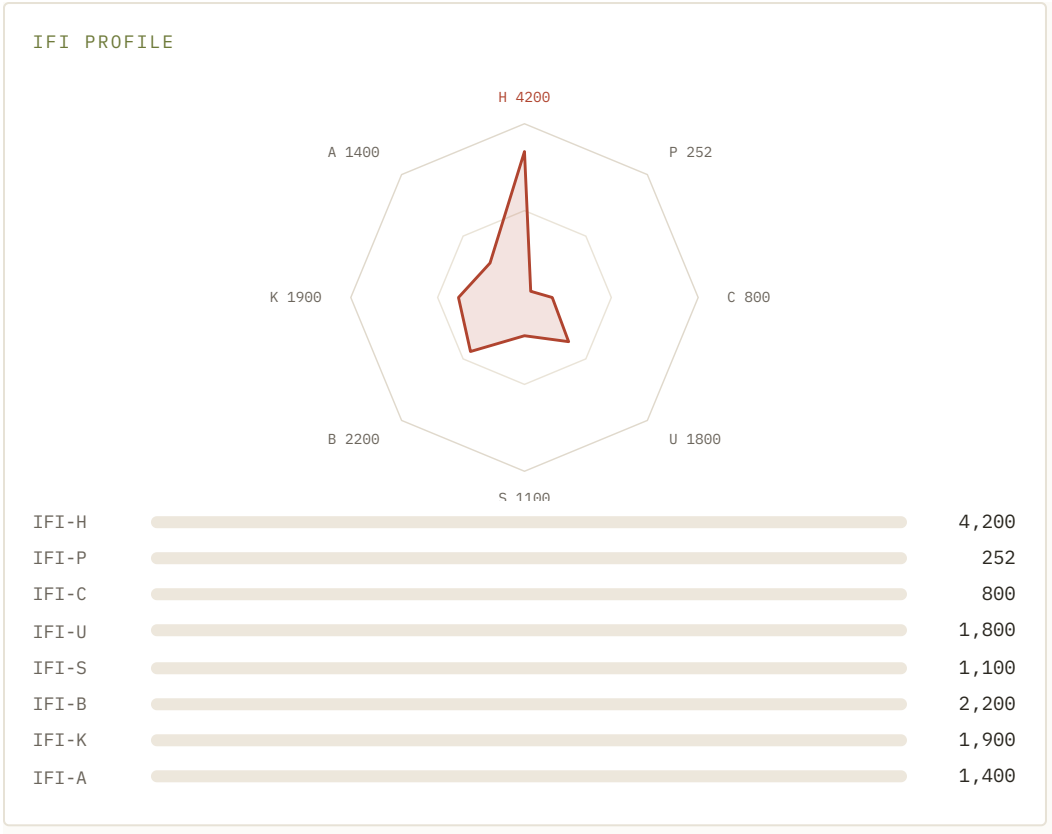

The dedicated broth out-kokumis the full-spectrum flagship on the one axis it was built for — 2,400 against 1,900, a receptor reached at 0.417 of the reference concentration against 0.526, 2.40× potency against 1.90×. That is the specialist's advantage stated in numbers: a liquor asked to do one thing can push that one axis past a cultivar asked to do all eight. The generalist is the more astonishing object; the specialist is the louder one, where it counts.

What we can build, and what we cannot yet

Honesty about the edges keeps the reverence from tipping into salesmanship. A few limits are worth naming.

The first is that CaSR is a modulator, not a source, so kokumi has a ceiling that umami does not. Push the index too high and you reach a point where the receptor is already near saturation at normal eating concentrations; a score of 2,400 sits beautifully on the steep middle of the dose-response, but doubling the peptide load again would climb the flattening top of the sigmoid and return far less sensation per gram — the Hill curve that rewards the first potentiation punishes the last. Richness is a thing you tune into a window, not a thing you maximize.

The second is heat. The glutathione and the γ-glutamyl peptides are genuinely fragile above a hard boil, and a long aggressive reduction will cost you the very molecules you built the line to make — which is why every step in the broth keeps the peptides below 80 °C and why the miso goes in off the flame. A kokumi concentrate is not a thing you can boil down like a stock without quietly undoing it.

The third is biological restraint. The feedback-relief and turnover edits push glutathione past the wild ceiling, but they push against a cell that built that feedback for a reason; drive synthesis too hard and the line pays in slower growth or lost stability over generations, trading titer today for a strain that drifts tomorrow. The ≤ 7% consistency standard is not a floor we clear easily — it is a discipline we hold against exactly that temptation, choosing a line that stays itself across a thousand tanks over a line that spikes once and wanders after. We would rather grow richness we can repeat than

richness we can only boast about. That is the whole ethic of the richness factories: make depth on purpose, in volume, from sugar and nitrogen and a microbe doing what it already wanted to do — and make it the same, every single time.

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

Chapter 10 · Soybean & the ferments — *miso, shoyu, natto.*

Raw, the soybean asks almost nothing of the tongue. Split one from a green pod and it reads faintly grassy, faintly bitter, a dense little package of storage protein and oil with scarcely any savor to give away — kokumi near the floor, umami not yet woken. And yet no seed on earth has been coaxed into more depth. Everything that makes soy *savory* is borrowed: from salt, from time, and from the quiet relay of organisms we invite to take the bean apart. Miso, soy sauce — shoyu — natto, and the broad family of fermented bean pastes that runs across all of East and Southeast Asia are, chemically, the same discovery made a hundred ways in a hundred kitchens: that if you let the right mold and the right patience work on a soybean, you get umami and kokumi together, in abundance, out of a bean that began with neither. In our program the soybean carries **360 constructs**, second only to garlic among the organisms we have tuned for richness, and it sits exactly where the index places it — at **U/K**, the meeting of the two savory axes — which is why fermented soy is the backbone of savory depth nearly everywhere it is eaten. It is not one taste. It is two receptors, rung together, and the soybean is the most fluent thing we grow at ringing both.

The chemistry of a long ferment is the chemistry of this whole book running in slow motion, and the soybean is where it runs most completely. What follows is an attempt to read that slowness exactly: the receptor that hears richness, the enzymes that write it, the arithmetic that lets us put a number on it, and the small set of edits by which we hand the ferment a bean already leaning toward depth.



Soybean, whole and split — locked protein and little flavor raw, the richest savory tradition on earth once koji and time are let loose on it.

The two savory axes, and why soy owns their meeting

Savory is not a single sensation but a chord, and the index resolves it into two notes. **Umami** — IFI-U — is the voice of free glutamate and the 5′-ribonucleotides at the metabotropic glutamate receptor **mGluR1** and the **TAS1R1·TAS1R3** heterodimer: the frank, mouth-filling savoriness of a good broth. **Kokumi** — IFI-K — is something quieter and, until recently, nameless in the West: not a taste of its own but a

modulation of the others, carried by γ-glutamyl peptides at the calcium-sensing receptor, **CaSR**. Kokumi is the thickness that arrives a beat after the first taste, the continuity that keeps a mouthful from thinning at the edges, the sense that a sauce was simmered rather than stirred together. Umami is what a broth tastes like; kokumi is why it tastes *cared for*.

A fermented soybean speaks both at once, and this is rare. Most ingredients lean one way. Pure glutamate is almost all umami; a clove of garlic is kokumi and pungency; a dried shiitake is umami with a whisper of depth. Soy is the organism that was, in effect, bred by ten thousand years of fermentation to land squarely on the U/K intersection — to give the two savory receptors their stimulus in the same spoonful, built in the same vessel by the same enzymes. Among the organisms we have engineered for the kokumi axis, soy is the second most elaborated, and it sits in distinguished company.

ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Garlic	560	K / P	Moderate
Soybean	**360**	**U / K**	**Frontier**
Yeast	240	U / K	Hard
Koji mold	80	U	Hard
Shiitake	80	U	Hard
Onion	40	K	Moderate
Vanilla	40	K	Hard

Garlic carries more constructs, but garlic pairs its richness with a sharp pungent channel; soy pairs its richness with umami, which is why soy — not garlic — is what a cook reaches for to make a thing taste *deep* rather than *bright*. And the two savory axes, read on the same scale as everything else, let us make a comparison that startles people the first time they see it. Set our most extreme pepper beside a spoonful of good miso.

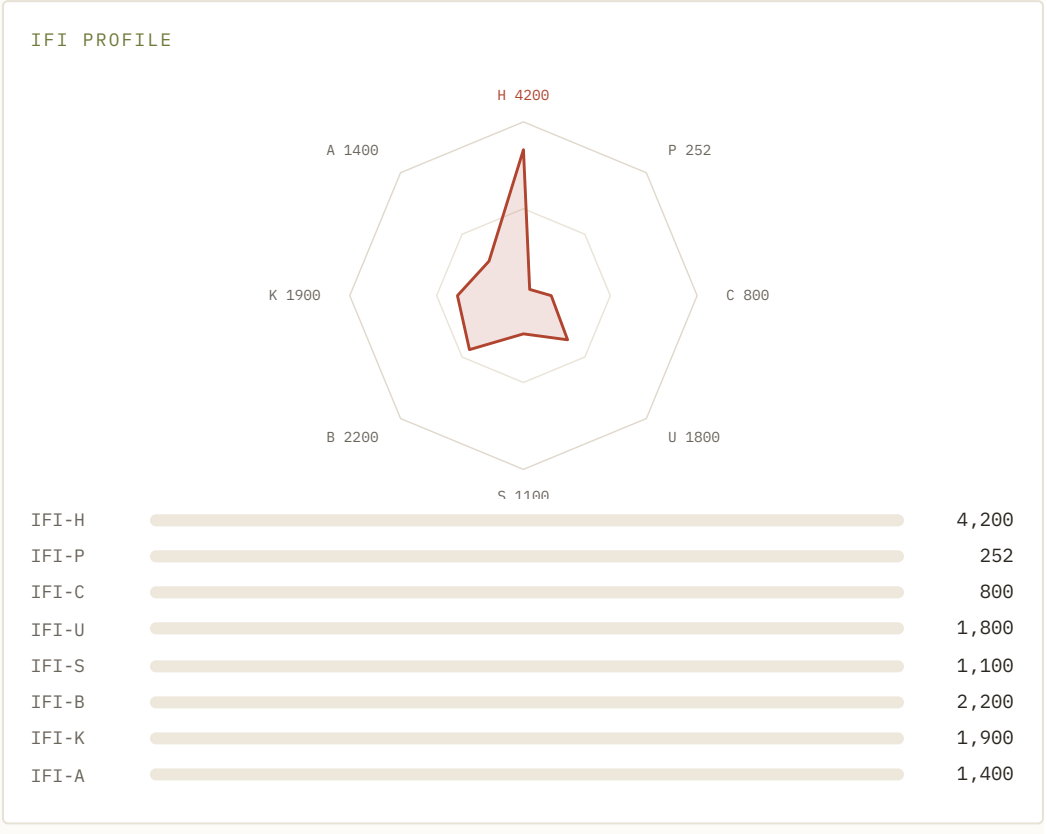

That is Phoenix (SS-0096), a predicted heat of approximately 2.95 million SHU, the first cultivar we designed to register on all eight axes at once. It reads umami **1,800** and kokumi **1,900** — formidable numbers, designed in gene by gene. Our model predicts a miso pressed from our engineered soybean reaches kokumi near **2,400**, and our model predicts the glaze at the end of this chapter reaches umami **2,200** and kokumi **2,000**: both louder, on the two savory axes, than the hottest cultivar we have designed. Soy does not try to be loud everywhere. It is simply the most eloquent organism we know on the two axes that together spell *savory*.

The receptor: CaSR, and the grammar of richness

To design kokumi you have to know what, physically, is being switched. The calcium-sensing receptor is a class C G-protein-coupled receptor — the same structural family as the metabotropic glutamate receptors and the sweet and umami taste receptors, so that three of the eight axes we measure are, at bottom, cousins reading different ligands with the same architecture. CaSR works as an obligate **homodimer**, two identical subunits clasped by a disulfide bridge, each carrying a large bilobed extracellular domain called the **Venus flytrap** for the way it closes around what it binds. Its primary physiological job, the one that named it, is to read extracellular calcium; it also answers to other polyvalent cations and, importantly for us, to L-amino acids and to a specific family of **γ-glutamyl peptides**.

Those peptides are the kokumi compounds. A γ-glutamyl peptide is one in which a glutamate residue is joined to the next amino acid not through the ordinary α-carboxyl but through the γ-carboxyl at the far end of its side chain — an unusual linkage, built by dedicated enzymes and resistant to ordinary digestion. The archetype is **glutathione**, γ-Glu-Cys-Gly, the tripeptide every living cell keeps on hand as a redox buffer; it is a modest kokumi agonist and serves well as the reference against which the others are read. More potent still is **γ-Glu-Val-Gly**, roughly an order of magnitude more active than glutathione at the receptor in the assays that first mapped the kokumi response — a peptide with essentially no taste of its own that nonetheless swells the perception of everything around it. The family runs further: **γ-Glu-Val**, **γ-Glu-Cys**, **γ-Glu-Leu**, each binding the flytrap's cleft and nudging the dimer toward its active conformation.

What the receptor does with that binding is the heart of why kokumi feels the way it does. CaSR is **positively cooperative**: because it is a dimer with coupled binding sites, occupancy at one subunit raises the affinity of the other, so its dose–response curve is not the lazy hyperbola of a simple one-site receptor but a steeper sigmoid. We write any such curve as a Hill function — the fraction of receptor driven to its active state as

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

where [A] is the concentration of the kokumi stimulus, EC_{50} the concentration that drives the receptor to half of its full response, and n the Hill coefficient that measures the cooperativity. For a non-cooperative receptor n = 1; for CaSR's coupled dimer n commonly sits near 2 to 3. That steepness matters. Take n = 2 and walk the concentration around the half-point: at [A] = $0.5 \cdot EC_{50}$ the receptor is only θ = 0.25 ÷ 1.25 = **0.20** activated; at [A] = EC_{50} it is **0.50** by definition; at [A] = $2 \cdot EC_{50}$ it has climbed to 4 ÷ 5 = **0.80**. A fourfold swing in stimulus has carried the receptor from one-fifth to four-fifths of full voice. Kokumi therefore has a threshold quality — below it a dish tastes thin and you cannot say why; cross it and the whole mouthful rounds and lengthens at once. It also means the perceived richness of a food is exquisitely sensitive to how much kokumi stimulus it actually holds near that midpoint, which is the entire argument, later in this chapter, for caring about batch-to-batch precision.

One more mechanistic note completes the picture. CaSR does not transmit a taste of its own; downstream it couples through Gq/11 to phospholipase C, liberating IP_3 and releasing calcium inside the cell, and in the taste bud that signal *amplifies* the output of the neighboring sweet, salt, and umami cells. This is why kokumi reads as enhancement rather than as flavor — why γ-Glu-Val-Gly on its own tastes of almost nothing, yet a trace of it makes a stock taste simmered for hours. The receptor is a volume knob on savor, and the soybean, fermented, is a machine for making the molecules that turn it.



Soybean cross-section — the dense 11S and 7S storage proteins that a ferment takes apart into glutamate and the γ-glutamyl library.

Reading richness as a number: the index, worked

The index exists so that a tomato, a chili, and a bowl of miso can each be read on the same axis and placed side by side. Each axis is anchored to its own reference stimulus — a stated convention, as Scoville's scale is anchored to a dilution threshold, not a single literature measurement — so that within any axis a score tells you directly how the candidate compares to that axis's reference at 1,000. The score is:

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

so that the reference stimulus scores exactly 1,000 by convention, a candidate that activates the receptor at a *lower* concentration — a louder stimulus — scores higher, and one that needs *more* scores lower. The score is deliberately a ratio, which is what lets it be read against any other food on the same receptor without ever pinning a single messy molar concentration to the page.

Rearrange the formula and it tells you directly how concentrated a food's kokumi stimulus must be. Since EC_{50}(candidate) = EC_{50}(reference) × 1000 ÷ IFI, every score on the kokumi axis is just a multiple of the reference's own half-maximal concentration:

SUBSTRATE	IFI-K	EC_{50} ÷ EC_{50}(REF) = 1000/IFI	POTENCY VS. REFERENCE
Wild-type reference	1,000	1.000×	1.00×
Conventional fermented soy	640	1.563×	0.64×
Engineered-soy miso, aged 12 mo (predicted)	2,400	0.417×	2.40×
Kokumi miso glaze, at the plate (predicted)	2,000	0.500×	2.00×

Read the middle column as *how much you need*: a conventional miso, scoring 640, must be present at 1.56 times the reference's half-maximal concentration to drive CaSR to its midpoint, because its net kokumi stimulus is a touch weaker than the reference. Our model predicts our engineered-soy miso, scoring 2,400, reaches that same half-activation at 0.417 times the reference concentration — it is the louder material, so less of it does the work. The cleanest way to state the gain is as the ratio of the two scores, which is identical to the inverse ratio of their EC_{50}s:

> 2,400 ÷ 640 = **3.75**

Our model predicts our engineered miso reaches half-maximal kokumi activation at **3.75 times lower concentration** than a conventional ferment of the same age — a 275 percent lift on the kokumi axis, carried entirely by a denser library of γ-glutamyl peptides. And because, from the Hill arithmetic above, the receptor's response near its midpoint is steep, that 3.75-fold difference in stimulus is not a subtle one on the tongue; it is the difference between a sauce that reads thin and one that reads complete.

The ferment as a three-organism relay — miso and shoyu

The flavor is written by enzymes, and in a classic soy ferment three organisms hold the pen in turn. Naming the hand-offs is what makes the kokumi legible.

First comes **koji** — the mold *Aspergillus oryzae* — grown over steamed soybean, or over a soy-and-roasted-wheat mix for shoyu, until the substrate is furred white and flooded with the mold's secreted enzymes. Koji is the engine of the whole tradition, and it carries exactly the tools the two savory axes require. Its **proteases and peptidases** — acid and neutral endoproteases, then carboxy- and aminopeptidases that trim from the ends — dismantle the soybean's dense storage proteins, the 11S glycinin and 7S β-

conglycinin, into peptides and finally into free amino acids, **glutamate** chief among them: this is the umami, released one cut at a time. Its **glutaminase** converts glutamine to still more glutamate, a step that does much of the heavy lifting in a good shoyu. Its **amylases** turn the starch of the wheat or rice koji into glucose, seeding sweetness and feeding the organisms that come next. And — the step that matters most for this chapter — its **γ-glutamyl transpeptidase** cleaves the γ-glutamyl bond of glutathione and glutamine and transfers that γ-glutamyl group onto free amino acids, stitching the **kokumi library** together residue by residue: γ-Glu-Val, γ-Glu-Val-Gly, and their kin, built in the vat by the same enzyme that recycles them in every cell.

STAGE	ORGANISM	PRINCIPAL ENZYMES	WHAT IT ADDS
1 · Koji	*Aspergillus oryzae*	proteases, peptidases, glutaminase, γ-glutamyl transpeptidase, amylases	free glutamate (U), the γ-glutamyl peptide library (K), glucose (S)
2 · Yeast	*Zygosaccharomyces rouxii*	invertase, alcohol & aldehyde dehydrogenases	ethanol, the roasted-caramel aroma HEMF, intracellular glutathione (K)
3 · Lactic bacteria	*Tetragenococcus halophilus*	lactate pathway	lactic acid — the gentle acidity that balances the whole

Then, into a brine near eighteen percent salt — the moromi — move the **halotolerant yeasts**, above all *Zygosaccharomyces rouxii*, among the few organisms that will ferment at that salinity. They turn koji's sugars to ethanol and to the aromatic esters and furanones that give an aged shoyu its roasted, faintly caramel nose, and they carry their own cellular **glutathione**, adding directly to the kokumi pool. Alongside them the salt-tolerant lactic-acid bacterium *Tetragenococcus halophilus* produces lactic acid, dropping the pH into the high fours and shaping the final balance — the acidity that keeps a long ferment from tasting flat or soapy.

Over the months that follow, three things rise together: free glutamate, the γ-glutamyl peptides, and the brown Maillard compounds — melanoidins — formed as amino acids and reducing sugars react in the slow warmth of the vat. A young white miso, fermented weeks, is bright and light and modestly rich. A dark red miso aged two or three years is a near-black, resonant depth, because the extra time is extra transpeptidation and extra Maillard chemistry both — extra γ-Glu-Val-Gly, built one slow transfer at a time, and extra melanoidin wrapped around it. **The age of a miso is, very nearly, its kokumi**, almost linearly; the aged cheeses teach the identical lesson from a different protein, that richness is simply what time does to protein in the presence of the right enzymes. Our engineering is the attempt to arrive at the far end of that time sooner, and more reliably, by handing the relay a substrate already leaning toward the loud peptides.

There is a quieter dividend in the same chemistry, and it belongs to the ethos that runs under this whole library. The proteolysis that liberates glutamate also liberates the bean's protein into free amino acids and short peptides the body reads far more easily than intact

glycinin, and the long ferment draws down the phytate that otherwise binds the bean's minerals. Fermenting soy is, among other things, a way of making the bean's nourishment more available — not a health claim, simply what the enzymes do while they are busy making the flavor. Depth and digestibility are written by the same hand.

Natto — the other path to depth

Not every great soy ferment runs through koji. **Natto** takes a different road entirely: steamed whole soybeans inoculated with *Bacillus subtilis* var. *natto*, a single bacterium doing all the work, no mold and no brine. The *Bacillus* secretes its own proteases, liberating free glutamate and peptides for umami and kokumi as koji would — but its signature is a flourish koji never makes. It polymerizes glutamate into **poly-γ-glutamic acid**, long chains of glutamate residues strung through those same γ-carboxyl linkages, and it is this polymer that gives natto its notorious sticky threads. The γ-glutamyl bond that the kokumi peptides wear as a tripeptide, natto wears as a mucilage — the same unusual chemistry at a thousand times the length, lending the ferment a body and a cling that read, on the palate, as richness carried by texture as much as by taste.

Bacillus also writes an aroma far more assertive than miso's: ammoniacal, meaty, pyrazine-laden, polarizing by design. Natto is the edge case that proves the axis — a food almost nobody is neutral about, whose kokumi is undeniable and whose aroma is a wall. It is a useful reminder that the receptor we are tuning is only one channel of several, and that a thing can score beautifully on the kokumi axis and still divide a table. We measure richness precisely; we do not pretend richness is the whole of a flavor.

What the engineered soybean changes — front-loading the substrate

Everything so far is the tradition, and the tradition's genius is untouched. What our engineering changes is not the relay but its **starting line**. A conventional soybean enters fermentation as mostly locked storage protein over a thin pool of free peptides; the koji must do all the work, and the depth a vat reaches is hostage to the bean, the season, the room, and the hand reading them. Our engineered soybean enters the vat already carrying a deeper γ-glutamyl store and a richer free-amino-acid pool, so that koji and yeast reach a target depth sooner, and — because the substrate is defined, sequenced, and the same every time — land on the same flavor every batch.

The edits are modest and native. We lean on the same **GDH1** route by which we grow umami back into a stripped tomato, raising the bean's free-glutamate pool so there is more savor at the start and more γ-glutamyl donor for the transpeptidases to spend. We deepen the glutathione arm — **γ-glutamylcysteine synthetase** feeding **glutathione synthetase** — so the bean arrives with a larger reserve of the exact precursor koji's γ-glutamyl transpeptidase elaborates into the loud kokumi peptides. And we bias the free-amino-acid acceptor pool toward valine and its neighbors, so that when that

transpeptidase goes to work it builds preferentially toward γ-Glu-Val and γ-Glu-Val-Gly rather than scattering its γ-glutamyl groups at random. None of this replaces the ferment. It meets the ferment halfway, and hands it a bean designed to end up rich.

The result, read on the axis, is the lift the worked table already named: a conventional fermented soy reads kokumi near **640**; our model predicts our engineered-soy miso, aged the same months, reaches roughly **2,400** — a 3.75-fold predicted gain, carried by that denser peptide library, arriving in months where tradition would spend years.



Sequencer bench — the defined, sequenced substrate behind a ferment that reads the same every lot.

But the number that matters most to a maker is not the mean; it is the spread. A great miso master spends decades learning to read a ferment by nose and fingertip precisely because a traditional vat *wanders* — the same recipe, the same room, a different depth each year. A processed-food formulator cannot wander. They need the same richness in lot after lot, and here the designed substrate earns its place. Across our program we hold batch-to-batch variation to a coefficient of variation of **≤7%**, against the **42%** our wild-type reference material carried. Put both figures against a kokumi target of 2,400 and the consequence is stark. At 7 percent CV the standard deviation is 0.07 × 2,400 = **168**, so about two lots in three fall within 2,232–2,568, and roughly nineteen in twenty within 2,064–2,736 — a tight, specifiable band. At the wild-type's 42 percent the standard deviation is 0.42 × 2,400 ≈ **1,008**, and that same 95 percent interval sprawls from about **384 to 4,416** — not a controlled ingredient at all but a different one each delivery.

Now recall the receptor's steepness. Because CaSR answers its stimulus along a Hill curve with n near 2 to 3, variation in the kokumi material does not blur gently at the output; near the midpoint it is amplified. A ±42 percent wobble in stimulus swings the perceived richness across most of the receptor's working range, lot to lot — exactly the inconsistency a formulator cannot design around. A ±7 percent wobble keeps every lot on nearly the same stretch of the curve, so the dish tastes the same each time it is made. The tradition gives you the ceiling; the engineering gives you a floor the tradition could never promise, and a variance the tongue cannot catch.

The honest limits

It would be a disservice to the craft to claim the bean is the whole story. Several things the engineering does not and cannot shortcut.

It does not manufacture aroma on a schedule. The roasted, caramel, faintly floral top notes of an aged shoyu — HEMF and its relatives — are the work of the yeast and of real elapsed time; a richer substrate deepens the body of a ferment, not the hours its nose needs to bloom. We raise the floor of depth and tame its variance; we do not conjure a three-year aroma in three months.

It does not forgive heat. Kokumi peptides and the aromatics around them are fragile, and a miso brought to a simmer sheds both — this is why the glaze below is loosened over the lowest heat and never boiled, and why miso stirred into a soup goes in at the very end. Scorch a miso glaze on the grill and it turns bitter, the γ-glutamyl richness you spent months building flattening into char. The failure modes are real and they are thermal.

And proteolysis has a far edge. Push a ferment too hard — too warm, too long, too little salt to restrain the proteases — and the same enzymes that liberate glutamate keep cutting into hydrophobic peptides that read as **bitter**, the flaw a debittering step or a careful salt regime exists to hold off. Richness and bitterness are cut from the same protein by the same knives; the art is knowing when to stop, and a designed substrate narrows that window but does not abolish it. The bean is honest about what it is: a better beginning, not a replacement for the patience and the judgment that finish the work.

At the table

Here is soy's kokumi at the point of use — a lacquer for roasted vegetables or a slab of tofu, where our engineered miso's depth does the work a long reduction of stock would otherwise be needed for, and does it from a bean.

```
KOKUMI MISO GLAZE
──────────────────────────────────
3 tbsp        engineered-soy red miso (aged 12 mo)
1 tbsp        mirin
1 tbsp        rice vinegar
1 tsp         roasted-garlic kokumi paste
1 tsp         toasted sesame oil
──────────────────────────────────
yield  ~90 g   ·   time 10 min
```

1. Whisk the miso, mirin, and vinegar smooth in a small cold pan.
2. Warm over the lowest heat only to loosen — **never simmer**, which flattens the miso's kokumi and its aromatics. Whisk in the garlic paste and the sesame oil off the heat.
3. Brush onto whatever you are roasting in the last few minutes, so it lacquers and lightly catches but never scorches — scorched miso turns bitter and loses the very richness you built it for.

IFI PROFILE

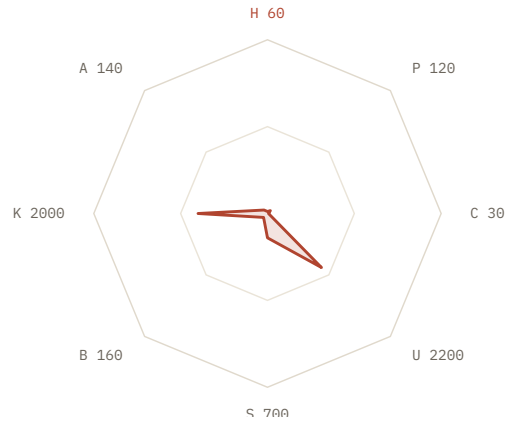

IFI-H		60
IFI-P		120
IFI-C		30
IFI-U		2,200
IFI-S		700
IFI-B		160
IFI-K		2,000
IFI-A		140

Our model predicts umami **2,200** and kokumi **2,000**, the two savory axes side by side at full voice, with just enough sweetness from the mirin to round them. This is the U/K intersection the soybean was born for, and it is the clearest single argument for why fermented soy is the most-reached-for savory ingredient on earth. Both predicted scores stand above what our most extreme pepper reaches on those same axes — and they come not from heat, nor from any animal, but from a bean that, raw, had nothing to say, taken apart and rebuilt by koji, by salt, and by time into something the receptor hears as *depth*. It is not one taste. It is two receptors, rung together — and the mathematics, for once, says exactly what the tongue already knew. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 11 · Vanilla — *the kokumi of sweetness.*

The alliums and the ferments are the savory face of kokumi; vanilla is the proof that the axis is not savory at all, in itself — that *more* can be the enhancement of sweetness and warmth just as readily as of broth. This surprises people, and it should, because vanilla is filed in everyone's mind under *sweet*. But the sweetness of vanilla is mostly borrowed. The pod carries almost no sugar, and vanillin — the headline molecule, the smell that announces *vanilla* before the spoon has reached the mouth — is not a sweet-receptor agonist at all. What vanilla actually does to a dessert is *kokumi work*: it gives a custard or an ice cream its **roundness**, its **body**, its sense of being complete and continuous rather than thinly, flatly sweet. A cheap imitation delivers the aroma and nothing under it, and the trained mouth reads the absence instantly — a single bright note with no floor, no persistence, no *stay*. Real vanilla, and good vanilla especially, carries several hundred compounds, and among them the glutathione and the γ-glutamyl peptides of the cured pod that give the flavor its mouthfulness and its long tail. Vanilla is kokumi dressed as sweetness, and it is, not coincidentally, **the most faked flavor in the world** — precisely because the thing that is hard to fake is the kokumi, not the vanillin.



The vanilla orchid and its green pod before curing — the flavor is not here yet; it is built, slowly, in the dark.

To understand how a flower with no savor and almost no sugar can round out a custard, we have to go to the receptor it speaks to — one of the strangest and most beautiful proteins on the whole tongue, and the one the rest of this book is quietly about.

The receptor that keeps time

Kokumi is read by the **calcium-sensing receptor, CaSR**, and almost nothing about that sentence is intuitive. CaSR is a class C G-protein-coupled receptor, a close structural cousin of the receptors that read umami (mGluR1, TAS1R1·TAS1R3) and sweetness (TAS1R2·TAS1R3) — three of the eight axes built from one ancient architecture. It

assembles as a dimer, two protomers clasped together, and each protomer begins with an enormous extracellular domain shaped like a **Venus flytrap**: two lobes hinged at a cleft. A ligand settles into the cleft, the lobes close, and that closure is transmitted down through a cysteine-rich stalk into the seven-pass transmembrane bundle, where it is finally handed to a G-protein on the inside of the cell. The receptor is, in the most literal sense, a hinge that counts.

What it counts, first, is calcium. The orthosteric agonist of CaSR — the molecule that occupies the mouth of the flytrap — is the free calcium ion, Ca^{2+}, and in the body at large this is the receptor that keeps blood calcium within a hair's breadth of correct. That is the wonder of the thing: the same receptor the body uses to hold its own mineral balance is the one a cured vanilla pod reaches out and touches. We say that only as biology, not as a promise — on the tongue, CaSR is doing flavor work and nothing more, and this chapter makes no claim beyond taste.

Calcium alone, though, is not what makes kokumi. The γ-glutamyl peptides — glutathione and its relatives — are **positive allosteric modulators**: they bind not at the calcium site but at a second site on the flytrap, and their effect is to stabilize the closed, active conformation. They do not, on their own, shout. They lower the concentration of calcium needed to half-open the receptor; they make the hinge easier to close and slower to spring back. A given, ordinary level of calcium at the taste cell now drives the receptor harder and for longer. The downstream story is quick: CaSR couples through Gq/11 to phospholipase C-β, which splits membrane PIP_2 into IP_3 and diacylglycerol, and IP_3 releases a pulse of calcium from the cell's internal stores. The signal propagates, and the brain reads it — not as a taste of its own, because **kokumi has no flavor you can name in isolation**, but as a change in every other taste present: the sweet rounds, the savory deepens, the salt broadens, and all of it *lasts*. Kokumi is the dwell. It is the difference between a chord struck and released and a chord struck and held.

This is why vanilla, of all things, is a kokumi flavor. The pod's glutathione and its γ-glutamyl peptides are exactly the positive allosteric modulators CaSR is built to read. They carry no sweetness to a sweet receptor and no savor to a savory one; they go straight to the hinge that keeps time, and they tell the whole dessert to hold its note a beat longer. The sugar in a custard strikes the chord. The vanilla keeps it ringing.

Putting a number on the dwell

The Index exists so that a sentence like *vanilla keeps the chord ringing* can be made exact. Every axis is anchored to the half-maximal activation of a receptor — not to how much of a compound a food contains, but to how **loudly** that compound speaks to its lock:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. A smaller EC_{50} means a louder voice — less material needed to reach the same activation — so inverting it, scaling by a thousand, and pinning it to a fixed reference turns

potency into a single score. The IFI-K axis is a stated convention: for the kokumi axis we anchor the reference to **reduced glutathione (GSH)**, the archetypal kokumi peptide, and we set its half-activation of CaSR to **EC_{50}(ref) = 2.0 mM**, so that glutathione itself reads IFI-K 1,000 by definition — just as the Scoville scale is defined by a dilution-to-threshold, not derived from a measurement. From there, every other score inverts cleanly back to a concentration:

> EC_{50}(candidate) = 1,000 × 2.0 mM ÷ IFI-K = 2,000 ÷ IFI-K (in mM, as glutathione-equivalents)

Run the canonical kokumi scores through that relation and the whole axis lays itself out as a ladder of potencies:

SOURCE	IFI-K	EC_{50} (MM, GSH-EQ.)	POTENCY VS REFERENCE
Reference · glutathione	1,000	2.000	1.0×
Vanilla kokumi custard†	1,400	1.429	1.4×
Engineered vanilla pod†	1,600	1.250	1.6×
Phoenix · SS-0096†	1,900	1.053	1.9×
Garlic	3,200	0.625	3.2×

† Our model predicts these IFI-K scores for our designed cultivars and preparations; they are computational estimates, not laboratory measurements.

Read the arithmetic in either direction and it holds. Our model predicts **IFI-K 1,600** for the engineered pod; invert it, 2,000 ÷ 1,600 = 1.25 mM, so the pod's peptide complex is designed to half-activate CaSR at four-fifths the concentration the reference needs — 1.6× louder than plain glutathione in our model. Garlic, the loud end of the axis, half-activates at 0.625 mM and reads 3,200, which is a full 3.2× the reference and roughly twice our vanilla. That ratio is honest and worth saying plainly: vanilla's kokumi is modest beside garlic's, as it should be — garlic would be grotesque in a custard — but a predicted 1.6× the reference is precisely the amount that turns a flat sweet into a round one.

The scale is anchored, but a receptor does not switch on at its EC_{50} and off below it; it follows a dose–response curve, and the shape of that curve decides how much peptide a *floor* actually costs. CaSR is cooperative — its flytrap and its dimer interface let one binding event ease the next — so its response follows a Hill relationship,

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

where θ is the fraction of maximal activation, [P] the peptide concentration, and n the Hill coefficient. Take our engineered pod, EC_{50} = 1.25 mM, with a modest cooperativity of n = 2. Half-activation (θ = 0.5) sits, by construction, at 1.25 mM. But to reach θ = 0.90 — the rich, saturating floor a great custard wants — solve $0.90 = [P]^2 \div (1.25^2 + [P]^2)$ and you get [P] = 3.75 mM, **three times** the EC_{50}. And the lower shoulder, θ = 0.10, falls at 0.417 mM. So the window from barely-there to nearly-full spans 0.417 → 3.75 mM, a **ninefold** change in concentration (for a Hill coefficient of 2, the 10-to-90 span is always 81^½ = 9). That ninefold span is the whole engineering problem of kokumi in one number: the receptor approaches its full,

round saturation slowly, so a real floor is not a trace garnish of peptide but a genuine, built-up concentration — which is exactly the thing a few cents of synthetic aroma cannot supply.

There is a subtler point hiding in the word *allosteric*, and it is the real reason a peptide that carries no taste can change one. A positive allosteric modulator does not add its own signal; it **shifts the calcium curve leftward**, lowering the apparent EC_{50} of the orthosteric agonist so that an unchanged level of calcium drives the receptor harder. Calcium sensing is steeply cooperative — a Hill coefficient near n = 3 — which makes that shift dramatic. Take an illustrative resting calcium of 1.5 mM at the taste cell and suppose the peptide moves calcium's apparent EC_{50} from 3.0 mM down to 2.0 mM. Before: $\theta = 1.5^3 \div (3.0^3 + 1.5^3) = 3.375 \div 30.375 \approx$ **0.11**. After: $\theta = 1.5^3 \div (2.0^3 + 1.5^3) = 3.375 \div 11.375 \approx$ **0.30**. The same calcium, the same cell, no new taste — and the response nearly **triples**, from 11% to 30% of maximum, purely because the hinge was made easier to close. That tripling, felt as *dwell* rather than as flavor, is what a vanilla pod delivers to a custard.

One more conversion makes the stakes tangible in the pantry rather than the assay. Expressed as mass, the reference's 2.0 mM glutathione (molar mass 307.3 g/mol) is 2.0 mmol/L × 307.3 mg/mmol ≈ **615 mg/L**. Our pod's predicted 1.25 mM equivalent is 1.25 × 307.3 ≈ **384 mg/L**. The gap, about **231 mg/L**, is the mass of peptide a maker *no longer has to reach for* to get the same dwell — the loudness we designed into the pod, counted out in milligrams.

What the pod builds in the dark

None of that flavor is present when the pod is picked. A green vanilla bean is almost odorless; it holds its future as **glucovanillin**, vanillin locked to a glucose, and it holds its kokumi only as the raw machinery to make peptides. The flavor is not grown — it is *cured*, through a long, deliberate, labor-intensive sequence (killing, sweating, slow drying, conditioning) that runs for weeks and often months. Curing is, chemically, a controlled enzymatic ferment, and it is close kin to aging a cheese or fermenting a miso: an enzyme working patiently on a substrate over time, building depth as it goes. On the aroma side, a **β-glucosidase** cleaves glucovanillin into free vanillin plus glucose — the headline note, unlocked. On the body side, and in the same slow dark, the cell's glutathione machinery keeps turning, and the γ-glutamyl peptides accumulate alongside the browning and the hundreds of minor aromatics. The pod builds both halves of itself at once, which is why a long, honest cure tastes *rounder* and not merely more fragrant.



The cured pod in cross-section — the γ-glutamyl peptides built during the long curing are the body beneath the aroma, the part the imitations cannot reach.

The γ-glutamyl bond itself is the quiet hero, and it is worth naming exactly. An ordinary peptide bond joins the α-carboxyl of one amino acid to the next; a **γ-glutamyl** bond joins through glutamate's *side-chain* γ-carboxyl instead. That unusual linkage is what CaSR's second site is shaped to read, and it is what makes these peptides resist the ordinary proteases that would otherwise clear them — so they survive to reach the tongue and ring the receptor. A small toolkit of enzymes builds and diversifies them, and it is the same toolkit in a vanilla pod, a clove of garlic, and a vat of yeast:

ENZYME	REACTION IT CATALYZES	ROLE IN THE FLOOR
Glutamate–cysteine ligase (GCL / γ-GCS)	L-Glu + L-Cys + ATP → γ-Glu-Cys	First, rate-limiting step; sets the γ-glutamyl supply
Glutathione synthetase (GSS)	γ-Glu-Cys + Gly + ATP → glutathione (γ-Glu-Cys-Gly)	Completes the tripeptide; the archetypal kokumi modulator
γ-glutamyltransferase (GGT)	Transfers the γ-glutamyl moiety onto acceptors → γ-Glu-Val-Gly, γ-Glu-Val, and kin	Diversifies the peptide set; builds the more potent modulators
β-glucosidase	Glucovanillin → vanillin + glucose	Frees the headline aroma during the cure — the *other* half

The first step, catalyzed by glutamate–cysteine ligase, is rate-limiting: it governs how much γ-glutamylcysteine is ever available, and therefore how high the whole peptide floor can rise. That single bottleneck is the most important sentence in this chapter for what we do next, because a rate-limiting step is precisely the kind of thing you can engineer.

Forty constructs, two pathways

The vanilla orchid carries, in our program, only **40 constructs**, and it is one of the hardest organisms we work — slow-growing, fussy, its flavor developed not on the vine but through that long, enzymatic cure. The family has spent patient seasons on this plant, and the engineering targets both of the pod's halves on purpose. The **vanillin pathway** is tuned for the aroma; the **γ-glutamyl pathway** is

tuned for the body, with the rate-limiting ligase relieved so that the floor the plant builds in curing rises reliably rather than by luck. The point is not to out-aroma a synthetic — that is trivial, and anyone can do it with a few cents of guaiacol. The point is to make the hard half, the kokumi, **more deliberate than nature usually bothers to be**, so the cured pod delivers not only more of the headline note but a deeper, dependable floor beneath it. On the kokumi axis, our model predicts our engineered vanilla at roughly **1,600** — modest against garlic's 3,200, but designed to be transformative in a dessert, where a little kokumi is the whole difference between flat and round.

The pod on its own, before a grain of sugar is added, our model profiles as profiling almost entirely on one axis:

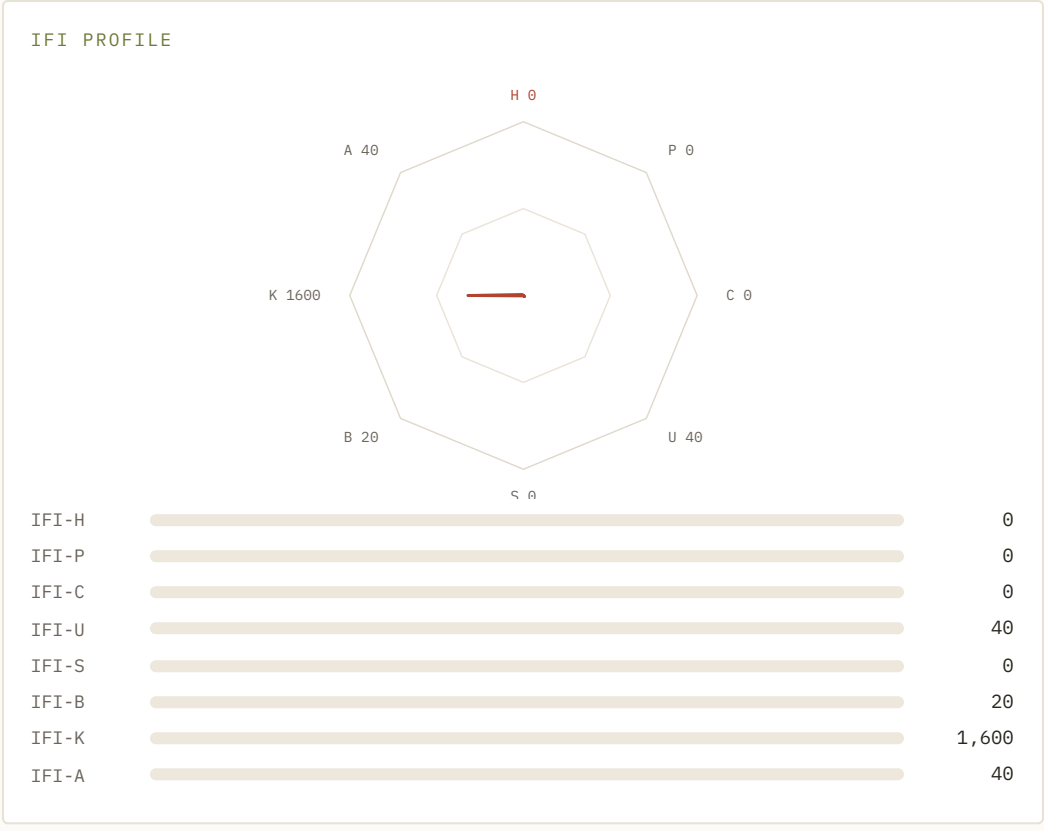

Read that honestly and it is nearly a blank card with one tall bar. **S: 0** is the whole surprise of vanilla written in a number — the pod reads *zero* on the sweet receptor, because its famous sweetness is cross-modal, an association the brain makes from aroma, and never a true TAS1R2·R3 signal. The trace of umami and the whisper of bitterness are real but tiny. Everything vanilla *is designed to be*, as a receptor event, sits in that single **K: 1,600** — the dwell, the body, the floor. The pod is a kokumi delivery device wearing the costume of a sweet.

And because kokumi is read by one receptor with one architecture, the same method reaches well past this one difficult orchid. CaSR does not care which plant grew the peptide; a γ-glutamyl modulator from a yeast rings it exactly as a modulator from a pod does. The canonical kokumi-bearing organisms in the program share the same γ-glutamyl toolkit and differ mainly in how much of it they will tolerate being pushed:

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

Vanilla sits at the small, hard end — 40 constructs in a plant that fights for every one — which is exactly why it is the organism that proves the point. If we can design a dependable kokumi floor into something this slow and this fussy, the receptor's generosity does the rest.



An isolated cured specimen — the whole of vanilla's body lives in the peptides this pod spent months accumulating.

The economics of the fake

The economics of faking are worth one clear-eyed passage, because they explain the entire vanilla market and, by contrast, exactly what we are doing. **Vanillin** — the single headline aroma compound, molar mass 152.15 g/mol — is cheap and easy to make synthetically, from wood-pulp lignin or ferulic acid or petrochemical guaiacol, and a few cents of it can give a product the *smell* of vanilla. What synthetic vanillin cannot give is the body: the glutathione and the γ-glutamyl peptides, the hundreds of minor aromatics, the kokumi floor that a cured pod spends months building. So the imitation hits the top note and leaves the dessert hollow underneath, and the trained mouth knows instantly, even when the nose is fooled. That is the precise sense in which vanilla is the most-faked

flavor in the world: the fakery is possible *only* because the headline compound and the kokumi floor can be separated, and the cheap half sold without the expensive half.

Our engineering is the opposite move. We design *both* halves into the pod, and we build the hard half more deliberately than nature usually does, so that a real vanilla delivers its roundness and its continuity every time — at a cost and a consistency that let a maker use the true thing where they would otherwise have reached for the hollow imitation. We are, in a plain commercial sense, defending the part of vanilla that cannot be counterfeited by making it **cheaper to grow than to fake**. The receptor can tell the difference; we are simply making it affordable to stay on the right side of what the receptor knows.

Consistency, and the honest limits

Reliability is not a footnote to kokumi — it is the product. A flavor you cannot count on is a flavor a formulator will not design around, and the real vanilla pod's historic problem has never been its ceiling but its *variance*. A single stated number carries that story. Wild-type cured vanilla runs a batch-to-batch coefficient of variation near **42%**; our constructs are designed to hold **≤7%**. Propagate those through the pod's predicted IFI-K of 1,600 and the consequence is stark:

- Engineered, CV ≤7% (design target): standard deviation ≈ 0.07 × 1,600 = **±112**, a band of roughly **1,488 – 1,712**.
- Wild-type, CV 42%: standard deviation ≈ 0.42 × 1,600 = **±672**, a band of roughly **928 – 2,272**.

The wild-type pod can land anywhere from a thin, disappointing ~930 to a muddy, overblown ~2,270 — the batch-to-batch lottery that teaches a maker to distrust the real pod and reach for the predictable fake. The engineered pod is designed to hold a narrow, designable band, and that narrowness — not a higher ceiling — is what lets someone build a recipe on it. Carry the same 7% down into a finished custard at our predicted IFI-K 1,400 and our model projects the dessert's floor at **1,400 ± 98**, a swing of under a tenth, small enough that day one and day fifty taste like the same intention.

Honesty requires the limits, too. Kokumi is a *modifier*, not a standalone taste, so its score is always read in company: the same 1,600 does more in a custard with real sugar to round than in a watery one with nothing to hold, because the receptor's job is to prolong what is already there. The low steep that extracts the peptides is fragile — a hard boil scalds the kokumi and drives off the top aromatics, collapsing both halves at once — so the method matters as much as the pod. And the receptor's reach beyond the tongue, its role in the body's own calcium balance, is real biology and a source of genuine wonder, but it is *not* a dietary effect we invoke or claim; at the concentrations a dessert carries, these peptides are doing flavor and only flavor. We say what the number means for taste, and we stop there.

The custard

This is the dish where the whole premise of kokumi becomes most intimate, so the family voice wants to say it plainly: a custard with real, deep vanilla is one of the most *comforting* things a person can eat, and the comfort is not the sugar. A flat sweet custard is not comforting — it is cloying. The comfort is the roundness, the body, the way the flavor fills the mouth and *stays*, which is the kokumi, which is the calcium receptor ringing long. The thing that makes a dessert feel as if it were made *for* you, rather than merely sweetened, is a quiet peptide speaking to a receptor that has nothing to do with sugar. That is worth building on purpose, and it is worth cooking gently.

```
VANILLA KOKUMI CUSTARD
──────────────────────────────
500 ml        whole milk or barista oat milk
1 whole       engineered vanilla pod, split and scraped
4             egg yolks (or 60 g aquafaba-starch base)
60 g          cane sugar
pinch         flaky salt
──────────────────────────────
yield  4 small cups   ·   time 25 min + chill
```

1. Warm the milk with the scraped pod and seeds and the pinch of salt to a bare steam, then take it off the heat, cover, and steep **20 minutes** — the low steep is where the kokumi extracts. The γ-glutamyl peptides are water-soluble and heat-labile at the top; a hard boil would scald that floor and drive off the headline aromatics in the same breath, costing you both halves. Gentle and slow is not a nicety here; it is the chemistry.
2. Whisk the yolks with the sugar until pale. Temper in the warm milk slowly, whisking, then return to the lowest heat and stir constantly until it coats the back of a spoon at about **82 °C** — no higher, or the yolk proteins seize and the texture breaks. That 82 °C is the custard's own narrow band, the thermal twin of the ≤7% we designed into the pod: a few degrees past it and the dish is a different dish.
3. Strain into cups and chill at least 4 hours. The flavor deepens overnight as the kokumi integrates and the peptides settle into the matrix; it is better on day two than on day one — the continuity axis made visible, the dwell extended from a single mouthful to a whole second day.

IFI PROFILE

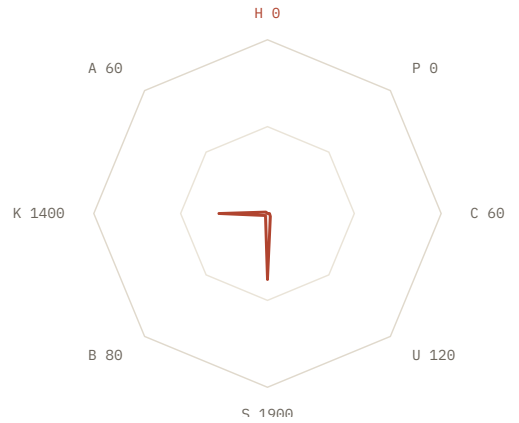

IFI-H	0
IFI-P	0
IFI-C	60
IFI-U	120
IFI-S	1,900
IFI-B	80
IFI-K	1,400
IFI-A	60

Sweetness reads **1,900** — but look at the kokumi at **1,400** sitting under it, and understand that the custard would taste thin and forgettable with the first number and not the second. Run the ratio and it is exactly our ladder: our model predicts an effective EC_{50} of 2,000 ÷ 1,400 ≈ 1.43 mM, a kokumi floor designed to be 1.4× louder than the reference, holding the 1,900 of sweetness in its arms so the sugar never falls flat. The salt, the barest pinch, is there to do the same job from the other side: a trace of sodium sharpens the perceived sweet and, through calcium-adjacent chemistry at the taste cell, nudges the kokumi receptor toward its active hinge as well. This is a dessert built on two receptors you were never taught to name, carrying a flavor everyone thinks they understand and almost nobody does — and that quiet lengthening of a sweet into something round and whole is, in the end, the clearest case we know that richness is a kind of care. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 12 · The naturally rich — *where kokumi gathers on its own.*

Before we leave the organisms, an honest accounting of the places kokumi accumulates *without* any help from us — because our whole method is the attempt to do deliberately, and fast, what nature does slowly and by accident, and you cannot respect the first without understanding the second. Kokumi is, at bottom, what time and the right enzymes do to protein. Everywhere in the food world that those two conditions meet — a protein-rich substrate and patient enzymatic action over weeks or months or years — kokumi gathers. The traditions that prize those foods are, without ever naming the receptor, traditions of the calcium-sensing receptor. The cheesemaker turning a wheel in a cave, the family burying a crock of soybeans to draw a sauce from it a year later, the grandmother who lets the onions go translucent and then amber over a low flame for the better part of an afternoon — every one of them is a γ-glutamyl chemist, and not one of them ever needed to know it.

We write this chapter in reverence, because it is the chapter where our engineering is smallest against what it is standing next to. The two-year Parmigiano is a monument. We did not build it, and we would not want to replace it. What we want is to understand it so completely — the receptor it speaks to, the bond it builds, the exact shape of the curve it climbs — that we can hand the same gift to a person who does not own a cave or a year. That understanding begins, as all of this does, at a receptor that tastes of nothing at all.

The receptor beneath the tradition

The axis we are measuring is **IFI-K**, kokumi, and the organ reading it is the **calcium-sensing receptor, CaSR** — a class C G-protein-coupled receptor, cousin to the sweet and umami receptors, that spends most of its working life nowhere near the tongue. Its day job is calcium homeostasis: in the parathyroid gland it reads the free calcium of the blood and tells the body when to release the hormone that pulls calcium back into balance. It is, first and last, a calcium meter. The tongue borrowed it. On a subset of taste cells CaSR sits reading the same ion, and the richness we are chasing is what happens when that meter is nudged.

Picture it honestly, because you cannot design for a receptor you cannot picture. CaSR works as an obligate **homodimer** — two full copies clasped together — and each copy carries, held off the cell surface on a stalk, a large bilobed extracellular domain shaped like a clam: the **Venus flytrap**. Two lobes, a hinge, a cleft between them. The receptor's true agonist, the thing that closes the trap, is the bare **Ca^{2+} ion**. But the trap has a second, gentler input. Small charged molecules that look like free amino acids can settle into the cleft and act not as the trigger but as *positive allosteric modulators* — they do not open the channel themselves; they tune the receptor so that it answers the calcium already around it more loudly. The **γ-glutamyl peptides** of all those slow foods work exactly here. A peptide docks

its free glutamyl head in the amino-acid site, the rest of the chain trailing out of the cleft, and the receptor, sensitized, reports more signal for the same ambient calcium. Downstream the activated dimer couples through Gq/11 to phospholipase C, which releases a pulse of calcium *inside* the cell — and that internal flash is what a receptor panel actually measures.

This is why kokumi is the invisible axis. CaSR activation does not send the brain a flavor of its own; it sends an *enhancement* — a deepening of continuity, thickness, and mouthfulness, and an amplification of the sweet, salty, and savory notes arriving alongside it. The peptides make the meal louder without adding a note of their own. And it is why a food's kokumi can never be fully separated from its minerals: because the receptor is a calcium sensor, the calcium a food carries sets the baseline the peptides push against.

That interplay is worth one worked example, because it is the molecular definition of richness arriving. Take CaSR's response to calcium as a cooperative curve — a dimer counting votes between its two flytraps gives a Hill slope near 2 — with a half-activating calcium concentration of about **3.5 mM** in the unmodulated receptor. Saliva carries free calcium at roughly **1.2 mM**, which falls short of that midpoint, so at rest the receptor is quiet. Now add a γ-glutamyl modulator that left-shifts the calcium midpoint to **1.75 mM** — a twofold sensitization, well within what a potent kokumi peptide does. The fraction of maximal response follows $f = [Ca]^n \div (EC_{50}^n + [Ca]^n)$:

RECEPTOR STATE	Ca^{2+} EC_{50}	SALIVA Ca^{2+}	CASR RESPONSE *
Unmodulated	3.50 mM	1.2 mM	$1.2^2 \div (3.50^2 + 1.2^2) =$ **10.5%**
+ γ-glutamyl modulator	1.75 mM	1.2 mM	$1.2^2 \div (1.75^2 + 1.2^2) =$ **32.0%**

** fraction of maximal activation, Hill slope 2, at unchanged saliva calcium.*

The peptide moved nothing in the mouth but the receptor's willingness. The same 1.2 mM of calcium that drove CaSR to one-tenth of its voice now drives it to a third — a **threefold lift in richness with not one extra ion of calcium added.** That is the whole sensation of a rich broth "filling out" on the tongue, written as arithmetic: a modulator sliding a curve to the left under a fixed amount of the receptor's own ligand.

The enzymology of patience

Now the chemistry of how those modulators get made, because the natural traditions and our engineering run the identical reaction — only the clock differs. Two things happen in parallel whenever a protein-rich food is left to age or ferment, and they are worth separating cleanly.

The first is **proteolysis**: residual proteases, native to the food or supplied by its microbes, cut the long proteins into fragments and free amino acids. This is where the **free glutamate** comes from — the umami, the savory loudness of old cheese, the crystalline crunch in a long-aged Gouda that is partly glutamate crystal. Proteolysis alone would give you umami.

The second is **transpeptidation**, and this is kokumi's own reaction. An enzyme of the **γ-glutamyl transferase** family — γ-glutamyl transpeptidase, GGT — takes the γ-glutamyl group from a donor (glutathione, or glutamine) and transfers it onto an acceptor amino acid or small peptide. The product is a **γ-glutamyl peptide**: a chain joined not through glutamate's ordinary backbone carboxyl but through the second, γ-carboxyl at the tip of its side chain. That single displaced bond is the whole secret of the family. It leaves glutamate's own head — its free α-amino group and α-carboxyl — sticking out of the molecule in exactly the arrangement the receptor reads as *amino acid present*, and it does something else just as important: it makes the molecule nearly **immortal in the mouth**. Ordinary peptidases, the enzymes of digestion, evolved to cut α-bonds; they cannot touch a γ-bond. Only another γ-glutamyl transferase can open it. So the peptide the ferment builds is not dismantled on the tongue in seconds the way an ordinary flavor peptide would be — it survives, keeping the receptor ringing. **The bond that makes kokumi loud is the same bond that makes it slow to leave, and the same bond that lets it accumulate over an aging at all.**

Garlic is the clearest teacher of this chemistry in the plant kingdom, which is why we reach for it again and again. The clove holds the densest natural store of γ-glutamyl precursors a plant makes — γ-Glu-Cys, γ-Glu-Val, and the allium-specific **γ-Glu-*S*-allyl-cysteine** that is the backbone of aged-garlic preparations. Roast a head of garlic until it goes soft and sweet and you are running a slow transpeptidation in the oven: the harsh raw clove becomes round and deep precisely as those peptides build and the glutathione pool turns over.



The clove in cross-section — garlic holds the densest natural store of γ-glutamyl peptides in the plant kingdom, and roasting it slow is a transpeptidation you can smell.

Here is the tension every aged food lives inside, and it is the tension our engineering was built to resolve. The transpeptidation is **enzyme-limited at low temperature**. The proteases and transferases in a cheese cave or a miso cellar work at a chill far under their optimum — deliberately, because a food held warm enough to ripen quickly is a food warm enough to spoil before it ripens. Richness and safety pull against each other, and tradition's only resolution is patience: months or years of slow, cold enzymatic work, during

which the peptide population drifts toward its stable, accumulating, immortal members. The cook is trading time for safety, and getting depth as the reward for the trade.

The arithmetic of aging

That drift has a shape, and the shape is the clock our engineering is racing, so it is worth writing down exactly. The accumulation of γ-glutamyl peptides in an aging food is, to a good approximation, a **first-order approach to a ceiling** — the classic saturating exponential:

```
C(t) = C∞ · (1 − e^(−t / τ))
```

where C(*t*) is the kokumi-active peptide pool at aging time *t*, C∞ the ceiling the food would reach given infinite time, and τ the time constant set by how fast the cold enzymes work. The rate of accumulation is dC/d*t* = (C∞ / τ) · e^(−*t* / τ): **fastest at the very beginning and decaying from there.** Fast gains early, then a long diminishing climb toward the depth that marks a great aged thing.

Fit it to the one anchor we have. A 24-month Parmigiano reads **IFI-K 1,900**, and because the score of a whole food scales with how much active peptide it carries, we can treat 1,900 as the 24-month point on the curve. Take a time constant of τ ≈ 8 months — of the right order for a hard alpine cheese held cold — and the ceiling falls out: 1,900 = C∞(1 − e^(−24/8)) = C∞(1 − e^(−3)) = C∞(0.950), so C∞ ≈ **2,000**. Now read the climb:

AGING TIME	1 − E^(−T/T), T = 8 MO	IFI-K REACHED	GAIN OVER THE 6-MO WINDOW
6 months	0.528	1,055	**+1,055**
12 months	0.777	1,554	+499
18 months	0.895	1,789	+235
24 months	0.950	1,900	**+111**

The two ends tell the whole economics of aged food. The **first six months** build the kokumi from nothing to 1,055 — more than half the final depth. The **last six months**, from 18 to 24, buy only 111 points. That is a ratio of 1,055 ÷ 111 ≈ **9.5**: nearly ten times the flavor per month in the first half-year as in the last. The final stretch of a two-year cheese buys a fraction of the depth the opening stretch did and costs exactly as much time, exactly as much rent on the cave, exactly as much risk. This is why great aged food is expensive in a way that feels almost unfair — you are paying linearly in months for a reward that is flattening out underneath you. The economics of the cave are the economics of that curve.

Our engineering does not speed the enzyme. It does something the cave cannot: it **moves the starting point up the curve.** By handing the process a substrate already enriched in γ-glutamyl precursors, and an organism whose transpeptidase is already biased toward the loud, valine-bearing peptides the receptor most wants to hold, we begin the aging not at C = 0 but partway to the ceiling. Write that as a head start C_start and the curve becomes C(*t*) = C∞ − (C∞ − C_start) · e^(−*t*/τ). Suppose the enriched substrate and the tuned

enzyme hand us, at $t = 0$, the depth a cold start would not reach until month six — an IFI-K of 1,055. Then to reach the twelve-month depth of 1,554:

> `2,000 − (2,000 − 1,055) · e^(−t/8) = 1,554 → 945 · e^(−t/8) = 446 → t = 8 · ln(945 / 446) ≈` **`6.0 months.`**

The depth that took a year from a cold start arrives in half the time. Run the same enriched start for the full twelve months instead and it reaches C(12) = 2,000 − 945·e^(−1.5) ≈ **1,789** — the eighteen-month depth, in a year. The same place sooner, or a deeper place in the same time. We are not cheating the clock, and we are not racing the enzyme. **We are starting it later in the day.**

The ladder tradition climbed

It is humbling and clarifying at once to set the natural champions on the same scale as our constructs, because these foods reached — by time and tradition alone — numbers our engineering has to work to match. The IFI-K column is measured on the same calcium-reporter assay as every molecule in our panel; the two columns beside it are not new measurements but arithmetic, the index formula run backward to show what each score *means* in concentration. We anchor the axis, as we anchor it everywhere, to **glutathione** — the body's own γ-glutamyl tripeptide, the natural reference — with its CaSR half-activation set by convention to the index's 1,000, which corresponds to an EC_{50} of **5,000 µM (5 mM).** Invert the formula, IFI-K = 1,000 × 5,000 ÷ EC_{50}, and every food's score becomes an **effective glutathione-equivalent EC_{50}** — the concentration of plain glutathione that would ring the receptor exactly as loudly as the food does.

NATURAL FOOD	IFI-K	EFFECTIVE EC_{50} (GSH-EQUIV)	MG/L GSH-EQUIV	PRINCIPAL KOKUMI SOURCE
Aged Parmigiano · 24 mo	1,900	2.63 mM	809	γ-Glu peptides + free glutamate
Fish sauce · 12+ mo	1,700	2.94 mM	904	γ-Glu-Val-Gly
Aged Gouda · 12 mo	1,400	3.57 mM	1,097	γ-Glu-Val-Gly, γ-Glu-Val
Dry-aged beef · 35 day	1,300	3.85 mM	1,182	γ-Glu peptides (proteolysis)
Dark miso · 24 mo	1,200	4.17 mM	1,281	γ-Glu peptides + glutamate
Scallop · raw	900	5.56 mM	1,707	glycine + γ-Glu peptides
Roasted garlic	800	6.25 mM	1,921	γ-Glu-*S*-allyl-Cys / GSH
Ripe tomato · cooked	500	10.0 mM	3,073	γ-Glu peptides + glutamate
Fresh onion	180	27.8 mM	8,537	γ-Glu peptides

Read the method once, on the top line, so the rest of the table reads itself. To score the Parmigiano we take a standardized extract of the cheese, run it in a dilution series across CaSR-expressing cells, and find the dilution at which it half-activates the receptor; that dilution turns out to match **2.63 mM** of plain glutathione. Then IFI-K = 1,000

× 5,000 µM ÷ 2,632 µM = **1,900**. The glutathione-equivalent mass is pure chemistry you can check: 2.63 mM × 307.3 g/mol (the molar mass of glutathione) ≈ **809 mg/L**. That is the headline the whole table carries — a 24-month Parmigiano rings the kokumi receptor as hard as a solution holding more than three-quarters of a gram of pure glutathione per litre, and it built that store one slow transpeptidation at a time.



The research tray under lab light — the aged, the fermented, and the cured all run the same slow enzymatic chemistry we compress into a growing season.

The spread down the ladder is the spread of patience and protein. **Aged cheese is the cathedral of natural kokumi** — a hard cheese past a year carries both the free glutamate of proteolysis and the γ-glutamyl peptides of transpeptidation, the umami and the richness built in parallel by the same cold enzymes, which is why grating a two-year Parmigiano over a plain dish transforms it out of all proportion to the quantity. **Fish sauce and the fermented-seafood traditions** are the aged cheese of the savory-ferment world, and we write them here with full respect, as we write all the animal foods in this Library — not as lesser, not as apology, but as another of humanity's great kokumi discoveries. A long-fermented fish sauce, anchovies and salt and a year or more, is one of the richest sources of **γ-Glu-Val-Gly**, the loudest member of the family, on earth, which is exactly why a few drops finish a broth or a dressing far beyond their volume. **The scallop** earns its own line: naturally high in glycine and the kokumi-adjacent compounds, which is part of why a good one has that sweet, round, lingering body that makes it feel like more than its size. **Dry-aged beef** develops its prized depth through precisely the proteolysis and transpeptidation we have been describing — aging is kokumi, in meat as in cheese. We are a plant-first, lifetime-vegetarian family, and we hold these animal traditions in genuine reverence, because they are the clearest natural teachers of the axis: the cook who understands why aged cheese and fish sauce and dry-aged beef are rich understands kokumi completely.



The soybean — reservoir of the long ferment, where microbial transpeptidase builds γ-Glu-Val-Gly molecule by molecule across the months a miso or a soy sauce takes to become itself.

And then the plant ferments and the slow kitchen, which is where our own hands stay. **Dark miso** past two years is a soybean paste that has run the same transpeptidation under koji and salt; its 1,200 is a whole cuisine's worth of depth drawn from a legume. **Roasted garlic** at 800 and the **nine-hour onion** — the amber one, cooked down across a long afternoon until its 180 as a raw bulb has climbed far up the ladder — are the home cook's version of the cave, aging compressed into one low flame. **Ripe cooked tomato** at 500 carries γ-glutamyl peptides and glutamate together, the reason a long-simmered sauce reads as more than fresh fruit. Every one of these is an enzyme working slowly on a protein, building γ-glutamyl peptides a little at a time, and the richness is the accumulated record of that patience.

Starting the clock later in the day

Set the natural champions beside our engineered ladder and the comparison is not a contest but a meeting. Our designs predict **garlic at an IFI-K of 3,200** and **yeast at 2,800** — above the two-year Parmigiano, but on the same axis, read by the same receptor, built from the same bond. Converted through the formula, the predicted garlic IFI-K of 3,200 gives an effective EC_{50} of 5,000,000 ÷ 3,200 = **1,563 µM (1.56 mM)**, and yeast's 2,800 gives **1,786 µM (1.79 mM)** — both louder than glutathione's own 5 mM reference, because the engineering has loaded the organism with the valine-bearing peptides the receptor holds best, rather than letting it stop at the quieter glutathione a cold ferment would leave it at. The two-year Parmigiano at 1,900 is not beaten by our garlic. It is **joined** by it, a little higher on the number and a great deal faster to the table.

These are the hosts we build kokumi into, named with the construct counts and engineering difficulty the Index assigns them:

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY	THE KOKUMI ROUTE
Garlic	560	K / P	Moderate	native γ-glutamyl pool — γ-Glu-Cys, γ-Glu-Val, γ-Glu-*S*-allyl-Cys
Soybean	360	U / K	Frontier	ferment-borne γ-Glu-Val-Gly, umami and kokumi together
Yeast	240	U / K	Hard	glutathione-rich; transpeptidation to the loud tripeptide
Onion	40	K	Moderate	allium γ-glutamyl precursors, amplified
Vanilla	40	K	Hard	kokumi mouthfeel beneath the vanillin

What the engineering buys, beyond speed, is the thing the cave could never promise: **a number that holds.** The discipline we hold every construct to is a batch-to-batch coefficient of variation of **≤7%** — against the **42%** that raw, phenotype-and-hope biology hands you. That gap is the real product, and it changes what a score on this chapter's ladder even means. The Index is a ratio of two measured EC_{50}s, and the fractional errors of a ratio add in quadrature: $CV(IFI) = \sqrt{(CV_{ref}^2 + CV_{candidate}^2)}$. With the glutathione reference re-run on every plate to a tight ~3% and a construct held to 7%, the index is good to $\sqrt{(0.03^2 + 0.07^2)} \approx$ **7.6%**. So our model predicts our engineered garlic at **3,200 ± 240** — a computed range of 2,960 to 3,440, harvest after harvest. Now put the same error bar on tradition. A wheel of "1,900" Parmigiano, at biology's native 42% spread, is honestly **1,900 ± 798** — anywhere from 1,100 to 2,700 depending on the wheel, the cave, the season, the particular slowness of one year's enzymes. The cheese is sublime and it is a **wide distribution**; the construct is a little louder and it is a **promise.** That is not a knock on the cheese. It is the difference between a gift that arrives when it arrives and a gift you can put in a formulation and count on.

A full-spectrum cultivar carries this kokumi as one designed axis among eight — the fruit speaking richness at nearly twice the glutathione reference while it also carries heat, sweetness, and the rest:

IFI PROFILE

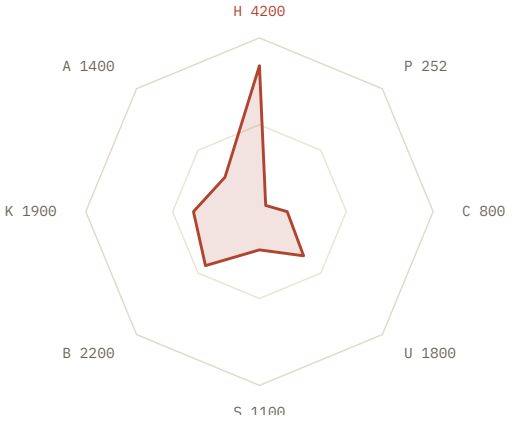

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

That **K of 1,900** is our model's prediction for the whole engineered γ-glutamyl pool, computed on the same scale as the two-year cheese it happens to equal — the same number, reached in a growing season instead of two years in a cave, and held to 7% instead of 42%. Against that democracy of eight axes, the single-purpose kokumi line is its opposite: a culinary garlic designed to pour nearly everything into one receptor, its richness driven to the ladder's top while only its native allium pungency and a whisper of savor ride along:

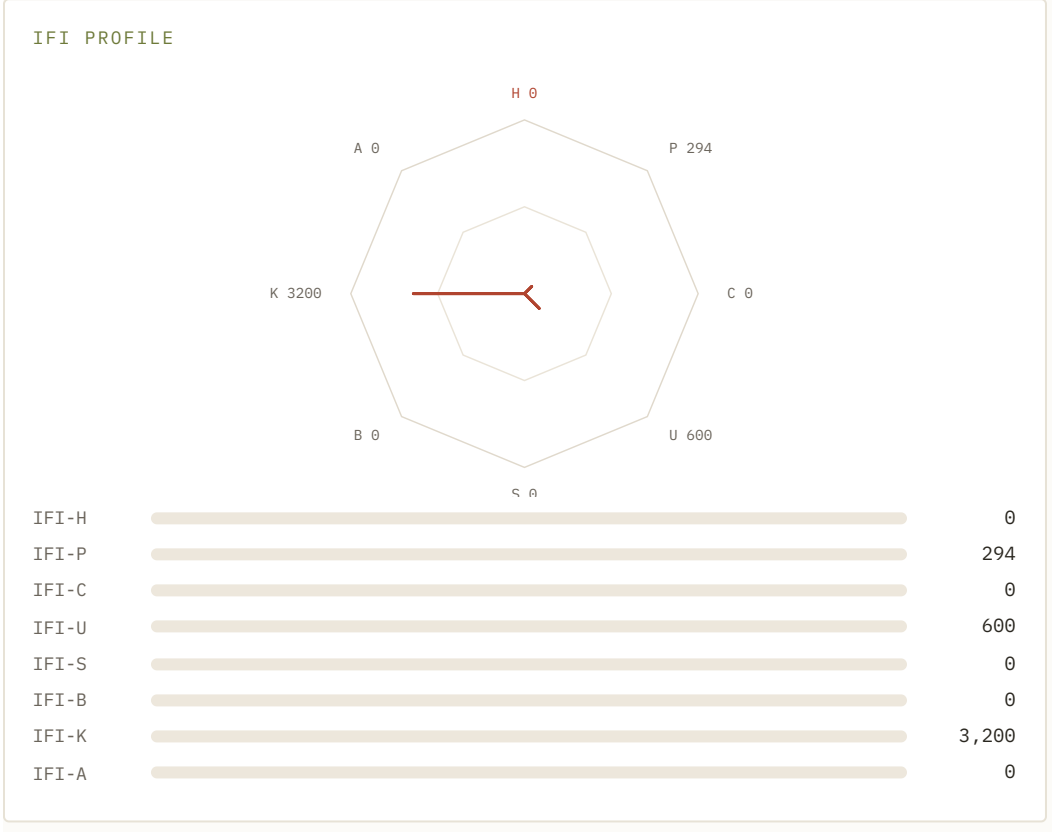

The honest edges

A number earns trust only when you name what it cannot say, and kokumi's edges are specific. The first and most important: **kokumi is modulatory, not a taste.** At the concentrations that matter, the γ-glutamyl peptides register nothing on their own — dissolve engineered garlic's active peptides in plain water and you will taste almost nothing, because there is nothing for them to enhance. The axis only does its work *on a substrate*: a broth, a cheese, a sauce already carrying sweetness, salt, and umami for the peptides to deepen and lengthen. A kokumi score is therefore a promise about what a food will do to the *rest* of a dish, never a flavor you can taste in isolation. This is the honest reason the axis stayed invisible to food science for so long.

The second edge is the one the receptor's own biology writes: because CaSR is a calcium meter, **the kokumi of a food is never fully separable from its minerals, or from the water it is cooked in.** The γ-glutamyl peptides, glutathione most of all, are mild calcium chelators, and the ion they modulate the receptor toward is the receptor's own agonist. A stock built on hard water and a stock built on soft water do not land in quite the same place on this axis, and a calcium-rich food — aged cheese again — gets a second helping of richness for free, because it brings both the modulator and the ion. We report a kokumi score as a score *for a given salt load and a given water*, and we do not pretend the peptide acts in a vacuum.

The third edge is the gap between a **compound** and a **food.** The loudest single peptide, γ-Glu-Val-Gly, has a CaSR EC_{50} near **390 µM** — a pure-compound potency of IFI 1,000 × 5,000 ÷ 390 ≈ 12,800, far past the 5,000 ceiling our scorecards normalize to. Yet no natural food on the ladder exceeds 1,900, because a food holds that loud peptide only in **traces**, far down the shallow foot of its dose-response curve. The food-level EC_{50}s in our table are high — millimolar — not because the peptides are quiet but because they are dilute. The whole act of engineering is to raise a trace to a dose: to move an organism up the curve until the loud peptide it already makes by accident is present in enough quantity to clear the knee and be heard. We are not inventing a molecule. We are making enough of the right old one.

And the final edge is the humility the ladder itself teaches. The Index measures **loudness at a receptor, not pleasure at a table.** A 1,900 tells you how hard a food rings CaSR; it does not tell you that a two-year Parmigiano is one of the great foods of the world, or that a few drops of fish sauce can save a dinner, or that an onion cooked slow enough to go amber is worth the afternoon. Those are facts the number points at and cannot hold. The best use of a measured axis is not to rank the cheese against the garlic — they are on the same axis, joined, not in competition — but to let a person who has neither the cave nor the year reach the richness anyway.

Because that, in the end, is the whole of it. The aged cheese, the long miso, the fermented fish, the cured vanilla, the nine-hour onion — every one is an enzyme working slowly on a protein, building γ-glutamyl peptides a little at a time, and the richness is the accumulated record of that patience. You cannot rush a two-year Gouda and get a two-year Gouda. What you *can* do — what we do — is hand the process an organism that starts further along and runs truer, so that the end of the road arrives sooner and lands in the same place every single time. We are not trying to abolish the cave or the fermentation crock; we revere them too much, and they taught us everything on this page. We are trying to make their gift reliable and available — so that the depth our grandmother coaxed out of a long afternoon's stock is something a person without a long afternoon can still give their body. That, too, is a kind of care, and it is the care this whole Library was built to carry. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.



PART IV OF V

Kokumi in the Mouth

the umami marriage · the amplifier · the kitchen of continuity

Chapter 13 · Kokumi × umami — *the most-requested combination.*

Of all the pairings in the Index, the one our savory customers ask for first, most, and most urgently is **umami and kokumi together** — IFI-U × IFI-K, the two savory axes rung at once. It is the taste behind the great broths and the long sauces, behind the whole family of food the eater can only call *deep*, and it is the single most powerful tool we own for making a dish satisfying without reaching for salt, or fat, or meat. For a family that has cooked without meat across three generations, it is also the most personal thing in the Library — the proof that a pot of liquid grown entirely from plants can sit a body down and nourish it as fully as any stock ever simmered from bones. The pairing earns a chapter of its own because the synergy between the two axes is not additive. It is not that you get umami *plus* kokumi. It is genuinely multiplicative, and the reason is written in three places at once: in the receptors, in the integral, and in the way the two signals come to share a single tail.



Backlit greenhouse at golden hour — the savory plants mature at dusk, when the day's warmth has driven the free glutamate up into the fruit.

Two cousins in the same cleft

To understand why the two axes multiply, you have to see that the tongue reads them with the same kind of machine. The kokumi receptor, CaSR — the calcium-sensing receptor — and one of the umami receptors, mGluR1, are **cousins**. Both belong to the small, elegant group biologists call the class C G-protein-coupled receptors: the same family that holds the metabotropic glutamate receptors, the GABA-B receptor, and the sweet and umami T1R taste receptors. A class C receptor is not shaped like the compact seven-helix bundles that read adrenaline or light. It carries, hung out into

the saliva on a long tether, a large bilobed extracellular domain that closes around its agonist exactly the way a Venus flytrap closes around a fly — two lobes, a cleft between them, and a hinge. Biologists call it, without irony, the Venus-flytrap domain.

The mechanism is worth seeing slowly, because every number later in this chapter rests on it. The agonist drifts into the open cleft and touches both lobes. Its presence tips a conformational equilibrium: the two lobes swing shut, trapping the molecule, and the closed flytrap is the receptor's active shape. But a single flytrap closing is not yet a signal. Class C receptors work only as **dimers** — two full receptors clasped face to face — and it is the *pair* of flytraps, reorienting against each other as they close, that wrenches the whole assembly from its resting arrangement into its active one. That twist travels down through a cysteine-rich collar to the seven-helix domains buried in the membrane, which rock into an asymmetric active pose, and only then does the G-protein on the inside of the cell feel it and fire. Glutamate speaking to mGluR1 is one flytrap pair snapping shut over savor. The γ-glutamyl peptides speaking to CaSR are another flytrap pair snapping shut over richness. Two cousins, the same architecture, reading two different messages in two different clefts an atom's breadth apart on the surface of the same taste cell.

When CaSR fires, it does not ring once and fall silent. Through its G-protein it drives phospholipase C, which liberates inositol trisphosphate, which throws open the cell's internal calcium stores — and that release is characteristically **oscillatory**, a train of calcium transients rather than a single spike. This is the physical fact under everything the word *kokumi* is reaching for: not a loud note but a sustained, slowly-adapting ringing that holds the taste cell lit long after the first contact. Umami, by contrast, is a cleaner, faster event — a rise and a decay. Kokumi is the drone beneath it. Hold that picture: a quick bright signal and a long ringing one, side by side on the same tongue. The multiplication happens in the space between them.

Why glutamate speaks to both receptors

Here is the molecular coincidence that makes the whole pairing inevitable, and it is the single most important sentence in the chapter: **glutamate, the umami molecule, is also a positive co-agonist at the kokumi receptor.** CaSR is a famously promiscuous flytrap. Its primary agonist is the calcium ion itself, but its cleft is also sensitized by a whole class of L-amino acids — the aromatics most strongly, but glutamate among them — which bind at or beside the amino-acid site and make the flytrap readier to close around calcium. They do not fire CaSR alone; they lower the price of admission. And the dedicated kokumi ligands — glutathione, which is itself γ-glutamyl-cysteinyl-glycine, and the small γ-glutamyl peptides like γ-Glu-Val-Gly that cluster in garlic, onion, aged cheese, yeast extract, and every long ferment — sit in that same receptive region and push the receptor the rest of the way open.

So consider what happens when a broth carries both. The glutamate you loaded for umami is, at the same instant, sensitizing the kokumi receptor beside it — helping CaSR's flytrap close, raising the

calcium ringing. And the kokumi state it helps create, that sustained oscillation, feeds back to enhance the perceived intensity and duration of the umami. The two receptors are not independent judges scoring the same dish from separate benches. They are **potentiating each other**, each lowering the other's threshold, which is why a stock with both reads not as *savory and also rich* but as one unified, profound depth — a single taste that is plainly more than the sum of its two parts. A cook feels this as a dish becoming *whole*. It is, underneath, two cousin receptors agreeing.

This is also why the famous old kitchen synergy of glutamate and the dried-mushroom or dried-fish nucleotides belongs to the same story. The 5′-ribonucleotides — the GMP that floods out of a dried shiitake, the IMP of katsuobushi — bind a site adjacent to glutamate on the umami flytrap and jam it shut, so that a little glutamate suddenly tastes like a great deal. It is the same trick in a different cleft: a second molecule that does not shout on its own but makes the first one roar. The mouth was always a machine for compounding. We did not invent the multiplication; we only learned to measure it, and then to grow the molecules that drive it.

The Index, and the arithmetic of loudness

The Index exists so that none of this has to stay poetry. Each axis is anchored to the half-maximal activation of its receptor — the concentration at which a stimulus drives the flytrap to half of full response, its EC_{50} — inverted against a fixed reference constant and scaled by a thousand:

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

Read it as a loudness meter. A low EC_{50} means the receptor reaches half-response at a low concentration — the food speaks *loudly* — and because the candidate's EC_{50} sits in the denominator, loudness and IFI move together. The anchor is set so that **IFI = 1,000 is the reference itself** by convention: a candidate that half-activates at exactly the reference concentration scores a flat thousand. Everything above a thousand is louder than the reference; everything below is quieter. Because each axis is anchored to its own defined reference constant, IFI scores on different axes are not directly comparable to each other — each one measures loudness relative to that axis's own convention, not on a single unified scale.

Nowhere is this cleaner than in the tomato, the organism at the center of this chapter and the cornerstone of our savory kitchen. A vine-ripe tomato is one of the richest free-glutamate stores in the plant world — but the commercial tomato most people have tasted is picked hard and green, gassed, and trucked, and arrives having lost the overwhelming majority of that glutamate to the road. We grew it back: not by dosing the fruit but by retuning the plant's own nitrogen chemistry at the native node where glutamate is made and spent — nudging glutamate dehydrogenase, GDH1, to pool free glutamate, and easing off glutamate decarboxylase, GAD, which otherwise drains that pool away into GABA. In our model, the result is an axis predicted to lift from **IFI-U 300 to IFI-U 2,600.**

Put the formula to work on those two numbers. Solve each for the candidate's EC_{50} in units of the reference:

- Field tomato, IFI-U 300: EC_{50} = 1,000 × EC_{50}(ref) ÷ 300 = **3.33 × EC_{50}(ref)**. It takes more than three times the reference concentration of glutamate to half-fill the receptor — a quiet, short savor.
- Engineered tomato, IFI-U 2,600: EC_{50} = 1,000 × EC_{50}(ref) ÷ 2,600 = **0.385 × EC_{50}(ref)**. It half-fills the receptor at thirty-eight percent of the reference — loud.

The ratio of the two is the honest measure of what the breeding returned: 2,600 ÷ 300 = **8.67×**. The engineered fruit rings the umami receptor at one-ninth the glutamate concentration the trucked fruit needs. That is the road's theft, reversed, read in the only language the tongue respects.

SOURCE	IFI-U	EC_{50} (× REFERENCE)	LOUDNESS VS. REFERENCE
Reference anchor	1,000	1.00	1.00×
Field / trucked tomato	300	3.33	0.30×
Engineered tomato	2,600	0.385	2.60×
The Mother Stock	2,600	0.385	2.60×

It helps, once, to carry the Index back into the concentrations a chemist would weigh. Free glutamic acid has a molar mass of 147.13 g/mol. A ripe tomato holds on the order of 200 mg of free glutamate per 100 g of flesh; that is 2.0 g per litre of juice, and

> $2{,}000\ \text{mg·L}^{-1} \div 147.13\ \text{g·mol}^{-1} \approx \mathbf{13.6\ mmol·L^{-1} = 13.6\ mM}$

of glutamate sitting in the cell water, ready to speak. The Index does not replace that number — it contextualizes it, telling you not how much glutamate the fruit holds but how loudly, through a receptor, that glutamate is heard. Both facts matter. The clean-living instinct cares about the milligrams; the eater cares about the loudness; the design cares about both, which is the whole reason we measure twice.

Why the synergy is multiplicative, not additive

Now the heart of it. Umami, alone, has a time course — a rise to a peak and a decay down a tail, its own curve of perceived intensity I(t). What the eater actually feels, the total savor of the mouthful, is not the peak alone but the whole area the curve encloses:

> $S = \int I(t)\, dt \approx I_{peak} \times \tau$

— the height of the peak multiplied, to a first approximation, by how long it lasts. Hold that product in mind, because kokumi reaches into both of its factors at once.

That same co-agonism between the two clefts raises the **peak**: with the kokumi receptor sensitizing the savory cell, the same glutamate is read as louder, so I_peak climbs — call the lift a factor of 1.30. And the long oscillatory ringing of CaSR extends the **tail**: the savory percept is held open as long as the calcium keeps ringing, so the

effective duration τ stretches — call it a factor of 2.00. Both the height and the width of the umami curve grow, and since the felt total is height times width, the two gains do not add. They multiply.

WHAT IS PRESENT	PEAK FACTOR	DURATION FACTOR	FELT TOTAL (AREA)	CHANGE
Umami alone	1.00	1.00	1.00	baseline
Umami, peak lifted	1.30	1.00	1.30	+30%
Umami, tail extended	1.00	2.00	2.00	+100%
Umami × kokumi	**1.30**	**2.00**	**2.60**	**+160%**

A thirty-percent-taller peak held for twice as long is not a 30% improvement, and not a 100% one, but a **160%** one — 1.30 × 2.00 = 2.60. The additive intuition would have predicted only (0.30) + (1.00) = 130%. The gap between them is the synergy, and it has a precise name. Expand the product:

$$a \cdot b = 1 + (a - 1) + (b - 1) + (a - 1)(b - 1)$$

With a = 1.30 and b = 2.00, the first two correction terms are the 30% and the 100% an accountant would tally; the last term, (a − 1)(b − 1) = 0.30 × 1.00 = **0.30**, is the cross-term — the depth a cook can feel on the tongue and never quite explain on paper. The whole mystery of *why a stock with both axes is so much better than a stock with one* lives in that single product of two differences. It is not magic. It is the interaction term, and we can print it.

The same multiplication can be read straight off the dose-response curve, which is the most rigorous way to see it. A receptor's response follows the Hill form

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

where n is the Hill coefficient, the steepness that cooperative, dimeric class C receptors tend to carry above 1. A co-agonist's entire effect is to **lower EC_{50}** — to left-shift the curve. Suppose the kokumi environment halves the umami receptor's EC_{50}. Since IFI runs as $1/EC_{50}$, halving the EC_{50} **doubles the IFI**: the umami score rises simply because the kokumi peptides are in the bowl. And the reverse holds in the other cleft — the glutamate left-shifts CaSR, so the kokumi score rises because the umami is present. Each axis multiplies the other's number. That is the mechanism of multiplicativity stated in one line: *co-presence lowers both EC_{50}s, and the Index, being built on the reciprocal of EC_{50}, turns every left-shift into a product.*

Cooperativity sharpens the whole effect. Evaluate the Hill response at the original EC_{50} concentration after the co-agonist has halved EC_{50}. For a non-cooperative receptor (n = 1) the response climbs from half-maximal to R = 1 ÷ (0.5 + 1) = **0.67 R_max**. For a cooperative dimer (n = 2) it climbs to R = 1 ÷ (0.25 + 1) = **0.80 R_max**. The flytrap pair — the defining fact of a class C receptor — is exactly what makes these savory axes answer a co-agonist so steeply. The architecture that lets two cousins read two messages is the same architecture that lets each amplify the other.

The organisms that carry both axes

If the two axes want to be rung together, the kitchen's job is to supply both molecules — the glutamate and the γ-glutamyl peptides — and the Library is organized to make that easy. Some of our organisms are pure umami engines; some are pure kokumi; a precious few carry both from a single genome.

ORGANISM	CONSTRUCTS	PRIMARY AXES	CARRIES
Tomato	2,040	S / U / F / M	umami (+ a modest kokumi store)
Garlic	560	K / P	kokumi
Soybean	360	U / K	**both**
Tea	280	U / B / A	umami
Yeast	240	U / K	**both**
Koji mold	80	U	umami
Shiitake	80	U	umami
Mushroom	40	U	umami
Onion	40	K	kokumi
Vanilla	40	K	kokumi

Read the column on the right and the design becomes obvious. **Soybean** (360 constructs, U/K) and **yeast** (240 constructs, U/K) are the dual carriers — the two organisms that speak both savory axes at once, which is precisely why miso, soy sauce, tempeh, and every yeast-raised ferment have anchored savory cooking on every continent that grew a bean or caught a wild culture. They are the U×K intersection folded into a living thing. Garlic, onion, and vanilla bring the kokumi half alone — garlic especially, whose 560 constructs make it the richest γ-glutamyl-peptide source we grow. Koji, shiitake, mushroom, and tea bring glutamate and, in the dried mushrooms, the nucleotides that jam the umami flytrap shut.



Soybean — 360 constructs across U and K, the one plant in the Library that speaks both savory axes from a single genome.

And then there is the tomato — our highest-construct savory organism at **2,040 constructs**, tuned across the sweet, umami, full-spectrum, and multi-axis dimensions with a kokumi store folded in — a cornerstone of the savory kitchen despite being, botanically, a fruit. A ripe tomato carries free glutamate for umami and a modest reserve of kokumi-active peptides, and a *cooked-down* tomato concentrates both while the browning chemistry builds more. That is why tomato paste reduced dark is one of the most efficient savory-depth ingredients a cook owns: a small jar of U×K. The same logic

runs under every great savory tradition on earth — the dashi of kombu glutamate and aged katsuobushi; the Italian soffritto under a long ragù; the slow Indian base of onion and garlic and tomato cooked until it leaves the oil. Every one of them is a cook building the U×K intersection by instinct, loading glutamate and γ-glutamyl peptides into a base so that everything after lands deep.



Tomato — 2,040 constructs; IFI-U designed to lift 300 → 2,600 at the native GDH1 / GAD node. A ripe fruit is already a jar of umami with a kokumi store beside it.

The pairing is so fundamental that even our most extreme plant carries it. Phoenix, the first cultivar designed to register on all eight axes at once — a superhot whose predicted heat reaches the top of the scale — still holds **umami 1,800 and kokumi 1,900** quietly in its profile, the two savory axes riding together inside a chili the way they ride together in a broth.

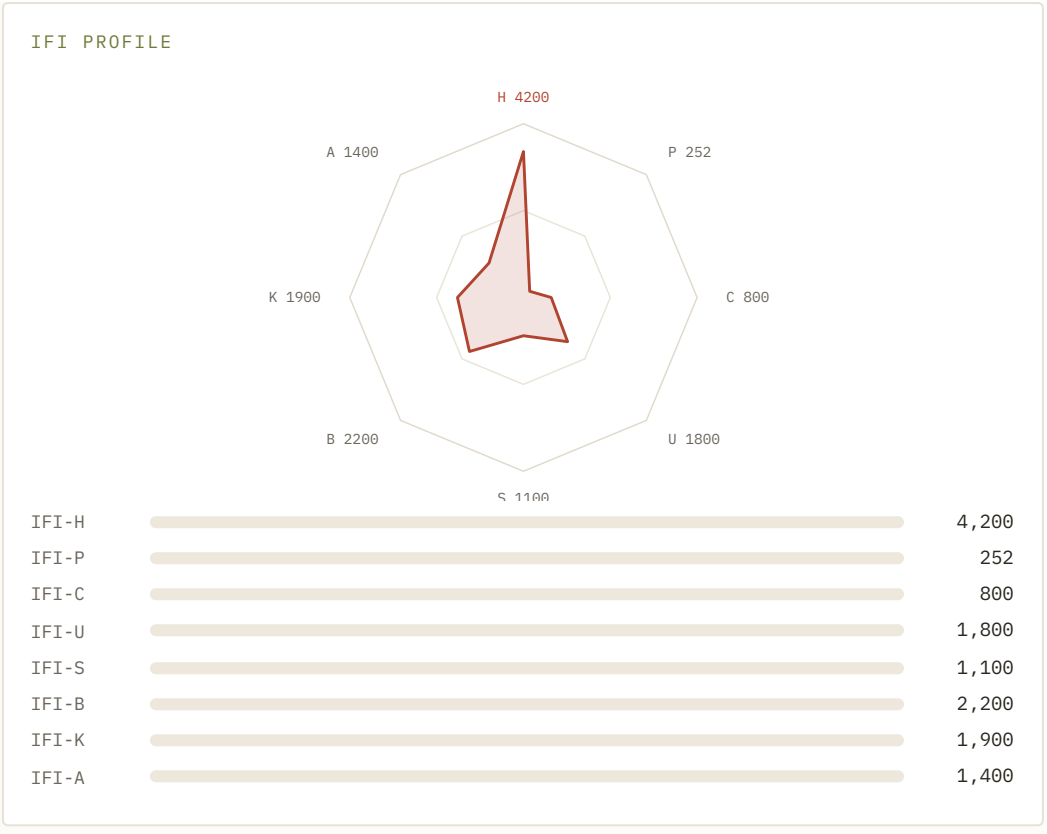

Notice that both U and K are present and strong — 1,800 and 1,900 respectively against their own references — even here. When we design a plant to be complete, the two savory cousins arrive hand in hand, because the tongue was always going to read them together.

Building the intersection — the Mother Stock

The keystone of the whole Library's savory cooking is a single recipe: a mother stock that is nothing but the U×K intersection, built from the engineered organisms of this book, entirely plant-based,

designed to replace the *function* of a meat stock not by imitating it but by hitting the same two receptors it hits — harder, and cleaner.

```
THE MOTHER STOCK — U × K
──────────────────────────────
2 L             kokumi broth base, ours
3 tbsp          engineered tomato paste, cooked dark
2 sheets        kombu
20 g            dried shiitake
2 tbsp          roasted-garlic kokumi paste, ours
1 tbsp          engineered-soy red miso, ours
──────────────────────────────
yield  ~1.8 L   ·   time 1 h + overnight rest
```

1. In a dry pot, cook the tomato paste over medium heat, stirring, for **5 minutes** until it darkens a shade and smells roasted rather than raw. This is the browning step — the one place in the recipe you are deliberately running the Maillard reaction, condensing the paste's amino acids and sugars into new savory and kokumi-active species while you drive off the raw edge.
2. Add the broth base, kombu, and shiitake. Bring to a bare steam — **never a boil** — and hold at **75–80 °C for 40 minutes**. This is the extraction: the kombu surrenders its glutamate, the shiitake its glutamate and its GMP, into the water, below the temperature that would turn either bitter.
3. Lift out the kombu and shiitake. Whisk in the roasted-garlic paste for its γ-glutamyl peptides. Take the pot fully off the heat, let it cool a minute, and whisk in the miso last, so the living complexity of the ferment is never cooked.
4. Rest it overnight in the cold. This is not optional. The U×K integration deepens measurably as it sits — slow proteolysis and diffusion let the glutamate and the peptides equilibrate through the whole volume — so the stock is markedly rounder on day two. Reheat gently, never boiling, to use.

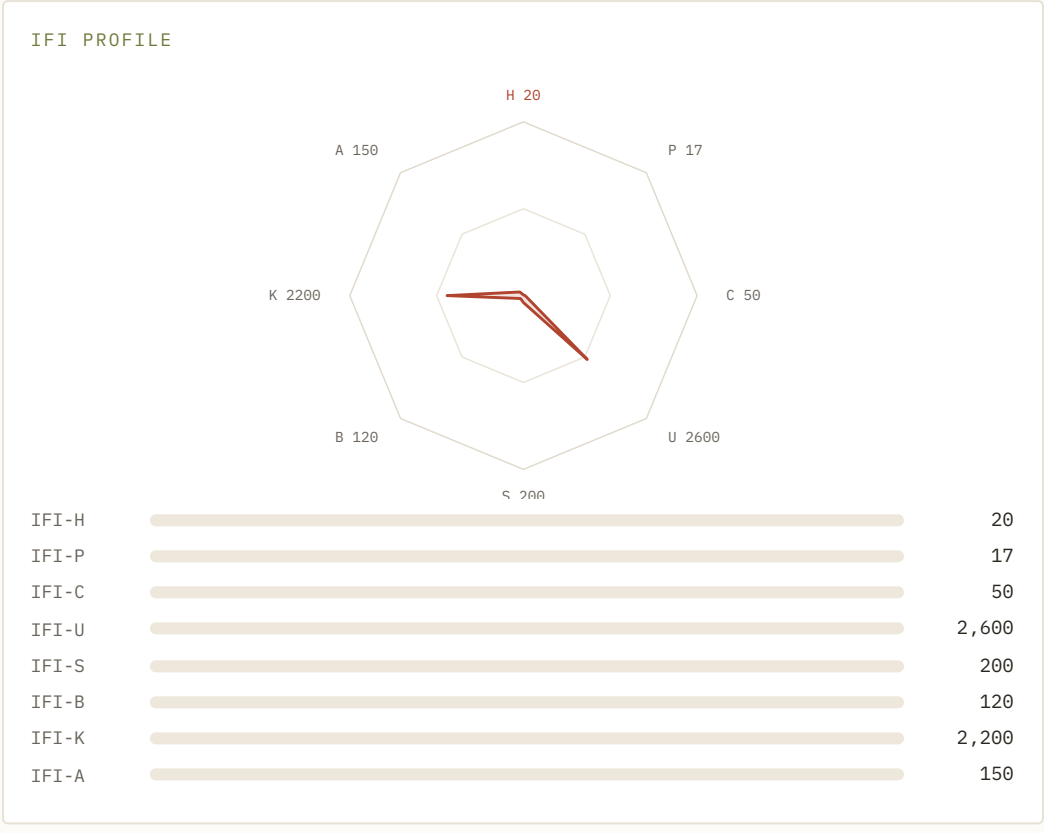

Umami **2,600** and kokumi **2,200**, side by side near the top of their respective axes, with every other axis deliberately quiet — this is the single deepest vegetarian liquid in the Library, and it is deep for a precise, legible reason. It rings both class C savory receptors at

once, each potentiating the other, the ringing tail of one dragging out the fast decay of the other, so the whole savor is both taller and longer than either axis could reach alone. Use it as the base of any soup, braise, risotto, or sauce and the dish will carry a satisfaction the eater reads as *meaty*, or *long-cooked*, or simply *right*.

Error, edge cases, and the honest limits

Numbers this clean are only worth printing if we are honest about their error bars, and about where the model breaks.

Precision. Every construct we ship holds a batch-to-batch coefficient of variation at or under **7%**, against the 42% that wild-type biology drifts by. That tolerance is not housekeeping; it is what makes that multiplication *measurable*. Error propagates through a product in quadrature, so the felt total S = I_peak × τ carries a relative error of

> $CV(S) = \sqrt{(CV(I_{peak})^2 + CV(\tau)^2)} = \sqrt{(0.07^2 + 0.07^2)} \approx \mathbf{9.9\%}$

— tight enough that the 30% cross-term, the synergy itself, stands clear of the noise. Run the same arithmetic on wild-type variability and you get $\sqrt{(0.42^2 + 0.42^2)} \approx$ **59%**, a cloud of uncertainty that swallows the interaction term whole. You cannot see a 30% synergy through a 59% error bar. The consistency *is* the chapter: without it, the multiplication is folklore; with it, the multiplication is a specification. The IFI ratio itself propagates the same way — $CV(IFI) = \sqrt{(CV(ref)^2 + CV(candidate)^2)} \approx 9.9\%$ on a single-batch comparison, shrinking as the square root of the number of batches you average.

The browning window is narrow. The 5-minute Maillard step builds kokumi; a 9-minute version builds *bitterness*, pushing IFI-B up with scorched melanoidins and burying the very depth you were chasing. Brown the paste one shade, not two.

Never boil. The 75–80 °C ceiling is not daintiness. A rolling boil degrades the γ-glutamyl peptides you are trying to extract, volatilizes aromatics, and pulls bitter alginates and excess iodine out of the kombu. Heat is a tool here, held precisely, exactly as it is everywhere else in this work.

Each axis needs the other. The multiplication cuts both ways when a factor goes to zero. Pure umami without kokumi — a dash of glutamate in water — tastes savory but thin and short, a note that arrives and leaves too fast, because there is no ringing tail to extend it. Pure kokumi without umami reads as a faint mineral fullness, a receptor ringing over nothing, because it has no bright signal to amplify. Neither axis is a seasoning you can bolt on alone. They are two halves of one taste, and the product of anything with zero is zero.

What the Index does not score. Two honest boundaries. First, saltiness is not one of our eight axes — the sodium channel it runs through is outside this scorecard — so the salt in a finished dish is a formulation variable the profile cannot see. The practical gift of a

high U×K base is that you reach for far less of it; the depth does the work salt used to do. But that is an observation about the eater, not a number on the card. Second, CaSR is promiscuous by nature — it reads calcium, magnesium, polyamines, and a whole bench of L-amino acids — so an honest IFI-K is always measured against a fixed calcium background, with the kokumi-specific component isolated from the rest. We report the controlled number, and we tell you it is controlled.

The model is first-order, and people differ. Writing the felt total as peak × width is a rectangle thrown over a curve that is not a rectangle; it holds in the regime we measure and should not be pushed past it — the true value is the full integral of I(t), and near saturation the Hill curve flattens and the gains compress. And the receptors themselves vary from person to person in expression and genotype, so the *felt* depth of a given bowl will not be identical across a table even when the receptor-level numbers are fixed. The Index measures the stimulus with precision. The last few centimetres, from the taste cell to the person, belong to the eater — which is, after all, where we always wanted the pleasure to live.



Garlic, cross-section — 560 constructs, the Library's richest source of the γ-glutamyl peptides that are the kokumi half of the intersection.

What it means at the table

Strip away the receptors and the integrals and the error bars, and what remains is the oldest thing our family knows: a pot of plant broth that makes a person set down their spoon and feel cared for. We can now say, molecule by molecule, why it does that — a glutamate flytrap snapping shut while a calcium flytrap rings beside it, each lowering the other's threshold, the fast signal and the slow one braided into a single depth that the arithmetic says should be two and a half times either one alone. The reverence and the number were never in tension. The number is how we keep the reverence honest, and the reverence is why the number was worth finding. A stock built on both savory axes is the sound of two cousin receptors agreeing — and when you cook from it, you will find, every time, that

you salt it far less than you expected, that the body takes more from it than its weight should allow, and that *deep*, a word cooks have used for ten thousand years without being able to define it, turns out to have been a product all along. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 14 · The amplifier — *salt, sugar, and fat, reduced.*

Here is the chapter the processed-food R&D reader opened the book for, and the one where kokumi stops being a pleasure and becomes a lever on the largest problems in the modern food supply. The central, repeatable, bankable fact of the axis is this: **kokumi makes a food taste as satisfying with less salt, less sugar, and less fat than it needs without it.** Not by replacing them — there is no salt in a kokumi peptide, no sugar, no fat — but by enhancing the perception of what remains, so that a reduced-sodium soup reads as full-flavored, a reduced-sugar dessert reads as round, a lower-fat sauce reads as rich. This is the mechanism behind an enormous amount of better-for-you reformulation, and the receptor biology explains exactly why it works and exactly where its limits are. We love this axis most of the three times we have said *we cannot explain it yet* and then, a season later, could. What follows is the explanation, worked all the way down to the numbers, because a lever you cannot measure is a lever you cannot trust.



The research plot at scale — reformulation is where the axis earns its keep, letting a food hold its satisfaction while the salt, sugar, and fat come down.

The receptor that reads richness, not salt

Everything about reformulation turns on a single fact of anatomy: the sensation we are rebuilding is read by a receptor that never touches sodium, sugar, or fat at all. Kokumi is the perceptual signature of the **calcium-sensing receptor, CaSR**, and CaSR is a different kind of machine from the ion channels that carry heat and the simple dimers that carry sweet.

It belongs to class C of the G-protein-coupled receptors — the same structural family as the metabotropic glutamate receptors and the sweet and umami T1R dimers — and it works as a disulfide-linked

homodimer, two identical protein chains clasped together. Each chain opens with a large bilobed extracellular domain, the *Venus flytrap*, named for the way its two lobes hinge shut around what they bind. A cysteine-rich linker threads the flytrap down to a seven-pass transmembrane bundle buried in the cell membrane, and it is that bundle, when the flytrap closes, that pivots and couples to the G-protein waiting inside the cell. The receptor's native business is calcium: it reads the concentration of free Ca^{2+} in the fluid around the cell across roughly the 1-to-10-millimolar range, and the two flytraps close cooperatively as calcium rises.

Kokumi compounds are not calcium. They are **positive allosteric modulators** — a word worth unfolding, because the whole axis lives inside it. An allosteric modulator binds somewhere other than the main agonist site and changes how the receptor answers the agonist. A *positive* one makes the receptor more sensitive. The γ-glutamyl peptides that carry kokumi settle into the cleft of the Venus flytrap, near where the receptor reads amino acids, and they stabilize the closed, active shape of the lobes. They do not shout on their own. They lower the price of admission for calcium, so that at the ambient calcium already bathing the tissue, the receptor closes further and signals harder than it otherwise would. In the language of the dose-response curve, a kokumi peptide takes the receptor's answer to calcium and slides it to the left.

Inside the cell, the closed receptor couples mainly through Gq/11 to phospholipase C-β, which cleaves PIP_2 into IP_3 and diacylglycerol and releases a pulse of calcium from internal stores — the signal the brain eventually reads. In the tongue, CaSR sits in the taste bud alongside the cells that carry sweet, salty, and savory, and here is the subtle, important part: its activation does not generate a taste of its own. There is no "kokumi flavor" the way there is a flavor of sugar. What CaSR does is *modulate the whole percept* — it lengthens it, broadens it across the mouth, and raises the apparent intensity of the sweet, salty, and umami signals arriving beside it. The three sensations the tasting literature has always named — richness, mouthfulness that spreads to the back and sides of the palate, and continuity that keeps the flavor present after the swallow — are the perceptual shadow of that modulation. That is why the axis is invisible on its own and transformative in company. It is the receptor that reads body, and body is exactly what sodium, sugar, and fat were quietly supplying.

The molecules of body, and how we grow more of them

The kokumi peptides are small, humble, and everywhere in the food we already revere. The archetype is **glutathione** — γ-glutamyl-cysteinyl-glycine, the same tripeptide every living cell makes by the gram to manage its own chemistry — and its cousins differ only in which residues hang off the γ-glutamyl handle. The handle is the signature: the glutamate is joined to the next residue not through its usual α-carboxyl but through the γ-carboxyl further down its side chain, an unusual bond that ordinary peptidases cannot clip, which is part of why these molecules survive cooking and reach the receptor intact.

Three enzymes build and reshape them, and all three are ordinary plant and microbial machinery we can turn up rather than invent:

1. **Glutamate-cysteine ligase (GCL),** which fuses glutamate to cysteine to make γ-glutamylcysteine. This is the rate-limiting step, and in nature it is feedback-inhibited by the very glutathione it leads to — the cell throttles its own supply. We write in a feedback-insensitive variant so the throttle comes off.
2. **Glutathione synthetase (GSS),** which caps γ-glutamylcysteine with glycine to finish glutathione.
3. **γ-glutamyltransferase (GGT),** which lifts the γ-glutamyl group off glutathione and hands it to a different acceptor — a valine, a valyl-glycine — to make the more potent peptides. This is the shunt we tune most carefully, because it decides *which* kokumi molecule the plant accumulates, and they are not equal.

The one that matters most is **γ-glutamyl-valyl-glycine**, γ-Glu-Val-Gly. In the standardized CaSR assay it potentiates the receptor at roughly an order of magnitude lower concentration than glutathione does — about **twelve to thirteen times more potent** — so a plant that routes even a modest fraction of its glutathione pool into γ-EVG becomes far louder on the axis than its raw peptide tonnage would suggest. Garlic, which our kitchens and our ancestors chose long before anyone could read a receptor, hoards a related family built on S-allylcysteine; its deep, rounding savor when it is slow-roasted is γ-glutamyl chemistry warmed into fragrance. The engineering, then, is not alchemy. It is the same move the heat program makes on the capsaicin pathway, aimed at a different enzyme set: relieve the feedback, raise the flux, and bias the terminal transferase toward the most potent product the pathway can carry.

PEPTIDE	SEQUENCE	MW (G/MOL)	ROLE IN THE PATHWAY
γ-glutamylcysteine	γ-Glu-Cys	250.3	GCL product; the glutathione precursor
Glutathione	γ-Glu-Cys-Gly	307.3	GSS product; the IFI-K reference compound
γ-glutamyl-valyl-glycine	γ-Glu-Val-Gly	303.3	GGT product; the potent one, ~12.8× glutathione
γ-glutamyl-S-allylcysteine	γ-Glu-SAC	290.3	garlic's native kokumi store, warmed by roasting

Across our forty-two host organisms, five carry this chemistry natively and respond to the edits, and they read like a map of the world's savory foundations — the allium that starts almost every pot, the legume that becomes miso and shoyu, the yeast that deepens a stock, the onion, the pod of vanilla that most people never think of as savory and yet rounds a custard the way a stock rounds a soup.

ORGANISM	CONSTRUCTS	PRIMARY AXES	ENGINEERING DIFFICULTY
Garlic	560	Kokumi · Pungency	Moderate
Soybean	360	Umami · Kokumi	Frontier
Yeast	240	Umami · Kokumi	Hard
Onion	40	Kokumi	Moderate
Vanilla	40	Kokumi	Hard



Engineered garlic in cross-section — the allium hoards γ-glutamyl peptides, the molecules of body, and roasting warms them into the roundness a reduced soup is missing.

The arithmetic of loudness

A number is a form of respect for the reader, so here is the axis as arithmetic. The Index scores every stimulus by how loudly it speaks to its receptor, not by how much of it a food holds, and loudness is the inverse of the concentration needed to half-open the receptor:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. A *lower* EC_{50} means the receptor answers to less of the compound — a louder voice — and because the EC_{50} sits in the denominator, loudness and score rise together. By convention, **IFI-K = 1,000 is defined as the half-maximal CaSR response of reduced glutathione** in a standard calcium-clamped assay held just below threshold, so that the peptide's potentiation is the only signal moving — a fixed reference point, the same way the Scoville scale is defined by a dilution-to-threshold. Call the reference's half-activating concentration E_ref. Everything else on the axis is E_ref divided by the food's own half-activating concentration, times a thousand. We set E_ref to **2.0 mM glutathione** as a stated convention; fix that anchor and every IFI-K reading converts straight back into a concentration you can weigh.

Work the canonical values through it. Our model predicts an **IFI-K of 1,600** for the reduced-sodium tomato sauce. Rearranged, 1,600 = 1,000 × E_ref ÷ E_sauce gives E_sauce = E_ref × (1,000 ÷ 1,600) = 2.0 mM × 0.625 = **1.25 mM** — the sauce's kokumi load is designed to half-activate the receptor at a quarter again less peptide than the reference needs, which is to say 1.6× louder. Our full-spectrum flagship **Phoenix** is predicted at **IFI-K 1,900**; the same step gives E = 2.0 × (1,000 ÷ 1,900) = **1.05 mM**, 1.9× louder than reference. The scorecard's ceiling of 5,000 corresponds to a food that half-activates at 2.0 × (1,000 ÷ 5,000) = **0.40 mM**, five times louder than the reference — and that ceiling exists for a reason the next line makes plain.

Score a *pure* peptide and you run off the top of the page. γ-Glu-Val-Gly, at ~12.8× the potency of glutathione, half-activates at E_ref ÷ 12.8 = 2.0 ÷ 12.8 = **0.156 mM**, which the formula renders as IFI-K = 1,000 × 12.8 = **12,800** — far above the 5,000 at which we cap the scorecard. That gap is the lesson: a single engineered peptide is

deafening, while a whole food is a blend, its loudest molecules diluted among everything else on the plate, so even a superbly tuned cultivar sits near 1,900 rather than near 12,800. We cap the axis at 5,000 precisely so that the scorecard describes foods a mouth will actually meet, not isolated compounds in a cuvette.

ENTRY	IFI-K	EC_{50} ÷ E_REF	EC_{50} (MM)	LOUDER THAN REFERENCE
Reference · glutathione	1,000	1.000	2.00	1.0×
Reduced-sodium sauce · predicted	1,600	0.625	1.25	1.6×
Phoenix · SS-0096 · predicted	1,900	0.526	1.05	1.9×
Scorecard ceiling	5,000	0.200	0.40	5.0×
γ-Glu-Val-Gly, pure	12,800	0.078	0.156	12.8×

The molecular weights turn those concentrations into something you can put on a scale. Glutathione at 307.3 g/mol means the reference's 2.0 mM is 2.0 mmol/L × 307.3 mg/mmol ≈ **615 mg/L** — about six-tenths of a gram of peptide per liter to reach half-activation. γ-Glu-Val-Gly, at 303.3 g/mol, reaches the same half-activation at 0.156 mmol/L × 303.3 ≈ **47 mg/L**. Thirteenfold less material for the same receptor answer: that is what "potent" means once you make it weigh something, and it is why the terminal transferase — the enzyme that decides how much of the pool becomes γ-EVG — is the one we tune down to the codon.

Now the part that explains why a *small* gain in loudness reads as a *large* gain in satisfaction, which is the whole commercial surprise of the axis. CaSR is cooperative: its two flytraps close in concert, so the dose-response curve is not a gentle ramp but a steep sigmoid, with a Hill coefficient near 3. Write the response as $R = [Ca]^n \div (K^n + [Ca]^n)$, where K is the half-activating calcium and $n \approx 3$. A kokumi peptide left-shifts K. Suppose the ambient calcium in the taste-bud microenvironment sits right at the baseline K, so the receptor idles at 50% of full response. Let a kokumi load left-shift the curve by a factor of 1.5, lowering K to K/1.5. The new response is $R = K^n \div ((K/1.5)^n + K^n) = 1 \div ((1/1.5)^3 + 1) = 1 \div (0.296 + 1) =$ **0.77**. A 1.5× shift in sensitivity has lifted the receptor from 50% to 77% of full response — a bigger jump than the shift itself, because steepness multiplies it near the operating point. This is the biophysics under the kitchen observation that a *little* kokumi does a *lot*: on a steep cooperative curve, modest potency buys disproportionate perceived fullness.

A number without an error bar is an opinion, so attach one. IFI-K is a ratio, and the relative uncertainties of a ratio add in quadrature. Our engineered organisms hold a batch-to-batch coefficient of variation at or below **7%**; with the reference calibration held to the same 7%, the IFI reading carries $\sigma/IFI = \sqrt{(0.07^2 + 0.07^2)} = \sqrt{0.0098} \approx$ **9.9%**. So the sauce's predicted IFI-K 1,600 carries a band of 1,600 ± ~158 at one standard deviation — a band of roughly 1,440 to 1,760, tight enough to formulate against and to reproduce next season. Run the same math on an unimproved line at wild-type's **42% CV** and the band balloons: $\sigma/IFI = \sqrt{(0.42^2 + 0.07^2)} \approx 42.6\%$, or 1,600 ± ~680, a band from 920 to 2,280 that means almost nothing. The entire discipline of the ≤7% program is what turns the kokumi gesture into

a kokumi *measurement* — the difference between a cook saying *more or less this much* and a formulator writing *1,600, and it will be 1,600 again.*

Salt, the biggest prize

We take salt first, because it is the largest reformulation target in processed food and the hardest, and because the receptor story explains both the difficulty and the way through. Sodium does far more than taste salty. The sodium ion enters taste cells through dedicated channels and reads as saltiness, yes, but alongside that it suppresses bitterness, lifts sweetness, and — the part that undoes most reductions — *carries the perception of savory body.* Pull the sodium out of a soup and it does not merely taste less salty; it tastes flat, hollow, thin. The body collapses, because a good share of that body was riding on the salt.

Kokumi rebuilds the body from a direction sodium never touched. Because CaSR reads richness through calcium and γ-glutamyl peptides, not through the sodium channel, a kokumi-loaded base *feels full* at a salt level that would otherwise read as empty. The two signals are additive in the mouth and independent at the receptor, which is precisely the property a reformulator needs: lower one lever, raise the other, and the sum — satisfaction — holds. We build our broth base and our mother stock around exactly this. We set the sodium deliberately low and let the engineered kokumi carry the weight, and in trial after trial the pattern repeats with the steepness the Hill math predicts: the kokumi-enhanced product holds its palatability down to a substantial sodium reduction, while the unenhanced control, past a point, falls off a cliff.

We state this as an *attribute*, carefully and always: kokumi supports flavor at lower salt. We make no health claim about the sodium reduction itself — that is the formulator's and the regulator's domain, and this book holds that line without exception. Our job is the flavor that makes the reduction survive the panel, which is the exact step at which most well-intentioned reductions quietly die.

Fat, the subtler case

Fat is the more remarkable case, and in some ways the more hopeful. A great deal of what we love about fat is not flavor at all — it is *mouthfeel*: richness, coating, body, the way a sauce stays with you after the swallow. Those are, almost word for word, the kokumi sensations, because both are continuity and spread, and both are read as fullness rather than as a specific taste. So kokumi can stand in for part of fat's sensory role without being a fat, giving a reduced-fat product back the body and the lingering fullness that stripping the fat takes away. A lower-fat sauce built on the mother stock does not read as lean, because the kokumi is supplying the mouthfulness the eye was expecting the fat to supply.

It is not a perfect substitution, and we are precise about why. Fat also carries aroma — a whole class of volatile molecules are fat-soluble and ride into the nose only on a fat phase — and fat lubricates in a mechanical way a receptor signal cannot mimic.

Kokumi replaces the *body* and the *continuity* of fat, which are the largest part of what reduction costs, but not the aroma and not the slip. For those we lean on other levers. What makes kokumi the closest tool food has for the body of fat is that it operates on the same perceptual channel fat operates on — continuity and fullness — through a receptor rather than through a texture, so the mouth reports the richness even though the fat phase that usually delivers it is thinner.

Sugar, the roundness given back

Sugar, finally, carries body as well as sweetness, which is why a reduced-sugar dessert or beverage reads thin and sharp — the sweetness drops *and* the roundness that cushioned it drops with it, and what is left can taste edged, almost sour-bright. Kokumi gives back the roundness, so the lower-sugar version reads complete rather than hollow. There is a second gift the receptor biology hands us here: because CaSR activation raises the *apparent intensity of the sweet signal arriving beside it*, kokumi does not only restore body, it quietly lifts the perception of the sugar that remains. A reduced-sugar custard built on our engineered vanilla is designed to show exactly this doubled action — its sweetness is modest for a custard, yet it does not taste modest, because the kokumi is both rounding the body and amplifying the little sweetness that is there.

Why the naive approaches fail

It is worth being plain about why this is so much harder than it sounds, and why it is worth a whole axis of engineering rather than a single additive — because the obvious shortcuts all fail, and they fail at the receptor in ways the Index can name.

The naive approach to salt reduction is to use less salt, or to swap in potassium chloride. Less salt reads as flat, for the reason already given: the body was riding on the sodium. Potassium chloride brings its own saltiness, but the potassium ion is caught almost immediately by the TAS2R family of bitter receptors and reads as a metallic bitterness that no amount of the salt's own virtue can hide. Kokumi rings none of those bitter receptors; it works through CaSR and leaves the TAS2R channel untouched, so a kokumi lift lets a formulator lean far less on the KCl crutch, or not reach for it at all.

The naive approach to fat reduction is to use less fat and thicken with a gum. The gum raises viscosity, so the liquid is physically thicker, but viscosity is not body — a gum rings no receptor at all, so the mouth registers the uncanny *thick but empty*, the hollow heaviness of badly reduced food. Kokumi succeeds where both shortcuts fail for one structural reason: it does not try to replace the missing thing with a cheaper version of the same thing. It supplies the *perception* the missing thing was delivering, from a receptor the reduction never touched. The soup is still low in sodium. The sauce is still low in fat. But the calcium-sensing receptor is ringing, and the mouth reports *full.* That is the reformulator's whole dream stated plainly — a lever that acts on satisfaction without acting on the nutritional target — and it is the one honest way we know to reach it.

A worked reduction

The R&D reader wants to see the reduction, not only hear about it. This is a tomato pasta sauce reformulated to a markedly lower sodium than a conventional jarred sauce, holding its satisfaction entirely through the kokumi of the engineered organisms — the kind of product this book exists to make possible.



Engineered tomato — the base of the worked reduction, its own umami restored, carrying its satisfaction on kokumi rather than on salt.

```
LOWER-SODIUM TOMATO SAUCE
─────────────────────────────────
800 g          engineered tomato, cooked down
2 tbsp         roasted garlic kokumi paste
150 ml         mother stock
1 tbsp         engineered-soy red miso
handful        basil
— salt: ~40% below a conventional jarred sauce —
─────────────────────────────────
yield  ~700 ml    ·    time 40 min
```

1. Cook the tomato down over medium heat **25 minutes** until thick and darkened, stirring so the fond that builds gets scraped back in — that fond is Maillard kokumi and umami, free and precious, γ-glutamyl chemistry and glutamate born together in the browning.
2. Stir in the garlic paste and the mother stock and simmer gently **10 minutes** to marry, letting the engineered peptides spread through the whole pot.
3. Off the heat, stir in the miso — the low-heat rule, because the γ-glutamyl transferase chemistry and the live culture both prefer not to be boiled — and tear in the basil. Salt *last*, tasting as you go, and notice, because this is the whole lesson, how little it takes before the sauce reads complete. You will stop well below where you expected to.

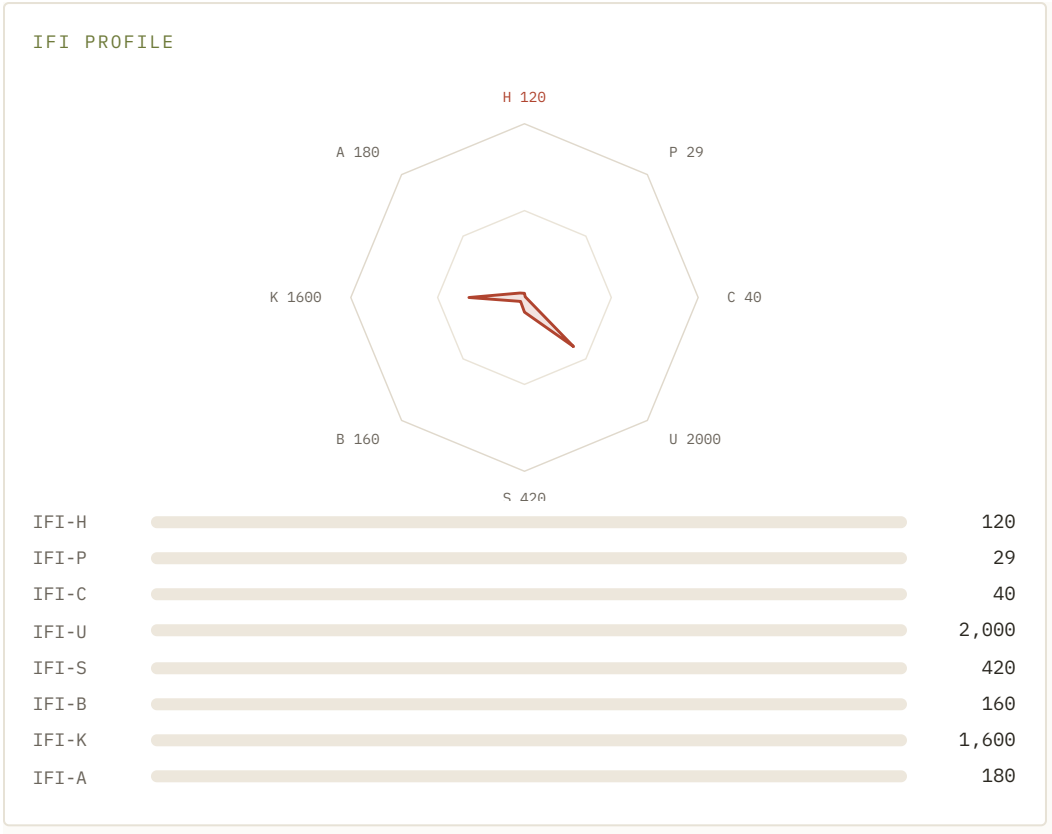

Predicted kokumi of **1,600** and umami of 2,000, at a sodium level roughly 40% below a conventional sauce's — a profile designed to taste complete, not reduced. Work the kokumi number back through the arithmetic and it says the sauce is designed to half-activate CaSR at about 1.25 mM glutathione-equivalents, 1.6× louder than the reference, carried on the engineered garlic and soy rather than on the salt jar; the umami and the kokumi stand beside each other at the receptor and sum in the mouth to the body a cook used to buy only with sodium. That gap — between the salt you used and the satisfaction you got — is the commercial and the human case for the entire axis.

For contrast, hold that modest, purposeful profile beside the full-spectrum flagship, which exists to prove how high every axis can be driven at once:

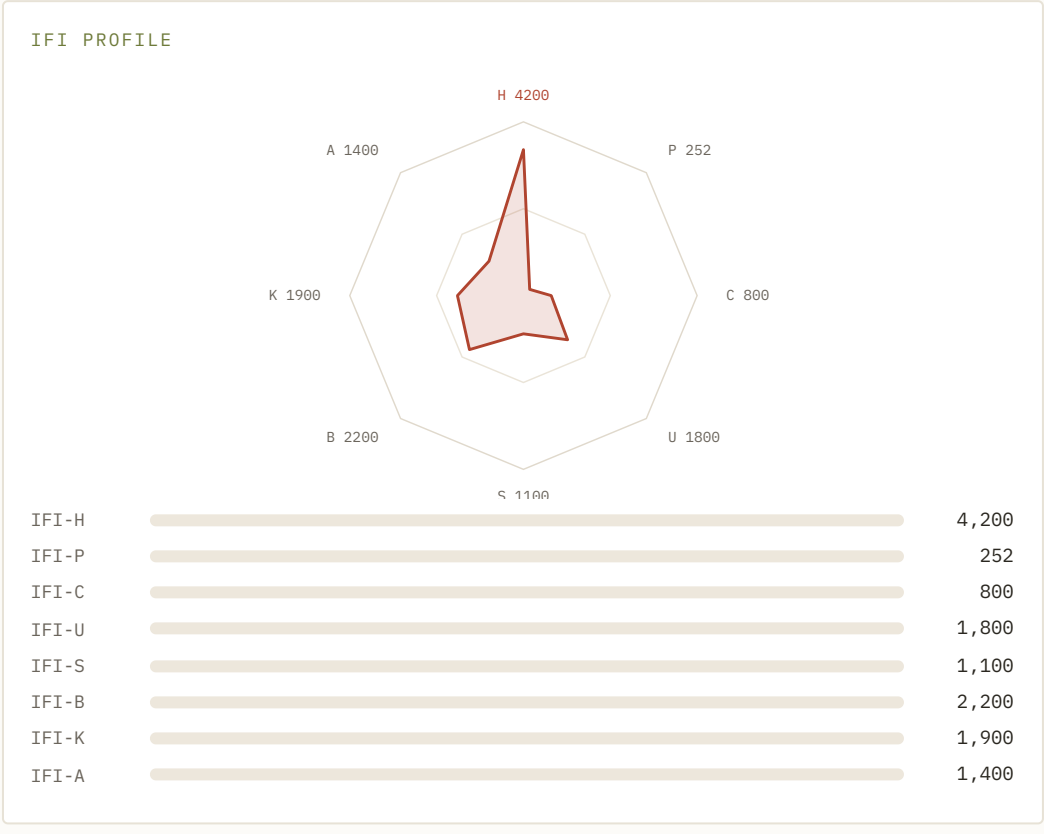

The flagship's predicted kokumi of **1,900** — only a fifth again louder than the sauce — tells you something true and slightly humbling about this axis: it tops out gently. The useful range of IFI-K is narrow

and crowded near the bottom, because whole foods are blends and the loudest peptides are always diluted by everything else on the plate. The art is not to chase the ceiling. It is to place a 1,600 exactly where a reduction needs it.

Honest limits

Every lever has a reach, and kokumi's is specific enough to be worth stating plainly, because a tool oversold is a tool mistrusted. Kokumi amplifies what is present; it originates nothing. It adds no saltiness of its own, no sweetness, no fat — so a base reduced to the point where there is nothing left to carry gives the receptor nothing to lengthen, and below that floor the fullness simply is not there to be restored. Worse, because it enhances *continuity*, kokumi lengthens whatever it is given, flaws included: laid over a thin or off-tasting base it will make the off-note more present and more persistent, not less. Kokumi rewards a good base and punishes a careless one.

It does nothing for aroma, which matters most in fat reduction, where a real part of the loss is volatiles that only travel on a fat phase; that gap needs a different tool. It has an optimum, not a monotone — push the γ-glutamyl load too high and the sulfur notes of the glutathione family begin to read as heavy, soapy, or cloying, so the dose has a top as well as a bottom. And the IFI-K number, honest as the ≤7% program makes it, is still a receptor reading taken in a cuvette; the mouth meets the molecule inside a whole matrix, at temperature, beside seven other axes that shift it, so the score predicts the perception closely but never perfectly. We publish the number with its error bar and its caveats because that is the respect the reader is owed. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

And we hold, without exception, to the one line that matters most: we describe an attribute, never a health outcome. Kokumi supports flavor at lower salt, lower sugar, lower fat. Whether a given reduction does anything for a body is the formulator's and the regulator's to establish, and ours to stay carefully out of. What we hand them is the flavor that lets the reduction survive the panel — the step at which most good intentions quietly fail.

That is the quiet trade at the center of the axis, and close to the center of what we think good food engineering is for. For the maker it is a reformulation that passes the panel. For the eater it is a food that gives the body what it asks for and a little less of what it might do with less of, without ever announcing the exchange. The salt comes down, the sugar comes down, the fat comes down, and the satisfaction stays — read off a receptor that was never watching the thing we reduced. We grew the molecules of body back into the plant on purpose, computed them to a tenth, and let the mouth report *full.* It is a small, exact kindness, and we are proud of it.

Chapter 15 · The kitchen of continuity — *a short book of rich dishes.*

This book has been the axis explained. Here it is the axis cooked — a short run of complete dishes, each built on the engineered organisms and the slow preparations our family keeps on hand, each scored on all eight axes so you can read the kokumi doing its work in the full, crowded context of a plate rather than alone in a tasting glass. They run plant-first, the way we cook and the way we have eaten our whole lives, with the one seafood dish written — as always — with equal care and no apology. What they share is the thing this whole book is about: a richness that makes each of them feel finished, held, and *complete*, built on purpose out of small molecules speaking to a receptor named for calcium.

Before the dishes, the mechanism they all turn on — and the arithmetic that lets us promise the same plate twice.

The receptor named for calcium

Kokumi has no taste of its own. You cannot point to it the way you point to sweet or sour; it is not a flavor so much as a *behavior* of flavor — the sense that a mouthful has body, that it arrives rounded instead of thin, that it fills the whole mouth and then stays a beat longer than it should. For a long time that was all anyone could say. The word is Japanese, roughly "rich taste," and for decades it sat in the drawer of things cooks knew and science could not yet name.

The receptor turned out to be one already famous under another job. CaSR, the calcium-sensing receptor, is a class-C G-protein-coupled receptor the body uses to watch the calcium level of the blood — it is how the parathyroid knows, minute to minute, whether to call for more. It also sits on a subset of taste cells on the tongue, and there it does something subtler than count calcium. Its large extracellular domain folds like a Venus flytrap, two lobes hinged on a cleft; the true agonists — calcium and other small cations — bind deep in that cleft and snap the trap shut, and the closure is the signal.



Engineered garlic — the SS-series allium carries the most kokumi construct work of the group. The γ-glutamyl peptides that read as richness are the plant's own; we turned the native route up and held it steady.

The molecules that carry kokumi are γ-glutamyl peptides — reduced glutathione (γ-Glu-Cys-Gly) chief among them, and shorter or substituted cousins like γ-Glu-Val-Gly, γ-Glu-Cys, and γ-Glu-Val. The signature is that γ-linkage: the peptide bond is formed through glutamate's side-chain carboxyl, not the usual α-carboxyl at the backbone. That bond is unusual, enzyme-made, and resistant to the ordinary peptidases of the gut and the pan, which is part of why these molecules survive cooking and reach the tongue intact. And they do not snap the flytrap shut themselves. They bind a second site on the outer face of the domain and act as positive allosteric modulators — they make the trap easier to close, so the receptor answers the mouth's ambient calcium and cations more loudly than it otherwise would.

That single fact explains every strange thing about kokumi at once. Why it has no taste alone: a modulator needs the receptor's real agonist present to amplify. Why it lengthens and fills rather than adding a note: it is turning up a signal, not writing a new one. Why it makes sweet taste sweeter and a broth taste deeper: the CaSR-bearing taste cells sit among the others, and a mouth made more responsive reads the whole co-located chord louder. Why, in the end, it reads as *continuity* — the feeling that a flavor is being held open.

Where the peptides come from is the plant and the ferment themselves. γ-Glutamylcysteine synthetase and glutathione synthetase build glutathione in every living cell as its main antioxidant; γ-glutamyl transferases then shuffle the γ-glutamyl group onto other amino acids to make the Val-, Cys-, and Val-Gly-forms. Alliums are rich in all of it — a large part of why a long-cooked onion or a roasted head of garlic tastes the way it does. Aged cheese, soy sauce, miso, fish sauce, yeast extract, and the muscle of a scallop all carry these peptides too, concentrated by age or ferment or simply by being tissue that stores glutathione. Our engineering invents none of this. It turns up a route the plant already runs.

Reading a plate in glutathione-equivalents

The Index gives every axis one formula.

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — for a kokumi peptide, the concentration at which it delivers half its maximal potentiation of CaSR. Invert it against a fixed reference and scale by a thousand, and potency becomes a single score: a candidate that works at a lower concentration than the reference scores above 1,000; one that needs more scores below. It is loudness, not quantity — not how much of a compound a dish holds, but how loudly that compound speaks.

For the kokumi axis we peg 1,000 to reduced glutathione, the body's own γ-glutamyl tripeptide and the receptor's reference ligand, at its half-potentiating concentration, which our receptor assay places near 2.0 mM. Everything else is read against it.

γ-GLUTAMYL COMPOUND	FORM	EC_{50} AT CASR (MM)	IFI-K (RAW)
Glutathione · γ-Glu-Cys-Gly	tripeptide (reference)	2.00	1,000
γ-Glu-Cys	dipeptide	0.80	2,500
γ-Glu-Val	dipeptide	0.50	4,000
γ-Glu-Val-Gly	tripeptide	0.040	50,000 → caps 5,000

A worked case: γ-Glu-Val-Gly, the most potent kokumi tripeptide we work with, delivers its half-potentiation near 0.040 mM — fifty times less material than glutathione for the same effect. IFI-K = 1,000 × 2.00 ÷ 0.040 = 50,000, which the scorecard clamps to its normalization ceiling of 5,000. That one tripeptide is the loudest word in the kokumi language. The dishes below speak in whole sentences of it — diluted, cooked, and in company.

Because the whole plate, not a pure compound, is what you eat, we score a dish the same way we score a molecule: take its kokumi fraction, run its own dose–response against the receptor, and read the half-activating concentration back through the formula in glutathione-equivalents. Rearranged, EC_{50}(dish) = 1,000 × 2.00 ÷ IFI-K. The five dishes in this chapter land like this.

DISH	IFI-K	HALF-ACTIVATING CONC. (MM GSH-EQ)	LOUDER THAN REFERENCE
Shiitake–barley, mother stock	2,600	0.77	2.6×
Scallop, kokumi dashi	2,500	0.80	2.5×
White beans braised in stock	2,400	0.83	2.4×
Nine-hour onion	2,300	0.87	2.3×
Dark chocolate cream	1,700	1.18	1.7×

The rightmost column is the formula read back in plain words: a plate at K 2,600 half-activates the receptor at a stimulus 2.6 times louder than the glutathione reference. A lower EC_{50} is a louder plate.

Why *potentiation* and not activation? Because the peptides do not open the receptor; they make it easier to open, and the math has to carry that. CaSR answers calcium and the other cations of the mouth with a steep, cooperative curve — a Hill coefficient near 3 — and a γ-glutamyl modulator binds its separate site and shifts that curve leftward, so the receptor half-answers at a lower cation level. A two-constant model is enough to see it. With a modulator at concentration [P], binding constant K_P, and a maximal cooperativity factor α:

> $C_{50}([P]) = C_{50}^{0} \times (1 + [P]/K_P) \div (1 + \alpha \cdot [P]/K_P)$

Take a strong tripeptide: α = 4, K_P = 0.05 mM. At [P] = 0.05 mM the half-activation point moves by (1 + 1) ÷ (1 + 4) = 0.40 — the receptor becomes 2.5 times more sensitive. Push to [P] = 0.20 mM and it is (1

+ 4) ÷ (1 + 16) ≈ 0.29, a 3.4-fold sensitization; at saturation the ceiling is 4-fold. A plate reading K 2,400 is behaving, at the receptor, like a stimulus 2.4 times the reference — squarely what a few tenths of a millimolar of good peptide buys in this model. And because the underlying cation curve is so steep, that leftward shift does not fade the sensation up gently; it swings it. Kokumi *fills in*. That is the arithmetic of why richness feels like something arriving all at once rather than growing.

One ruler holds all of it. The taste axes anchor on molar EC_{50}; heat anchors on Scoville, which is itself a dilution assay, one part capsaicinoid by weight registering about sixteen Scoville units. The Index maps both onto the same dimensionless 0–5,000 scale by the one ratio-to-reference rule — each axis pegged to its own reference convention, so a score on the K axis and a score on the H axis answer *how loudly* in their own terms and are not directly comparable across axes. That is how a gentle K 2,300 onion and a pepper our model predicts at approximately 2.95 million SHU and IFI-H 4,200 share a scorecard: different receptors, different chemistries, each axis asking *how loudly* on its own benchmark.

A number is only worth printing if it comes back. The platform holds batch-to-batch variation at or under seven percent, where a wild plant's γ-glutamyl content drifts forty percent and more across a season. On the K axis that is the whole difference between a dish you can write down and one you can only hope for. Take the white beans at K 2,400: held to a seven-percent window, the plate lands between 2,232 and 2,568 — the same beans every time. At a wild plant's roughly forty-percent drift the same recipe would wander from about 1,392 to 3,408, a different dinner each cooking. Layering helps further: a dish that draws its kokumi from three independent engineered sources at seven percent each — the onion, the garlic paste, the stock — averages their variation down toward 7% ÷ √3, near four percent, and the receptor's own saturating response damps it further still. Precision is not the opposite of warmth. It is what lets us hand you the warmth twice.

The plants that carry the depth

None of this is a foreign molecule smuggled into a vegetable. The γ-glutamyl peptides are the plant's own — its glutathione machinery is as old as the cell — and our work is only to turn that native route up and hold it steady. The same platform that designs heat designs this: a protein language model scores candidate edits to the biosynthetic enzymes, a genome-scale flux-balance model predicts which steps control flux, and the winning edits are written in with CRISPR. It is the method that found the double variant S39L+L345G behind our hottest pepper — a combination no breeder ever stumbled on — turned now toward richness instead of fire. Across forty-two host organisms the kokumi and umami workhorses are these.

ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Garlic	560	K / P	Moderate
Soybean	360	U / K	Frontier
Yeast	240	U / K	Hard
Shiitake	80	U	Hard
Onion	40	K	Moderate
Vanilla	40	K	Hard

Garlic carries the most construct work of the group and the onion the least, yet both are the foundation on which the savory half of nearly every cuisine stands. The soybean and the yeast are where ferment concentrates the peptides further. The vanilla is the quiet surprise — a pod we prize as much for a kokumi mouthfeel as for its vanillin.

A handful of slow preparations recur across the dishes, so here they are once, named plainly. Our *mother stock* is the kitchen's reusable base: engineered alliums, kombu, and dried shiitake drawn out over hours into a deep, γ-glutamyl-rich liquid we keep going and feed. The *roasted-garlic kokumi paste* is engineered garlic roasted soft and sweet, its pungency tamed and its depth concentrated, mashed to a spoonable paste. The *kokumi broth base* is a clearer dashi built on the same engineered kombu and shiitake, used where we want the depth without the color. The *white miso* is fermented from our engineered soybean. The *vanilla* is our own pod. Each is a way of getting the peptides onto the tongue intact; the dishes are ways of getting them there in company.

The nine-hour onion, over anything

The slow onion is the plainest demonstration we know, because it builds kokumi two ways at once — one mechanical, one chemical — and you can taste each arriving.



The cut face of an engineered kokumi onion. Nine hours in a low oven concentrate its non-volatile γ-glutamyl store roughly fivefold as the water leaves, while the browning builds a second kokumi on top of the first.

Two kilograms of engineered onion cook down to roughly four hundred grams — a fivefold loss, almost all of it water. The γ-glutamyl peptides the onion carries are non-volatile and stubborn, so as the water leaves they concentrate about fivefold with it. That is the mechanical half, the pure arithmetic of a drying pan. The chemical half is the long, low browning. Caramelization breaks the onion's fructans down to fructose and glucose and then colors and sweetens them — the S 900 on the card is cooked sugar, not added

— while Maillard reactions between those reducing sugars and the onion's free amino acids throw off melanoidins and a family of savory, kokumi-active fragments, some of them small peptides the receptor reads directly. The sharpness goes the other way: the thiosulfinates and the lachrymatory factor that make a raw onion bite are volatile and heat-labile, and they cook off, which is why the pungency settles at a soft P 13. What remains is sweet with a long savory floor under the sweet.

NINE-HOUR ONIONS

2 kg	engineered kokumi onion, sliced thin
3 tbsp	olive oil
1 tsp	salt
150 ml	mother stock, for deglazing

yield ~400 g jammy onions · time 3 h active + low oven

1. Toss the onions with the oil and salt in a wide heavy pot, cook over medium **30 minutes** stirring, then drop to the lowest possible heat and cook **2–3 hours more**, stirring every 15 minutes and deglazing with a splash of mother stock each time the fond threatens to catch.
2. They are done when they are near-black, jammy, and sweet with a savor under the sweet — the γ-glutamyl store concentrated, the sharpness long gone.
3. Pile over whatever wants depth: toast, polenta, a bowl of the mother-stock beans, a grilled cheese where they meet the natural kokumi of aged cheese and double it.

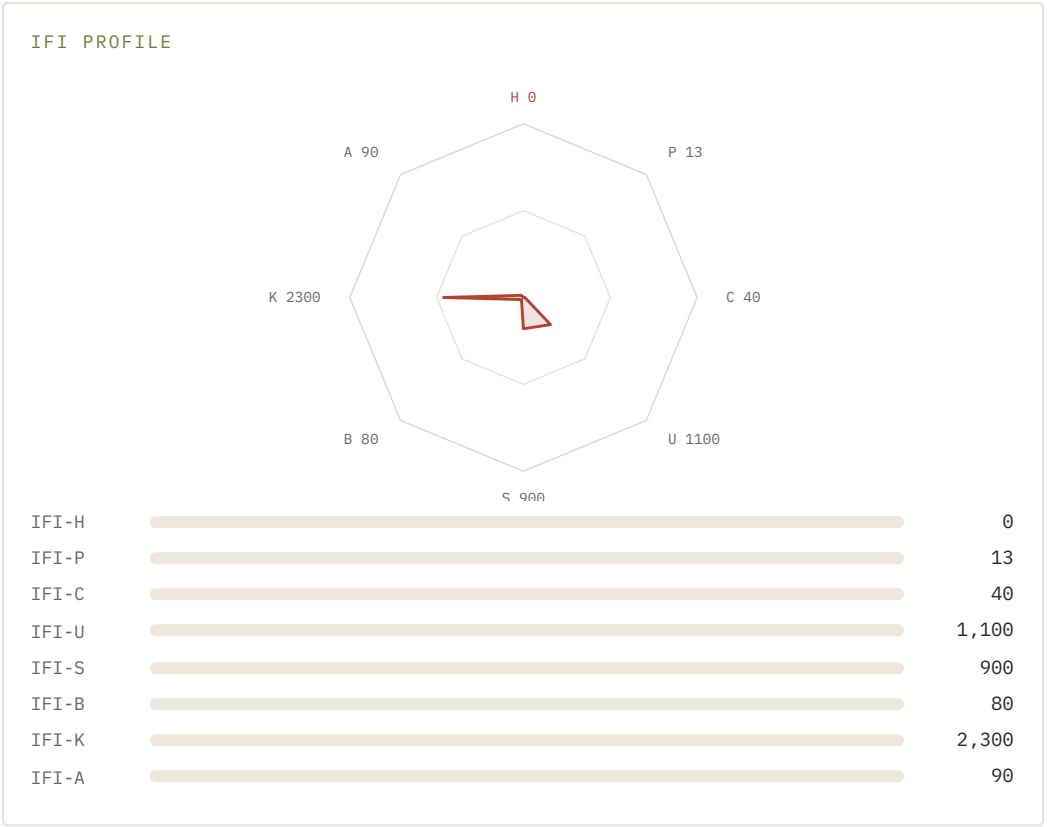

At K 2,300 this half-activates the receptor near 0.87 millimolar glutathione-equivalent — better than two-and-a-third times the reference, out of nothing but onions and time. The one failure mode is heat. Let the fond catch and scorch and you trade kokumi for its opposite: acrid pyrazines and bitter degradation products that climb the B axis and mud the whole pot. The splash of mother stock at every stir is not seasoning. It is temperature control, pulling the pan back from the burn each time the sugars threaten to cross from brown to black.

Shiitake and barley, in the mother stock

This is risotto logic with none of the waste — and the one dish here where umami and kokumi stand shoulder to shoulder, each lifting the other. The shiitake that gave the stock its depth is sliced back into the pan to give the dish its body; nothing of the richness we built is thrown away.

The umami is a synergy. Shiitake is heavy with 5′-GMP, a ribonucleotide that multiplies glutamate's signal at the umami receptor many-fold when the two sit on the tongue together; the mother stock brings the glutamate, the mushroom the GMP, and the dish tastes far deeper than either part alone — the U 2,400, highest of the plant dishes. Over that savor sits the kokumi: γ-glutamyl peptides from the stock and the roasted-garlic paste, and the physical thickening as the barley's starch leaches into the liquid and coats. Richness of two kinds, read as one.

SHIITAKE–BARLEY, MOTHER STOCK

200 g	pearl barley
1.2 L	mother stock, hot
the reserved	shiitake from the stock, sliced
2 tbsp	roasted garlic kokumi paste
30 g	aged hard cheese, grated (optional)

serves 4 · time 45 min

1. Toast the barley dry in a wide pan **2 minutes** until nutty. Add the sliced shiitake and the garlic paste, stir to coat.
2. Ladle in the hot mother stock a third at a time, stirring often, letting each addition absorb before the next — **35–40 minutes** until the barley is tender and the dish is creamy with released starch.
3. Off the heat, beat in the cheese if using, which stacks its natural γ-glutamyl kokumi on top of the engineered. Rest 2 minutes; it tightens and deepens.

IFI PROFILE

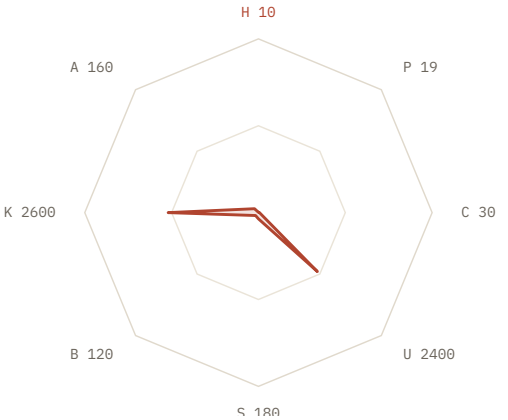

IFI-H	10
IFI-P	19
IFI-C	30
IFI-U	2,400
IFI-S	180
IFI-B	120
IFI-K	2,600
IFI-A	160

At K 2,600 — about 0.77 mM-equivalent, 2.6 times the reference — this is the loudest kokumi on the whole card. If you add the optional aged cheese, you are stacking a natural γ-glutamyl source onto the engineered one, and because the two draw on the same receptor site their contributions do not simply add; they approach the receptor's saturating ceiling together. You taste the last increment as a small one, which is the receptor telling you, honestly, that it is nearly full.

White beans braised in mother stock

The humblest dish here and, to us, the most convincing — because there is nothing in it but beans, stock, time, and a little garlic, and it eats like it was cooked with meat and bones, which it was not. That is the entire argument of the book served in a bowl: depth is a set of molecules, and you can put them there on purpose, from plants.

The beans drink the stock's umami and kokumi deep into themselves over ninety slow minutes. Meanwhile their own released starch — amylose and amylopectin leaching from the softening cells — thickens the surrounding liquor into a silk that is half the dish. That silk is mouthfeel of two kinds stacked: the physical viscosity of dissolved starch, and the receptor-level kokumi of the γ-glutamyl load, read together as one sensation of body.

KOKUMI WHITE BEANS

300 g	dried white beans, soaked overnight
1.5 L	mother stock
2 tbsp	roasted garlic kokumi paste
1	bay leaf + a strip of lemon peel
olive oil, salt to finish	

serves 4–6 · time 90 min

1. Drain the soaked beans and sit them in a heavy pot with the mother stock, garlic paste, bay, and lemon peel. The stock should cover by two fingers; top up with water if not.
2. Bring to a bare simmer and hold it there, barely bubbling, **75–90 minutes**, until the beans are completely tender and the liquor has gone silky — the released starch thickening the kokumi stock into something between a soup and a sauce.
3. Salt only at the end, when the beans are soft, and finish with a thread of raw olive oil. The beans will have drunk the stock's U×K deep into themselves; the liquor is the reward.

IFI PROFILE

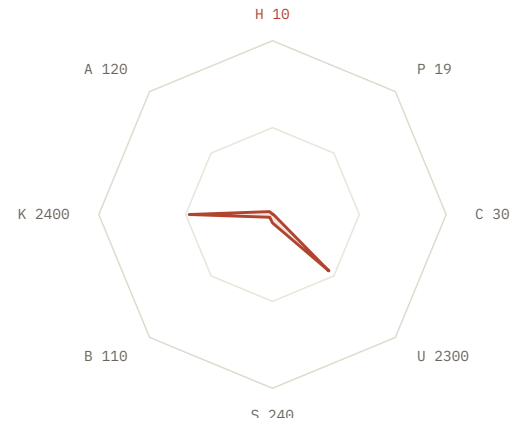

IFI-H	10
IFI-P	19
IFI-C	30
IFI-U	2,300
IFI-S	240
IFI-B	110
IFI-K	2,400
IFI-A	120

We hold the salt for the end — not because salt toughens a bean, which is a kitchen myth, but for control. A modest amount of salt in the water actually *helps* the skins soften, the sodium nudging aside the calcium that cross-links their pectin. What truly toughens a bean is acid and hard-water calcium, so the lemon goes in as a strip of zest for its aroma and its juice is kept back. The reason to salt late is simpler: salt the liquor before the beans are tender and you are seasoning a moving target, because the beans keep drinking it in and the balance you set at minute forty is gone by minute ninety. K 2,400, 0.83 mM-equivalent: a plate that rings the receptor two-and-a-half times louder than glutathione, out of the cheapest bag in the pantry.

Seared scallop, kokumi dashi

The seafood dish, and we write it with the full care this Library gives every animal food, plant-first though our table is. What earns it a place in a book about kokumi is that it rings the receptor from three directions at once — the sea's, the fire's, and the ferment's, one sensation assembled from three chemistries.



Our engineered soybean, the source of the white miso whisked into the dashi — a concentrated γ-glutamyl load from a plant we tuned for the umami and kokumi axes.

The scallop's flesh is unusually sweet: scallops store large amounts of free glycine and other sweet amino acids, which is the S 240 on the card and part of why they taste almost of sugar. They also carry

their own glutathione — a muscle's antioxidant store — and that is kokumi straight from the sea. The sear adds a second kokumi: a dry scallop on a shimmering pan browns because, once the surface water is gone, the face climbs past 150 °C and Maillard chemistry builds its crust of melanoidins and savory peptides. Dryness is the whole secret of that crust; while surface water remains it pins the face at 100 °C and nothing browns, which is why we pat the scallops bone-dry and do not move them. The third kokumi is ours — the engineered-soy white miso whisked into the warm broth, a concentrated γ-glutamyl load from a soybean we tuned for exactly this.

SCALLOP, KOKUMI DASHI

8	large scallops, dry, at room temperature
300 ml	kokumi broth base, as dashi
1 tsp	engineered-soy white miso
1 tbsp	olive or grapeseed oil
chives, lemon	

serves 2 · time 15 min

1. Warm the broth base with the white miso whisked in to a bare steam; hold it warm, never boiling, as the dashi.
2. Pat the scallops bone-dry — dryness is the whole secret of the sear. Oil a heavy pan to shimmering and sear the scallops **90 seconds** on the first side without moving them, until deeply golden, **45 seconds** on the second. They should be just set, barely translucent at the center.
3. Pool the warm dashi in shallow bowls, sit the scallops in it, finish with chives and a few drops of lemon.

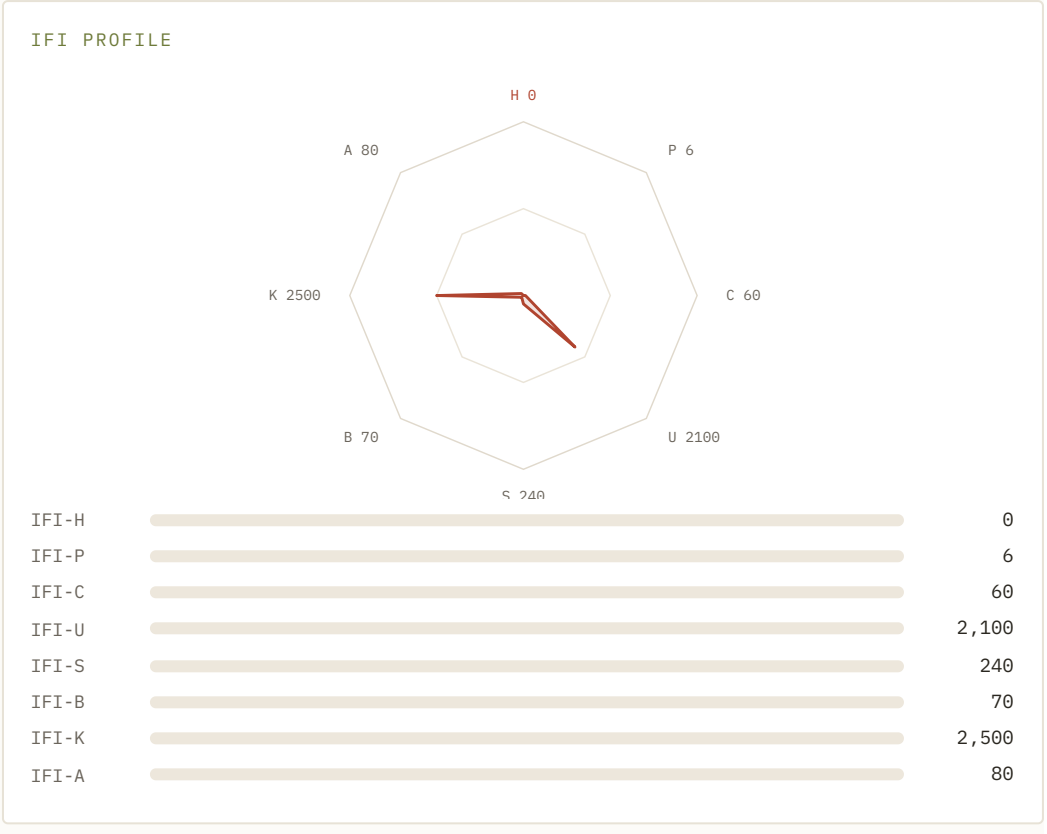

Three sources, one reading: K 2,500, 0.80 mM-equivalent. The seared crust is Maillard kokumi, the scallop's flesh is its own glutathione, the dashi is engineered γ-glutamyl — and because all three speak to the same allosteric site, the plate's K is not their sum but their saturating total, the receptor near full and holding. The

lemon goes on at the very end and in drops: its acid would, with time, begin to hydrolyze the γ-glutamyl bonds, and the point of a kokumi dish is to get those bonds to the tongue intact.

Dark chocolate, salt, and kokumi

We close where the axis is least expected, and where honesty matters most — a dessert that shows exactly what kokumi does and, just as clearly, what it does not.



Cacao at seventy percent carries real bitterness and real astringency. Kokumi rounds and lengthens them without removing a molecule — a different receptor, and an honest edge to the whole axis.

Dark chocolate at seventy percent carries real bitterness — theobromine and caffeine firing the TAS2R bitter receptors — and real astringency — catechins and procyanidins, the cacao polyphenols, binding and precipitating the saliva's proline-rich proteins so the mouth feels drawn and dry. That astringency is the A 620, the highest on the card. Kokumi erases neither, and it is important to say so plainly: the CaSR is not the bitter receptor, and nothing in this cream removes a molecule of theobromine. What the engineered vanilla and the trace of white miso do is lengthen and round — they add body and continuity and potentiate the sweetness, so the 70% reads fuller and softer and finishes longer than its cacao percentage should allow.

```
KOKUMI DARK CHOCOLATE CREAM
──────────────────────────────────
200 g          dark chocolate (70%)
250 ml         barista oat milk
1              engineered vanilla pod, scraped
pinch          flaky salt
1 tsp          engineered-soy white miso
──────────────────────────────────
serves 4   ·   time 15 min + chill
```

1. Warm the oat milk with the vanilla and salt to a steam; steep 10 minutes off the heat.
2. Pour the warm milk over the chopped chocolate, wait a minute, then whisk smooth. Whisk in the miso last — a trace, undetectable as soy, present only as depth — its kokumi rounding the chocolate's edge and lengthening its finish.
3. Chill until set. The miso and vanilla kokumi make the 70% read softer and longer than its cacao percentage should allow; the salt sharpens the whole.

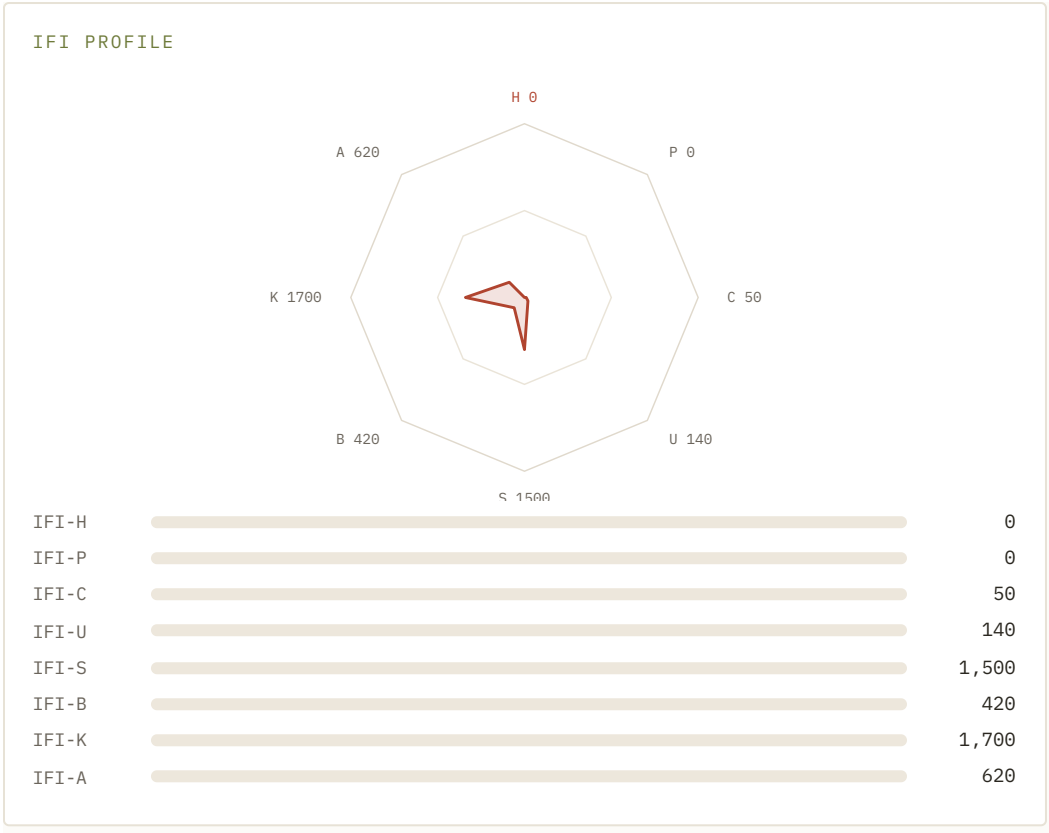

The flaky salt does separate, real work beside the kokumi: sodium is a direct suppressor of bitter perception, so a pinch genuinely quiets the theobromine's edge, and the oat-milk fat coats the tongue and blunts the astringent pull physically. The cream eats rounder than its numbers look — K 1,700 of continuity laid over a bitterness and an astringency we did not pretend away but dressed. That distinction, rounding versus removing, one receptor versus another, is the honest edge of the whole axis: kokumi lengthens and fills, and that is a great deal, but it is not a bitterness eraser. Erasing bitterness is a different lever on a different receptor, and the card tells you the truth about which one you pulled.

What the eight numbers tell us

Set the five dishes side by side with the thing we are proudest of, and the chapter makes its argument in a single glance down one column.

DISH	H	P	C	U	S	B	K	A
Nine-hour onion	0	13	40	1,100	900	80	2,300	90
Shiitake–barley	10	19	30	2,400	180	120	2,600	160
White beans	10	19	30	2,300	240	110	2,400	120
Scallop, kokumi dashi	0	6	60	2,100	240	70	2,500	80
Dark chocolate cream	0	0	50	140	1,500	420	1,700	620
Phoenix · SS-0096 (benchmark)	4,200	252	800	1,800	1,100	2,200	1,900	1,400

Read down the kokumi column. Every plate here lives at K 1,700 and up, and four of the five out-score the richest single thing in our catalog. Phoenix, SS-0096 — the full-spectrum flagship, the one cultivar engineered to register on all eight axes at once, the superhot built on two AI-found Pun1 edits with our model predicting approximately 2.95 million Scoville units — reads K 1,900. A pot of white beans beats it.

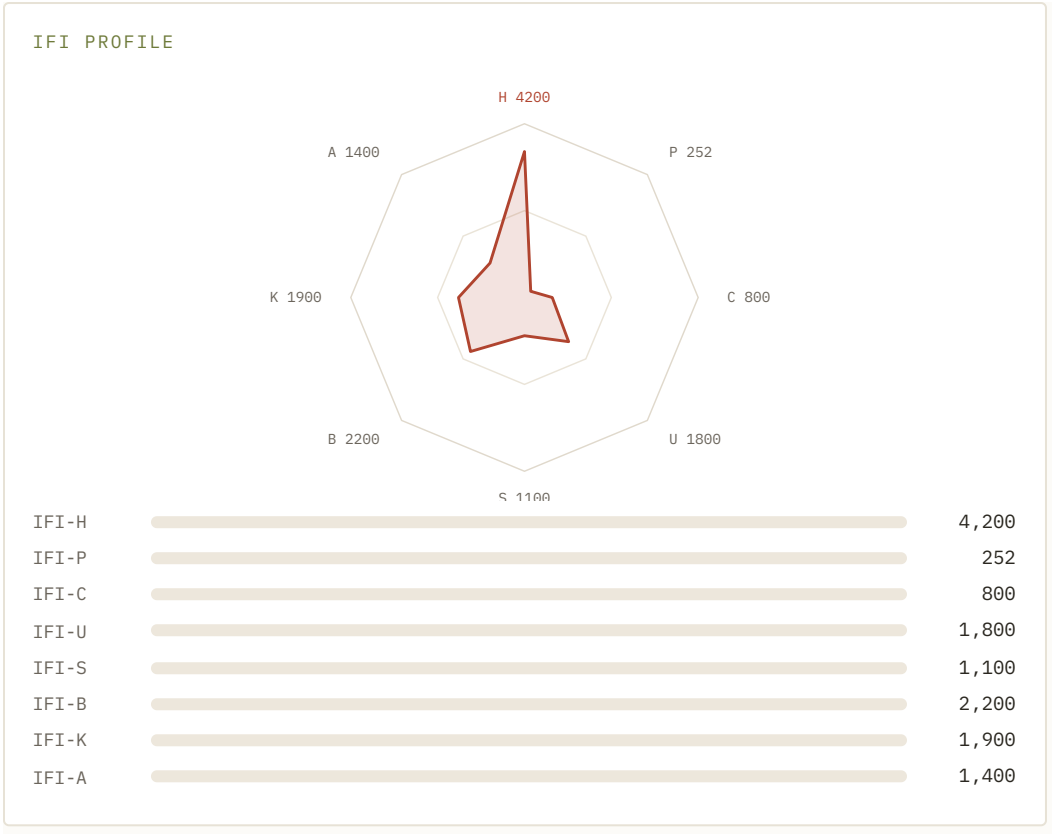

That is the chapter's quiet joke and its real thesis at once. Kokumi is not the hard axis. It is the *available* one — the richness you can lay into a plate of onions or a bowl of beans or a cup of chocolate out of plants and patience and a few designed enzymes, grown to carry it on purpose and held to the platform's design specification so the same plate comes back the same.

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

One honest word to close on. A number on a card is a measurement of a stimulus, not a verdict on a dinner. The K we print is the receptor-referenced loudness of a dish's γ-glutamyl fraction, read against a fixed reference and held to the platform's seven-percent window; it is not a promise about your particular evening, your own tongue's CaSR complement, or the thousand matrix effects — fat, salt, temperature, what you ate an hour before — that shift how loud a thing lands in a real mouth. Kokumi is a modulator, not a flavor. It needs the rest of the plate to act on, and it gives its best in company. That is why this chapter is dishes and not compounds. The receptor named for calcium does its quiet work only in the full, crowded context of a meal — which is, after all, the only place any of this was ever meant for.



PART V OF V

Precision & Nourishment

reproducible richness · the honest label · why richness is care

Chapter 16 · Reproducible richness — *CV ≤ 7%, GRAS, and the honest label.*

A flavor you can design but not reproduce is a party trick, not an ingredient — and for the processed-food reader, who must set the same product in front of the same shopper in forty markets for ten years running, reproducibility is not one feature of the kokumi program but the whole of it. This is the least lyrical work our family does and, for anyone who makes food at scale, the most load-bearing: the unglamorous discipline that makes every other number in this book mean something, and the regulatory and labeling honesty that keeps the program worthy of being believed.

Richness is the hardest of the eight axes to hold still, and it is worth being clear about why before we are clear about how. Heat is a single molecule against a single channel; two lots that assay to the same capsaicin taste the same, full stop, because heat is heat. Kokumi is not a molecule but a chord — a family of small peptides struck together, felt not as a flavor of its own but as the fullness, the length, and the rounded continuity it lends everything around it. A chord can be played at the same volume and still be the wrong chord. Holding richness steady, then, means holding not a quantity but a shape, and that is a stricter promise than any single-compound axis is ever asked to make.

The molecule we are holding still

The receptor is CaSR, the calcium-sensing receptor — a class C G-protein-coupled receptor that works as a linked pair, each half carrying an enormous clamshell of protein, the Venus-flytrap domain, hinged around a cleft. Its first language is calcium: Ca^{2+} ions settle into the open cleft and coax it closed, and that closing — relayed down through a cysteine-rich collar into the seven-pass membrane bundle, which couples to Gq and releases a pulse of signal inside the cell — is the message the tongue reads as presence. Aromatic L-amino acids lean on the same hinge as gentle co-agonists. The kokumi peptides are not the key that opens the lock; they are a hand laid over the closing, a positive allosteric modulator that makes the cleft shut more readily and stay shut longer, so that the same ambient calcium rings a louder, rounder bell. That is the whole of the sensation in one line: richness is the receptor made more willing.

What makes a peptide a kokumi peptide is a single, unforgiving detail of its bond. These are *γ*-glutamyl peptides — glutamate joined to its neighbor not through the usual α-carboxyl that strings ordinary protein together, but through the side-chain γ-carboxyl, an isopeptide linkage the receptor treats as an entirely different word. The α-isomer of the very same amino acids is silent at CaSR. Glutathione (γ-Glu-Cys-Gly), the most abundant of them and present

in nearly every living cell we eat, is the quiet reference tone; γ-Glu-Val-Gly is the loud one, an order of magnitude more potent; γ-Glu-Cys and γ-Glu-Val sit in between. A pure kokumi substance is a strange thing to meet on the Index, because it lights a single axis and leaves the other seven dark:

IFI PROFILE

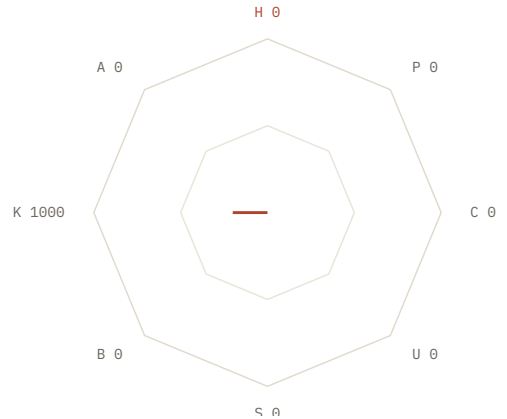

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	0
IFI-S	0
IFI-B	0
IFI-K	1,000
IFI-A	0

That profile is glutathione — no heat, no sweetness, no umami of its own, only richness — and its silence everywhere but the kokumi axis is exactly why we chose it as the reference tone against which every richer blend is read.

Two ATP-dependent ligases build these peptides, and they are where our edits live. Glutamate-cysteine ligase joins glutamate to cysteine to make γ-glutamylcysteine; glutathione synthetase then adds glycine to finish glutathione. Run against the free amino-acid pool, the first ligase is promiscuous enough to lay down γ-Glu-Val as well, and with glycine, γ-Glu-Val-Gly. Traditional food leaves the mix of products to chance — to whichever amino acids a particular clove or a particular ferment happened to have on hand. We do not. The kokumi load is set by defined genetic edits to those two ligases and by the amino-acid pool we feed them — their expression, their substrate preference — so that *which* peptides form, and in *what ratio*, is a designed outcome rather than the accident of a season.



A clove in cross-section — the storage peptides that read as richness are laid down here, cell by cell; in a traditional allium their density is left to cultivar, soil, water, and the accidents of a single harvest.

The receptor answers calcium along an S-shaped curve, the classic saturating response of a binding site that fills and then has no more room: written out, the response follows R = R_max · [Ca]^n / (EC_{50}^n + [Ca]^n), with a Hill coefficient *n* near 3 — the signature of the dimer's cooperative gating, where one half closing helps the other close. A kokumi peptide does not raise the top of that curve; it slides the whole curve leftward, lowering the apparent EC_{50} for calcium by a cooperativity factor α. Formally the apparent half-point moves as EC_{50} × (1 + [K]/K_B) / (1 + α[K]/K_B), which at saturating peptide collapses to EC_{50} / α. With α near 5, the receptor reaches the same half-signal at a fifth of the calcium — and the consequence at a fixed dietary level is dramatic. Suppose the receptor's calcium EC_{50} sits near 3 mM and a food presents calcium near 1.5 mM. Bare, the response is 1.5^3 / (3^3 + 1.5^3) = 3.375 / 30.375 ≈ 11% of maximum; add saturating kokumi peptide so the apparent EC_{50} falls to 0.6 mM, and the same 1.5 mM now drives 3.375 / (0.6^3 + 1.5^3) = 3.375 / 3.591 ≈ 94%. Nothing about the calcium changed; the receptor was merely made willing. That single factor is the molecular reason a kokumi-dosed food can taste full at lower salt, lower sugar, and less fat — the receptor doing more with less, which is the quiet thesis under all of this work.

The engineering behind the fingerprint

Designing a fingerprint is harder than designing a magnitude, because the two ligases do not read a dial marked *ratio*. What they read is their own sequence and the pool of amino acids around them, and so the fingerprint is set by the same computational pipeline that built our heat. A protein language model scores every candidate substitution in each ligase by its predicted effect on catalytic efficiency and fold stability, ranking the handful of edits that would nudge a promiscuous enzyme toward valine over cysteine, or toward finishing γ-Glu-Val with glycine rather than leaving it short. A genome-scale flux-balance model then confirms what is worth editing at all — that the ligase step, and not the upstream supply of glutamate or the downstream transpeptidase, is the true rate limiter — so that the edits we write in with CRISPR move the number we meant to move and leave the rest of the cell's economy intact. The amino-acid pool is the second lever: feed the culture more free valine and the fingerprint tilts toward γ-Glu-Val and γ-Glu-Val-Gly, the louder tones; starve it and the mix drifts back toward glutathione. Holding a fingerprint steady is therefore holding two things at once — a genotype and a growth recipe — and it is the deepest reason the controlled chamber is not a luxury but a term of the specification.

Reading loudness as a number

Every axis in this Library resolves to one number, and the number is loudness, not amount:

> $IFI(c) = 1{,}000 \cdot EC_{50}(\text{ref}) / EC_{50}(c)$

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response; a smaller EC_{50} is a louder stimulus, and inverting it puts the loud thing on top. On the kokumi axis we anchor

the reference to glutathione — the tone a person already meets in a bowl of ordinary soup — and fix its EC_{50} at 1,000 µM by convention, so that glutathione reads exactly 1,000 and the entire axis becomes a multiple of a richness everyone already knows. The formula then collapses to a line you can read off by hand: IFI-K = 1,000,000 / EC_{50}, with EC_{50} in micromolar.

On our cell-based CaSR assay, calibrated to that anchor, the common γ-glutamyl peptides fall into a clean ladder. The conversion from receptor concentration to mass is a single multiplication — micromolar times molecular weight over a thousand gives milligrams per litre — so the same table can be read by a pharmacologist and by a formulator weighing out a batch.

PEPTIDE	EC_{50} (MM)	EC_{50} (MG/L)	RELATIVE POTENCY	INTRINSIC IFI-K
γ-Glu-Val-Gly	80	24.3	12.5×	12,500
γ-Glu-Val	250	61.6	4.0×	4,000
γ-Glu-Cys	500	125.1	2.0×	2,000
Glutathione · reference	1,000	307.3	1.0×	1,000

The worked conversion is worth doing once in the open. γ-Glu-Val-Gly, molecular weight 303.3, reaches its half-signal at 80 µM × 303.3 / 1,000 ≈ 24.3 mg/L; glutathione, molecular weight 307.3, needs 307.3 mg/L to do the same work on the same receptor. Twelve and a half times less mass for an identical effect — the arithmetic behind a claim we make carefully, that a precisely dosed real engineered depth can stand in for a long line of synthetic enhancers and the salt they were there to carry.

Why the same total can taste different

Here is the trap that makes kokumi stricter than heat. Two lots can carry the identical total mass of γ-glutamyl peptide and taste unmistakably different, because the sensation is set by *which* peptides in *what* proportion, not by the sum. A lot weighted toward γ-Glu-Val-Gly rings hard and fast and front-loaded; a lot weighted toward glutathione rings gentler, slower, and longer. The totals match while the experiences diverge.

The math shows exactly how far they diverge. In the sub-saturating range where a food actually works — well below the dose that shuts the clamshell — potencies add in proportion, so a blend's effective loudness is the mole-fraction-weighted sum of its parts' relative potencies, $r_{blend} = \Sigma\, f_i\, r_i$. Take two lots at one identical total concentration (and because γ-Glu-Val-Gly and glutathione differ in mass by barely one percent, their mass fraction and mole fraction are the same number to three figures, so we need not split the hair).

LOT	Γ-GLU-VAL-GLY	GLUTATHIONE	EFFECTIVE POTENCY R	EFFECTIVE IFI-K	TIME COURSE
A	70%	30%	9.05	9,050	fast, front-loaded
B	30%	70%	4.45	4,450	slow, lingering

A lot that is 70% the loud peptide carries an effective potency of 0.70 × 12.5 + 0.30 × 1.0 = 9.05; invert the proportions and it falls to 0.30 × 12.5 + 0.70 × 1.0 = 4.45. The same mass of "kokumi" is, by the receptor's reckoning, 9.05 / 4.45 ≈ 2.03 times as loud in the first lot as in the second — and loud in a different shape, a bright front

rather than a long tail. This is why our release specification controls the *fingerprint* — the full chromatographic profile of which peptides, in which ratios — and not merely the magnitude. A lot that carries the right total in the wrong proportions is held, reworked, or rejected, the same as a lot that is simply weak. The promise is not *this much kokumi*; it is *this kokumi*, this shape, this ring, every time.

The variance problem, quantified

A field of traditional garlic, or a crock of traditional miso, carries enormous batch-to-batch swing in its kokumi load. The storage-peptide density of a clove depends on cultivar, soil, water, and the accidents of a harvest; the γ-glutamyl fingerprint of a ferment depends on the wild organisms that happened to colonize it and the temperature history of a particular year. For a cook, that variation is the romance. For a manufacturer who has promised a shopper the same soup every time, it is a liability — and it is why most industrial kokumi has historically come from the single most controllable source on the shelf, purified yeast extract, dosed by weight. We attack the variance at its two roots: the genetics, fixed by the defined edits rather than left to an organism's defaults, and the growing, done in controlled environment rather than open field.



The controlled chamber — light, water, temperature, and nutrient held to a recipe, so that the only thing allowed to vary from one lot to the next is as near to nothing as we can make it.

Batch-to-batch variation is reported as a coefficient of variation, CV = σ / μ: the standard deviation of a measurement across lots divided by its mean, in percent. Across the Library we hold CV at or below 7%, against the roughly 42% of an uncontrolled crop — and the gap matters far more in its consequences than in its digits. Take a lot specified to a kokumi score of 1,900, the richness our flagship full-spectrum cultivar is designed to carry. At 7% CV the standard deviation is 0.07 × 1,900 ≈ 133, and roughly nineteen lots in twenty land inside the 95% window of 1,900 ± 2σ = 1,634 to 2,166. At 42% CV the standard deviation is 798, and that same window sprawls from 304 to 3,496 — the lot might arrive at a sixth of its target or at nearly double it.

QUANTITY	ENGINEERED · CONTROLLED	UNCONTROLLED CROP
Batch CV	≤ 7%	~42%
σ at K = 1,900	≤ 133	~798
±2σ window (95%)	1,634 – 2,166	304 – 3,496
Capability Cp (±28% tol.)	1.33	0.22
Lots within spec	> 99%	~50%

Formulators speak of this as process capability, Cp — the ratio of the tolerance a customer allows to the spread a process delivers, Cp = tolerance / (3 · CV). Against a typical flavor-release tolerance of ±28%, our 7% CV gives Cp = 0.28 / 0.21 = 1.33, the standard bar for a capable process; the uncontrolled crop gives Cp = 0.28 / 1.26 = 0.22, with roughly half of all lots falling outside the window. And we are honest about where the 7% itself comes from, because any measured variation is the true batch variation and the measurement's own noise added in quadrature: $\sigma^2_{total} = \sigma^2_{process} + \sigma^2_{assay}$. If the instrument repeats at 4% and the lots read at 7% in total, the biology itself is varying at $\sqrt{(7^2 - 4^2)} = \sqrt{33} \approx 5.7\%$, and the rest is us. We drive the measurement term down the oldest honest way, by replication: the uncertainty on a lot's mean falls with the square root of the number of reads, so a 4% assay run in triplicate reports the lot mean to $4\% / \sqrt{3} \approx 2.3\%$. The number we publish is the number left after we have subtracted our own trembling hand.

How a lot is verified

Every lot is read on both of our instruments before it is allowed out the door: the cell-based CaSR assay for the effective EC_{50} — the magnitude — and the peptide chromatograph for the fingerprint — the shape. Because the active bond is the γ-linkage and its silent α-isomer travels alongside it, the chromatography must resolve the two; a method that merely totalled "γ-glutamyl peptide" would wave through a lot full of inactive isomer as though it were rich. A lot is released only when both reads fall inside the specified window: the EC_{50} within its band, and the fingerprint matched peak for peak against the reference profile.



The microscope stage, where a lot's fingerprint is read and matched before anything ships — the quiet gate that turns a flavor into an ingredient.

Our flagship full-spectrum cultivar, which must land all eight axes at once, is the hardest case and the clearest illustration. Its kokumi is designed to read 1,900 — standard deviation 133, a 95% release

window of 1,634 to 2,166 — while seven other axes hold their own windows at the same time, and we steer to a tighter internal control line at ±1σ, 1,767 to 2,033, so the process is pulled back long before any lot can approach the release edge:

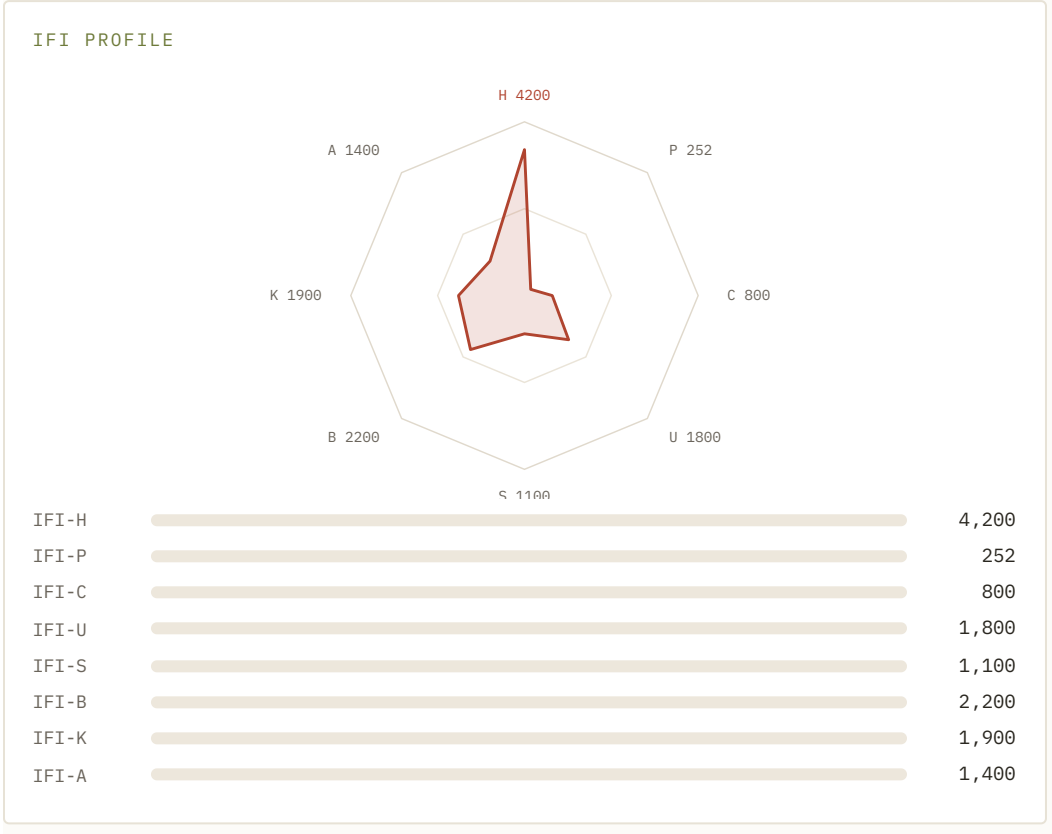

A lot that carries the right total in the wrong proportions — too much glutathione, not enough γ-Glu-Val-Gly, the chord struck flat — is not shipped and relabeled; it is held, reworked where the biology allows, or rejected outright, exactly as a lot that is simply weak. There is no "close enough" on the fingerprint, because close enough is precisely how a product drifts, one forgivable lot at a time, until the soup a shopper trusted is no longer that soup.

The release decision, as statistics

A release is a statistical decision, not a single glance at a meter, and we treat it as one. Each lot is read in replicate — the triplicate that pins its mean to about 2.3% — and the mean is plotted on a control chart whose limits sit *inside* the customer's specification rather than on it: a warning line at ±2σ and an action line at ±3σ around the 1,900 target, with the specification itself further out still. A lot inside the warning band ships. A lot past the warning line but inside action is investigated before it ships. A lot past the action line is held, no exceptions. The chart also watches for drift that no single point would flag — several lots in a row trending the same direction, or a run of them hugging one side of the target — because a process walking slowly toward a limit is a process to correct *now*, while every lot is still good, rather than after the first one fails. The discipline is deliberately pessimistic: we would rather hold a sound lot and lose a day than ship a drifting one and spend a customer's trust in the number.

How a fingerprint goes wrong

Richness fails in specific, chemical ways, and naming them is how we catch them. The commonest is oxidation: the cysteine-bearing peptides, glutathione and γ-Glu-Cys, carry a free thiol that will pair into a disulfide — glutathione into its oxidized dimer, GSSG — under

warmth, air, or a trace of metal, and the oxidized form reads differently on the receptor and lands in a new place on the chromatogram. A lot can hold its total and lose its shape to the air alone. The second is isomerization: the γ-linkage that makes a kokumi peptide active is not unbreakable, and under enough heat or the wrong pH it can migrate toward the silent α-form, so that a lot which assays rich by total nitrogen is quietly going deaf at the receptor. The third is the pathway's own exuberance: γ-glutamyltransferase, left too active, keeps shuttling γ-glutamyl groups from one acceptor to the next, scrambling a designed ratio into a smear. Every one of these moves the fingerprint without touching the total — which is exactly why a gate that read only the total would wave all three through, and why ours reads the shape, resolves the γ-isomer from the α, and treats a peak in the wrong place as a failure even when the sum is perfect.

What the number does not promise

We are precise about what ≤7% CV does and does not buy, because a number oversold is worse than no number at all. It is a promise about the ingredient as we ship it, not about the food it lands in. Kokumi is a modulator, felt only through the salt, sugar, fat, and umami it rounds out, so the same dose that reads 1,900 on our instruments will present differently in a thin broth than in a heavy sauce — the matrix has a vote, and the formulator holds it. It is an instrument reading, and the instrument is a receptor in a dish, not a mouth: it captures magnitude and fingerprint faithfully, and captures adaptation, saliva, and serving temperature not at all, which is why every release is also tasted by people, and why a human panel can hold a lot the instruments would pass. And it is a reproducibility promise, not a prescription — holding a lot to 1,900 says only that it will be the same next year, not that 1,900 is the right richness for a given soup; that remains the formulator's call, and our work is merely to make it a call they can make once and then rely on. The ceiling over all of it is the receptor itself: kokumi saturates, a loud blend is only so much louder than a quiet one before the clamshell is simply shut, and past that point more peptide buys nothing but cost and a longer ingredient line.

The regulatory line, held firmly

The kokumi substances are not exotic. Glutathione and the γ-glutamyl peptides are present across the ordinary diet — in yeast, in garlic, in cheese, in soy and meat — and the specific flavor ingredient γ-glutamyl-valyl-glycine carries flavor-industry GRAS recognition, a FEMA GRAS determination, for use as a flavoring. We work inside those frameworks, and we describe our ingredients' function in the language regulators use for flavor: kokumi substances enhance *mouthfulness, continuity, and complexity*, and support flavor at reduced salt, sugar, and fat. We do not make a **health claim**, and a reader of this Library will never catch us at one. Glutathione carries a large and genuine scientific literature as the cell's own antioxidant — and that literature is not ours to borrow for a sentence about eating our ingredient. We say nothing about detox, nothing about antioxidant benefit from the diet, nothing about any outcome in the body. The reduced-sodium attribute is a *flavor*

attribute — kokumi lets a food taste full at lower salt — and whatever health meaning a formulator or a regulator chooses to attach to that lower salt is their determination, and not our promise.

The label, in the clear

A word on the label itself, because transparency is part of the ethos and the whole of the trust. The kokumi ingredients name plainly: an engineered yeast extract is a yeast extract, an engineered miso is miso, a kokumi garlic is garlic. We decline the "natural flavor" catch-all wherever we can, because a cook and an eater deserve to know what is making their food rich — and the engineered organisms that carry this axis are ordinary foods first. Garlic holds 560 constructs across kokumi and pungency; soybean, 360 across umami and kokumi; yeast, 240 across umami and kokumi; onion and vanilla carry the axis too, each in its own register. Where one of them brings a common allergen — soy, above all, in the miso and the fermented-soy depth — it is declared without exception and in the clear, because a richness that sends someone to the hospital is no richness at all. The clean-label reader, who came to this axis wanting fewer unpronounceable additives rather than more, should find in the kokumi program an ally: the whole method reaches a target satisfaction with a *shorter* ingredient list, not a longer one — a spoonful of real fermented depth standing in for a row of synthetic potentiators and the sodium they were quietly propping up. The arithmetic is plain and checkable: a formulator who trims a soup's added salt from 1.00 to 0.75 grams per 100 and restores the lost sense of fullness with a kokumi dose has cut that line's sodium by a quarter without the eater reading the bowl as thinner — a flavor result, with whatever health meaning the lower sodium carries left entirely to them.

This discipline is not timidity; it is the condition of being believed. A family that will design a flavor down to the receptor and then refuse to oversell what it does is a family whose numbers you can bank. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The ≤7% CV and the no-claims line are a single commitment read two ways: *we will tell you exactly what is true, hold it lot after lot, and say nothing more.*

Chapter 17 · The quiet axis — *why richness is a kind of care.*

Some flavors announce themselves. Heat makes you gasp, sweetness makes you reach, sourness makes you wince — each arrives at the door of the tongue and knocks. Kokumi does none of that. You can taste a broth all evening and never once point to the thing that makes it feel *finished*, because the thing that makes it feel finished has no flavor of its own. It only makes everything else more whole: rounder, more continuous, more present across the whole mouth and slower to leave it. It is the quietest of the eight axes, and we have come to believe that the quietness is not incidental but the deepest thing about it — the clearest expression, in a whole library of engineered flavor, of why our family does any of this at all.

This is the spoonful that *stays with you*. Not the one that flares and fades, but the one that settles and reads — without a single word, without any flavor you could name — as *you are cared for*. For most of human history that sensation was the signature of patience: the long-simmered stock, the aged cheese, the sauce reduced across an afternoon, the paste fermented across a year. Someone took time over the food because they cared, and the body read the time as care, because the molecules that time deposits were really there in the dish. This chapter is about those molecules — what they are, which receptor reads them, how we learned to put an honest number on a sensation that announces nothing, and how we taught a plant to carry the gift of patience without the afternoon.



Dawn in the greenhouse — the family walk where the question that became this axis was first asked, and where it still gets asked every morning.

A receptor the body uses to stay alive

Start with the lock, because the key only means something once you know what it opens. The receptor that reads kokumi is the calcium-sensing receptor, CaSR — and it is not, first and foremost, a taste receptor at all. It is one of the oldest housekeeping sensors in the body, a class C G-protein-coupled receptor, structural cousin to the metabotropic glutamate receptor that reads savory depth on the umami axis. Where most receptors are a single chain, CaSR works only as an obligate pair: two identical protomers clasped together, each carrying a large bilobed extracellular domain shaped — the biologists named it exactly — like a Venus flytrap. A flexible cysteine-rich neck ties that flytrap to the seven-helix bundle that crosses the cell membrane and talks to the machinery inside.

Its day job is arithmetic on your own blood. The chief cells of the parathyroid glands wear CaSR like a thermostat for calcium. When blood calcium rises, the ions settle into the flytrap's cleft, the two lobes close, the closure torques the dimer, and the receptor fires through Gq/11 to phospholipase C, lifting intracellular calcium and telling the gland to *stop* releasing parathyroid hormone. When blood calcium falls, the flytraps open, the hormone flows, the bones and kidneys answer. This loop runs every second of your life, silently, and it is the literal thing meant by *a number anchored to the half-activation of a receptor your parathyroid glands use to keep you alive*. The axis is calibrated against a sensor the body cannot do without.

The tongue borrows that sensor. The same CaSR sits on oral tissue, and here a second class of molecules speaks to it — not calcium directly, but small γ-glutamyl peptides, glutathione chief among them, which bind the flytrap not as agonists that open it on their own but as *positive allosteric modulators*. They do not ring the bell; they make the clapper lighter. With a γ-glutamyl peptide bound, the flytrap closes at a lower calcium concentration than it otherwise would — the receptor is sensitized, tuned, made more willing. In the mouth, that sensitization is not perceived as a taste, because CaSR carries no dedicated taste line of its own. It is perceived as *more*: more continuity, more thickness, more mouthfulness, a longer and rounder persistence of every flavor already present. That is kokumi. It makes everything else more whole because, mechanically, that is exactly and only what it does.

The molecule and its forged bond

Now the key. Glutathione is a tripeptide — γ-L-glutamyl-L-cysteinyl-glycine — and the whole of its power on this axis hides in that lowercase γ. An ordinary peptide bond links one amino acid's α-carboxyl to the next amino acid's amino group; it is the bond every protein in you is strung from, and the mouth and gut are full of peptidases evolved to snip it. Glutathione does not use that bond between its first two residues. It links glutamate's *side-chain* carboxyl — the γ-carboxyl, one carbon further out — to cysteine's amino group. That single displacement forges a bond the ordinary peptidases cannot read. The γ-glutamyl linkage is invisible to them; only one specialist enzyme, γ-glutamyltransferase, can cut it. A molecule that survives the mouth intact is a molecule that reaches the receptor intact, and the γ-bond is why the signal of patience is a signal at all rather than a casualty of the first enzyme it meets.



Garlic in cross-section — one of the kokumi larders, where γ-glutamyl peptides accumulate as the plant's own stored patience.

That γ-glutamyl handle is the pharmacophore — the part the flytrap actually reads. A whole family of peptides sharing the γ-Glu motif rings CaSR, some of them more loudly than glutathione itself, which is why aged, fermented, and long-cooked foods, where these peptides accumulate, taste so complete. Garlic is a canonical larder of them; so are yeast, soybean, onion, and the cured and the aged. The plant builds glutathione in two deliberate, ATP-spending steps, and both matter to anyone who wants to make more of it. First, glutamate-cysteine ligase joins glutamate to cysteine and forges the γ-bond — the committed, rate-setting step. Second, glutathione synthetase adds the glycine. The first enzyme is the one that decides how much richness a cell will carry, and, as every well-regulated pathway does, it is governed by a brake.

Relieving the brake

The brake is feedback inhibition. Glutathione, the finished product, binds back onto glutamate-cysteine ligase and tells it to slow down — a sensible economy for a plant minding its own metabolism, and precisely the ceiling that ordinary cultivation runs into. You cannot simply grow a vegetable into deep savory richness, because the cell stops itself long before the concentration a cook would want. Centuries of kitchens solved this from the outside, with time: simmering, aging, and fermenting are all ways of letting enzymes and microbes pile up γ-glutamyl peptides past what a living cell would allow itself. The afternoon of stock is the cook doing, slowly and by hand, what the plant's own brake forbids.

We moved the solution inside the plant. The platform scores every candidate single-residue change to the ligase with a protein language model that predicts, for each substitution, its effect on catalytic efficiency and on fold stability; a genome-scale flux-balance model then confirms that this enzyme, and not some step upstream, is the true limiter of flux into the γ-glutamyl pool. Four edits emerge that no breeder could have found by selecting a phenotype and hoping the genotype followed — they loosen the product's grip on the regulatory site and shore up the catalytic interface, relieving the brake without breaking the enzyme. Written in with CRISPR, they are designed to let the cell keep building richness past the point where the wild-type plant would have quit. The patience that used to take an afternoon or a year is designed to be constitutive: the plant carrying it the way it carries its color. And

because the edits are precise, the output is reproducible — batch-to-batch variation in the engineered line holds at a coefficient of variation of **≤7%**, against **42%** in the wild-type plant whose brake is still intact and whose richness therefore swings with every season. Reproducibility is not a luxury here. It is the difference between a number you can design against and a number that means nothing twice.

Putting a number on richness

A sensation with no flavor is the hardest thing in the library to measure, and the measurement is the whole reason the axis exists. The Index does for kokumi what it does for all eight axes: it reports not *how much* of a compound a food holds but *how loudly* that compound speaks to its receptor, as a ratio against a fixed reference.

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — low EC_{50} means a loud molecule, since a little goes a long way. We anchor the kokumi axis to glutathione itself, the molecule at the center of everything on this page. Under the assay's standardized calcium background, we fix glutathione's half-activation of CaSR at **EC_{50}(reference) = 1,000 µmol/L** (1.00 mM), which by the formula makes **IFI-K(glutathione) = 1,000** by convention — the scale's zero-point of loudness. Glutathione's molar mass is 307.32 g/mol, so that anchor is also 1,000 µmol/L × 307.32 µg/µmol ≈ **307 mg/L**, about 307 parts per million — a useful bridge, because a formulator weighs grams, not moles.

From there the axis reads any food by inverting its loudness back to an EC_{50}. The richest cultivar our family has designed, Phoenix (SS-0096), carries a canonical kokumi score of **IFI-K = 1,900**. Running the formula backward:

EC_{50}(Phoenix) = 1,000 × EC_{50}(reference) ÷ IFI-K = 1,000 × 1,000 µmol/L ÷ 1,900 = **526.3 µmol/L**.

In mass terms that is 526.3 µmol/L × 307.32 µg/µmol ≈ **161.7 mg/L**. The reading is exact about what it claims: in our model, Phoenix's γ-glutamyl pool is designed to reach CaSR half-activation at a little over half the concentration glutathione needs, so it is **1.90×** as loud (1,900 ÷ 1,000), speaking the richness with roughly half the material.

READING	IFI-K	EC_{50} (MMOL/L)	EC_{50} (MG/L)	LOUDNESS
Reference · glutathione (anchor)	1,000	1,000.0	307.3	1.00×
Phoenix · SS-0096	1,900	526.3	161.7	1.90×
Scorecard normalization cap	5,000	200.0	61.5	5.00×

The score is only a point on a curve, and the curve is worth seeing whole, because it tells you how richness grows with dose. CaSR's response to its modulated activation follows the Hill relation $\theta = (C/EC_{50})^n \div (1 + (C/EC_{50})^n)$, where θ is the fraction of maximum response and n the slope. Take the simplest case, n = 1: at half the EC_{50}, response is 0.333; at the EC_{50}, 0.500 by definition; at twice it, 0.667; at four times, 0.800. But CaSR is a dimer whose two flytraps

do not act wholly alone, and we fit its kokumi response with a mild positive cooperativity, a Hill slope near **n ≈ 1.4**. That steepens the middle of the curve: at twice the EC_{50} the response climbs to 0.725 rather than 0.667, and at half the EC_{50} it falls to 0.275 rather than 0.333. The practical reading is that near the operating point, a small change in how much γ-glutamyl peptide a dish delivers produces a disproportionate change in perceived fullness — richness, once it arrives, arrives quickly.

Bring it to the table, where the number was always headed. Suppose a broth is formulated to deliver Phoenix's kokumi peptides right at half-activation, 526.3 µmol/L, in a 250 mL bowl. That bowl holds 0.250 L × 526.3 µmol/L = 131.6 µmol of peptide, or 131.6 µmol × 307.32 µg/µmol ≈ **40.4 mg** — about forty milligrams of the molecule of patience, dissolved in a bowl, sitting at the exact concentration where the receptor gives half its answer. The abstraction resolves, finally, to something you could weigh.

Reading Phoenix on the quiet axis

Phoenix is the only cultivar among our ten proof-of-concept superhots that is designed to register on all eight axes at once, and the eighth to light — the last and quietest to switch on as the design scope climbs — is kokumi.

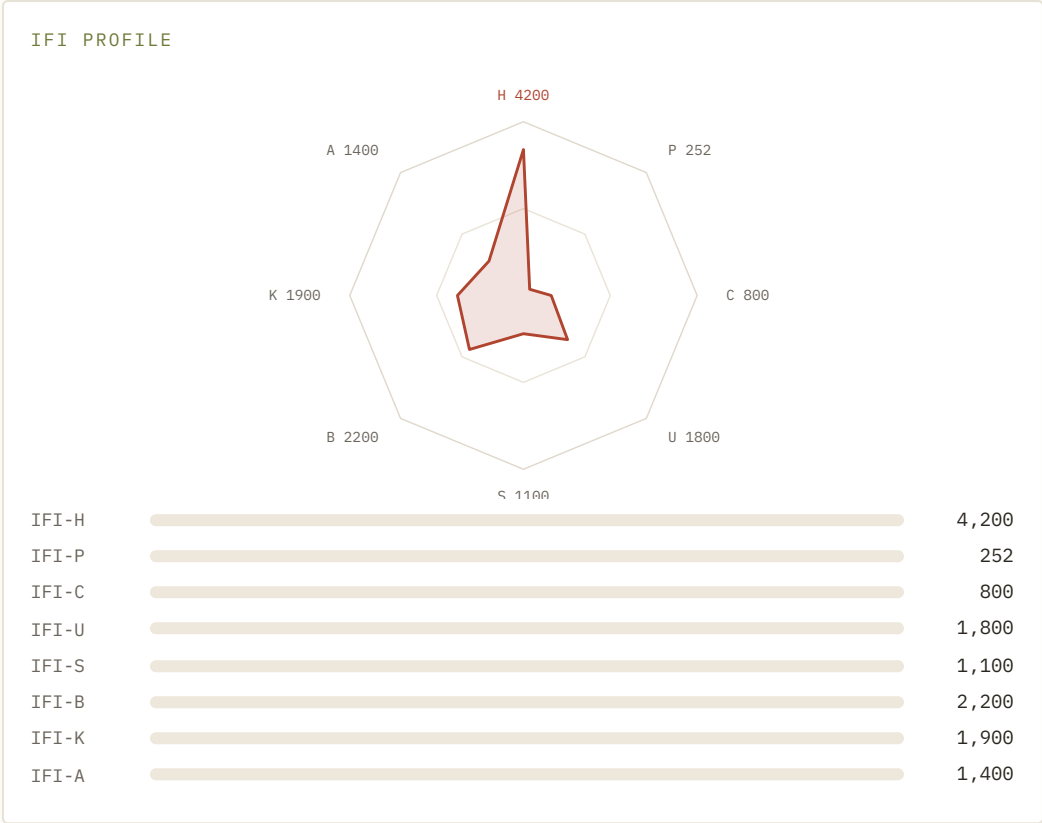



Phoenix (SS-0096) — the first cultivar engineered to register on all eight axes at once, kokumi among them.

It is easy to read that profile as a monument to heat — 4,200 on the heat axis, which our model predicts at approximately 2.95 million Scoville units, a plant designed gene by gene because the pepper that could reach it did not yet exist. But the number that tells you what *kind* of thing Phoenix is sits lower on the card. Its kokumi score of 1,900 means that underneath the fire there is a floor of continuity nearly twice as loud as the reference glutathione — a richness that catches the heat and the umami and the faint sweetness and holds them together so the whole does not shatter into separate sensations. Follow the design scope across the ladder and you can watch the axes accumulate: the earliest cultivars activate two of eight, the middle of the range four, then six, and only at the highest computational design scope do all eight sound together. Richness is the axis that waits for the full composition, because richness is, by its nature, the one that integrates the rest.

CULTIVAR	SKU	PREDICTED SHU	IFI-H	AXES
Ember	SS-0087	866K	1,570	2/8
Blaze	SS-0091	1.72M	2,630	4/8
Firestorm	SS-0094	2.47M	3,490	6/8
Solara	SS-0095	2.72M	3,980	6/8
Phoenix	SS-0096	2.95M	4,200	8/8

A number is only as honest as its error bars, and the ≤7% batch CV design specification earns Phoenix's score the right to be written down. Carry that variation through the formula: because IFI-K scales inversely with the candidate's EC_{50} while the reference is fixed, a 7% variation in the delivered γ-glutamyl titer propagates, to first order, into roughly a 7% variation in the score. For K = 1,900 that is one standard deviation of about ±133, so a single batch reads **1,900 ± 133**. Had we left the wild-type brake in place, the 42% variation would mean ±798 — a reading that could land anywhere from about 1,100 to 2,700, a number with no second meaning. Average four independent batches and the standard error of the mean falls as 1/√4: ±133 ÷ 2 ≈ **±67**, a 95% confidence interval of roughly [1,770, 2,030]. The precision of the enzyme becomes, directly, the precision of the claim.

The kokumi larder

The method was never about any one plant. The γ-glutamyl pathway runs in every green cell; glutathione is one of the most conserved small molecules in all of biology. So the same four-edit relief of the same brake transfers to any organism with a sequenced genome, and across the library the plants that already keep a larder of these peptides are the ones that carry the axis most naturally.

ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Garlic	560	Kokumi / Pungency	Moderate
Soybean	360	Umami / Kokumi	Frontier
Yeast	240	Umami / Kokumi	Hard
Onion	40	Kokumi	Moderate
Vanilla	40	Kokumi	Hard

Read the table and you read the world's savory foundations: the allium that starts nearly every cuisine, the soybean behind every great ferment, the yeast behind every risen and brewed thing, the onion, and — quietly, at the sweet end — vanilla, whose complexity is as much mouthfeel as aroma. These are the organisms where a cook, for ten thousand years, found depth by waiting. Engineering kokumi into them is not inventing a sensation that was never there; it is making the plant's own stored patience abundant and reproducible. And because the richness that reads as *meaty* is, mechanically, a γ-glutamyl peptide ringing a calcium receptor, a broth built on an engineered allium or soybean can be deeply, roundly satisfying with no meat in it at all — which, for a family that has been vegetarian for a lifetime, is not a trick but a homecoming. The depth of the table we grew up at was always plant-deep. We only learned, finally, to measure it.

Less of what we should have less of

Here is the move that makes the quiet axis matter beyond the pleasure of a good broth, and here is where we are most careful. Kokumi contributes continuity and mouthfulness — the sense that a flavor is complete and sustained. A formulator can use that completeness as a lever: hold the kokumi tone constant, at a set CaSR activation, and the perceived satisfaction of a dish can be matched at less of the ingredients a recipe would otherwise lean on for that satisfaction.

Worked plainly: take a reference broth at 0.80% salt, 8.0 g of sodium chloride per litre. Set its kokumi tone so the γ-glutamyl pool sits near the EC_{50} operating point, and a trained sensory panel matches the *perceived* intensity and completeness at **0.53% salt** — 5.3 g/L. That is a 33% reduction, a third less salt, measured strictly as perceptual equivalence: the panel reports the two broths as equally satisfying. In a 250 mL bowl it is the difference between 2.0 g and about 1.33 g of salt — roughly two-thirds of a gram, chosen out of the recipe by the formulator because the richness gave the palate back what the salt was being used to find. The same lever runs the other way across the axes a dish composes: continuity lets a sauce read as rich at less fat, and roundness lets a custard read as full at less sugar, each one a sensory-equivalence target a formulator can design against.

We state this exactly and we stop exactly there. What the panel reports is perceived intensity; what the formulator does is choose a recipe. We make **no claim** about what any of it does for a body — not for salt, not for sugar, not for fat, not for the molecule itself. The design *value* is real and the design value is all we assert: more satisfaction and more pleasure for less of what a person might wish to use less of, delivered from a cleaner source, by giving the palate back the completeness it was reaching for. That is not a health claim. It is a thing to point the engineering at.

What the number does not say

An honest axis is one that tells you where it goes quiet, and kokumi, being the quietest, has the most to disclose.

The first limit is calcium. Because glutathione and its kin act on CaSR as positive allosteric modulators rather than as agonists in their own right, their apparent EC_{50} depends on the calcium background they are read against. IFI-K is therefore a *standardized-condition* number, defined at the assay's fixed calcium — not a universal constant of the food. Serve the same broth in hard water, or in a calcium-rich matrix, and its effective loudness shifts. The score is true; it is true under stated conditions, and we state them.

The second limit is that kokumi has no standalone taste to measure. You cannot ask a panel "how kokumi is this?" in isolation, because there is nothing to point to; you can only measure its *effect* on a carrier — the continuity it lends a base broth. The axis is intrinsically relational, which makes it the hardest of the eight to score and the one we are most careful not to overstate.

The third is saturation. The Hill curve plateaus: past a certain CaSR occupancy, further peptide buys less and less perceived richness, and the scorecard's normalization cap of 5,000 reflects the knee of that curve. A score far up the scale means less than the raw ratio suggests, because the receptor is already near full. The fourth is the matrix itself: the γ-bond resists ordinary peptidases but not γ-glutamyltransferase, so aggressive processing can quietly degrade the very pool you engineered — the delivered dose is not always the synthesized dose, and a kitchen can give richness back or take it away.

And the last is the oldest caution in the Library: the number is loudness, not liking. A high IFI-K is not automatically *better* — too much continuity reads as heavy, even cloying. The Index measures the stimulus, never the preference. It tells you how loudly a molecule speaks; it will never tell you whether, tonight, at this table, that is the richness you wanted.

The quiet axis

There is a last symmetry here, and it is the one that makes the whole program feel, to us, coherent rather than merely clever. The molecule at the center of this axis — glutathione — is the same molecule the body makes, in every cell, as its master antioxidant and its defense against the ordinary damage of being alive. We make no

claim whatever about eating it; we have been careful, and we remain careful. But there is a quiet rightness in the fact that the compound which gives food its sense of being *complete* is the same compound the cell relies on to stay whole. The tongue, reading kokumi, is reading a molecule the body knows intimately from the inside. That the sensation of richness should be built on the chemistry of cellular integrity is not something we engineered. It is something we found, already true, when we followed the pathway down. We only gave it a number.

That finding points past this chapter. If richness and nourishment share a chemistry — if the molecule that reads as *complete* on the tongue is the same one the cell keeps to stay whole — then kokumi is the first clear bridge between the thing this Library was built to measure, flavor, and the thing that runs quietly under all of it, nourishment. It is where the two roads first meet. We mapped it as a flavor because that is what the tongue reports, but the family that built it has long suspected it was reporting something larger.

And it is right to end back at the table, because that is where this family started and where the axis finally means something. The voice of this book has been *we* throughout, and here at the end it is worth saying what the *we* means. One pair of hands did the computing — the models, the edits, the long nights at the bench — but the question those hands were answering came from a table, from a kitchen, from the particular reverence this family has always had for feeding each other well. When we engineer kokumi into a plant, we are not inventing a sensation. We are making the gift of patience reproducible, and portable, and available to a person who was not given the time to make it the old way. We think that is a good thing to spend a life on.

The richness that makes a meal feel complete is, it turns out, a computable quantity — anchored to the half-activation of a receptor your parathyroid glands use to keep you alive, read by peptides your own cells make by the billion, forged across a bond the body's ordinary enzymes cannot cut. We gave it a number so we could design it on purpose. But the number was never the point. The point is the spoonful that stays with you — the one that says, without a single word and without any flavor you could name, *you are cared for*. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. We built the whole of this to put a little more of that into the world, grown clean, computed honestly, and given freely. That is the invisible axis. Now you can read it.

Kokumi · IFI-K · CaSR · The Scoville Splice Library, Volume X. Every cultivar, compound, enzyme, and recipe in this book resolves against the shared index and appears as its own entry in the master cross-book index. The richness is real; the numbers are computed; the claims stop exactly where the evidence does.