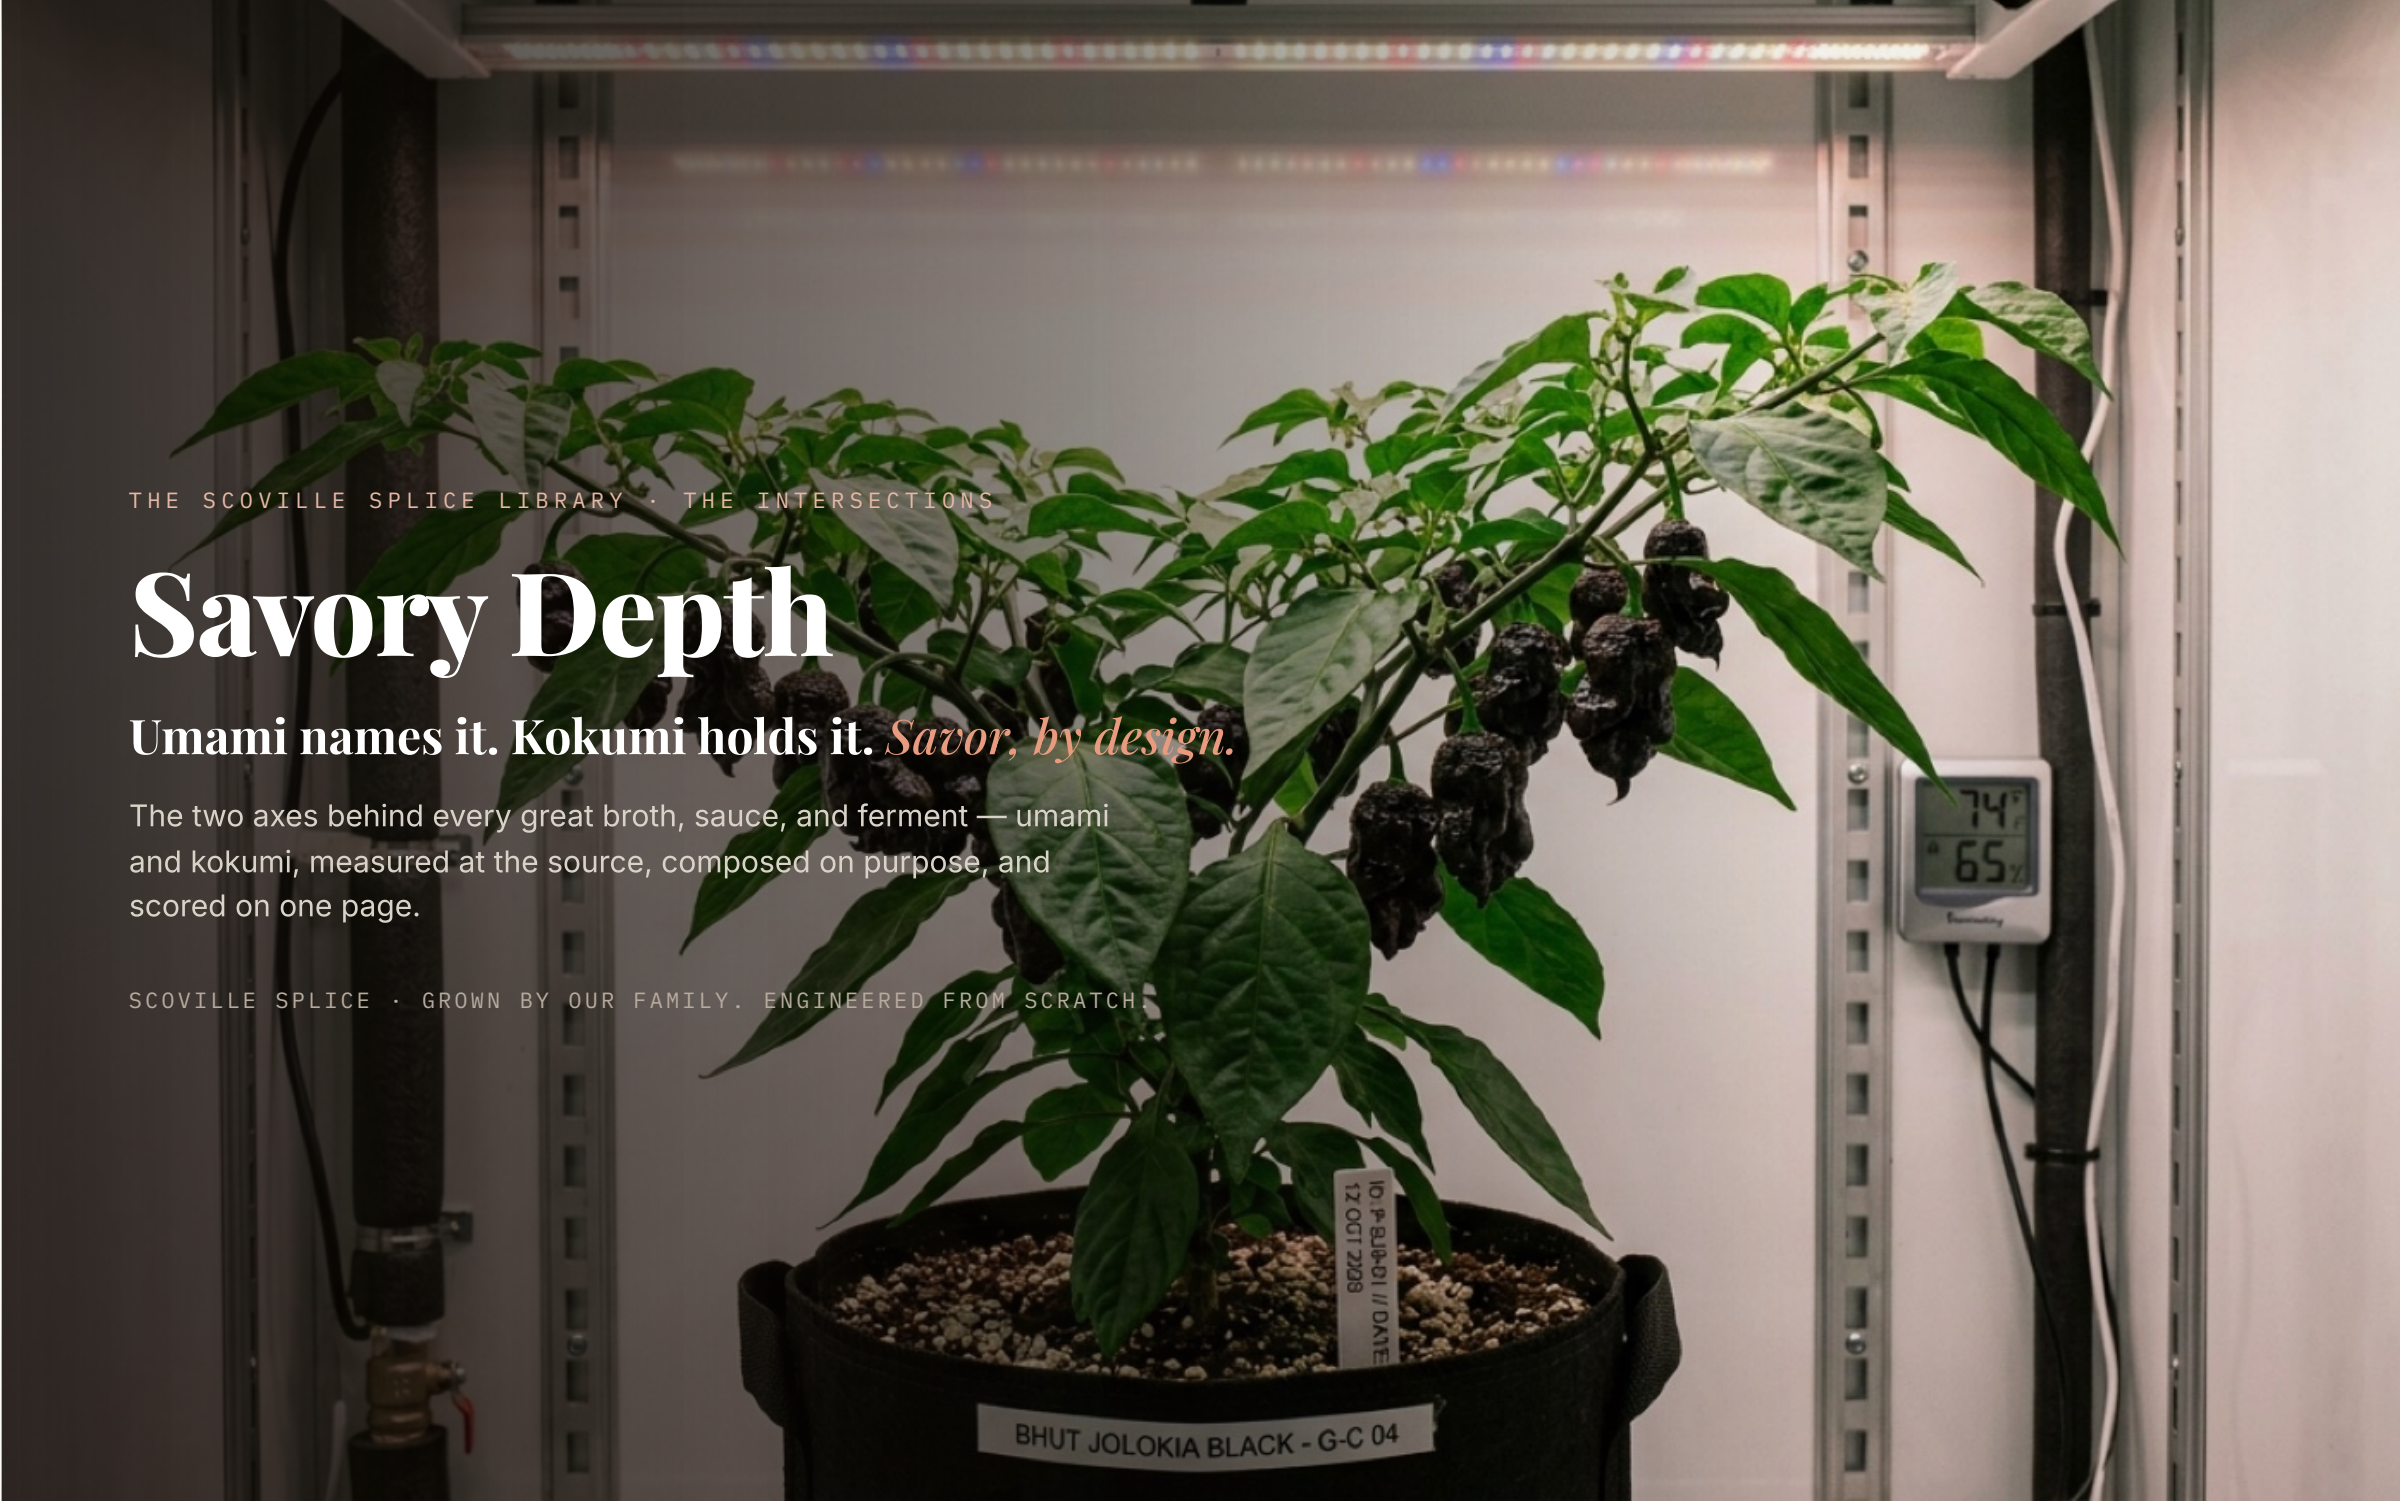

THE SCOVILLE SPLICE LIBRARY · THE INTERSECTIONS

Savory Depth

Umami names it. Kokumi holds it. *Savor, by design.*

The two axes behind every great broth, sauce, and ferment — umami and kokumi, measured at the source, composed on purpose, and scored on one page.

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

Contents

01 Opening essay · The taste of being fed

02 Chapter 1 · Umami — the note the body trusts.

03 Chapter 2 · Kokumi — the depth you feel, not taste.

04 Chapter 3 · The synergy — why savor multiplies.

05 Chapter 4 · Reading U×K on the Index — the number under the broth.

06 Chapter 5 · The glutamate engine — how a plant makes the note.

07 Chapter 6 · The ribonucleotide route — the gain on the note.

08 Chapter 7 · The γ-glutamyl pathway — where body is built.

09 Chapter 8 · The method — how we decide the edit.

10 Chapter 9 · Tomato — the savor bred out, grown back.

11 Chapter 10 · Soybean — one seed, both axes.

12 Chapter 11 · Yeast — the manufacturer's engine.

13 Chapter 12 · Koji, shiitake & the mushrooms — savor released by other life.

14 Chapter 13 · The alliums & the quiet kokumi — body, from the bulb and the leaf.

15 Chapter 14 · Designing a broth — the stack, assembled.

16 Chapter 15 · Sauces, pastes & the reduction — depth you can hold in a spoon.

17 Chapter 16 · The honest meat analog — savor that does not pretend.

18 Chapter 17 · Ferments — time as an ingredient.

19 Chapter 18 · Precision & the ≤7% CV — a savor you can promise.

20 Chapter 19 · Label, compliance & the clean line — honesty on the back of the pack.

21 Chapter 20 · The frontier & the table — where savor goes next.

Opening essay · The taste of being fed

There is a particular moment at our table — the lid comes off a pot that has been going a while, and before anyone has tasted anything, the whole kitchen leans in. No one engineered that reflex. It is older than cooking. The body reads savory the way it reads warmth from a fire: as a promise that it is about to be cared for. Of every sensation a plant can offer, this is the one we are least able to resist and least able to describe. We call it, loosely, depth.

Depth is two things pretending to be one. The first is **umami** — the bright savory note that says *protein, here, now* — free glutamate landing on the TAS1R1·TAS1R3 pair and on the taste-variant metabotropic glutamate receptors, the sensors that evolved to find nourishment and reward us for finding it. The second is **kokumi** — not a taste at all, but the body under the taste: the continuity, the roundness, the way a good broth seems to fill the whole mouth and stay. Umami you can point to. Kokumi you can only miss when it is gone. Together they are what a century of cooks have chased with bones and time and, when those ran short, with a shake of monosodium glutamate that delivered the note and left the body behind — a savor that arrives and then, flatly, leaves.



The research greenhouse before the lights come up — where the savory organisms are grown, not seasoned.

This is the most-requested intersection in our whole library, and it is the one manufacturers ask us for by feel rather than by name: *make it taste like it simmered all day.* What they are describing is **IFI-U × IFI-K** — the savory note and the savory body, composed in the right ratio so that one does not cancel the other. It is not a seasoning problem. It is an architecture problem, and it is solved most completely at the source, in a plant or a microbe grown to carry both.

We came to it the long way. The family work began with heat, because heat is the axis you can see — a pepper our model predicts at approximately 2.95 million SHU on the same Index that scores everything else. But the method that designs a burn is not a pepper method. It is a flavor method: measure the receptor, find the enzyme that feeds it, model the pathway, pull the flux, and lock the result to a batch-to-batch variation of **seven percent or less**. Savor turned out to be the axis where that discipline matters most — because savor is the one a kitchen spends the most hours chasing, and the one a factory has the least honest way to reproduce.

So this book is that architecture, drawn in full. The science of each axis, exactly — the receptors, the ligands, the mathematics that makes two savory compounds multiply instead of add. The handful of organisms — soybean, yeast, koji, the mushrooms, the restored tomato, the alliums — that carry savor natively, each with its real construct count and its designed profile. The biosynthesis, named enzyme by named enzyme, from the carbon skeleton the plant pulls out of its own respiration to the tripeptide that rounds a broth. And then the builds: broth, sauce, the honest savory of a meat analog that does not pretend, and the ferment where the plant keeps working after you have walked away. Every one of them scored, so you can taste the depth and read it on the same page.

A word on how this book is built, because savory depth rewards being read in order. It begins with the two receptors — the umami note and the kokumi body — because you cannot design what you cannot name, and these two have been mistaken for one for most of culinary history. It moves to the biosynthesis, naming every enzyme between a plant's respiration and the molecule that rounds a broth, because that is where the design actually happens. Then the organisms, one by one, each with its real construct count and its designed profile; then the builds, where the organisms become broth and sauce and ferment; and finally the discipline — precision, the honest label, the frontier — that turns a good result into a reliable product. It is a long book because savor is a deep one, and we would rather exhaust the real chemistry than pad a thin version of it. Read straight through, it is a single argument: that the oldest comfort a kitchen offers can be measured, and so can be grown.

We write this the way we cook — reverent about the living thing, exact about the chemistry, and certain that the two are the same act. A broth is a nourishment you give your body on purpose. We would like you to be able to design it.



PART I OF V

The Two Axes

TAS1R1·TAS1R3 · taste-mGluR · CaSR — the savory note and the savory body

Chapter 1 · Umami — *the note the body trusts.*

AXIS 03 · IFI-U · TAS1R1·TAS1R3 / taste-mGluR

Umami is the taste of being told the truth about food. When free glutamate settles onto the **TAS1R1·TAS1R3 heterodimer** — and onto the truncated, taste-specific variants of the metabotropic glutamate receptors that sit alongside it — the tongue is not sensing a flavor so much as delivering a verdict: *this carries protein; this is worth eating.* It is the one taste the body answers to with trust rather than caution, and it is the reason a broth can feel less like a dish than like an act of care.

Begin with the receptor, because everything the plant must do is shaped by it. TAS1R1 and TAS1R3 are **class C G-protein-coupled receptors**, the same architectural family as the sweet receptor and the calcium-sensing receptor we will meet in the next chapter. Each has a large extracellular **Venus flytrap domain** — two lobes hinged like a clamshell — connected through a cysteine-rich linker to the seven-pass membrane bundle. L-glutamate binds in the cleft of the TAS1R1 flytrap and pulls the two lobes shut. That closure is the signal. It propagates through the linker, tips the transmembrane bundle, and the receptor couples to its G-protein — **α-gustducin** at the taste cell — which activates phospholipase C-β2, releases IP_3, opens the TRPM5 channel, depolarizes the cell, and triggers a pulse of ATP onto the afferent nerve. From clamshell to nerve, savor is one conformational change amplified a billion times a second.

Spell the cascade out, because its length is why umami is slow to arrive and slow to leave — the quality a cook calls *lingering*. Gustducin's activated α-subunit frees its βγ pair, which drives **phospholipase C-β2** to cleave the membrane lipid PIP_2 into diacylglycerol and **inositol trisphosphate (IP_3)**; IP_3 opens the type-3 IP_3 receptor on the endoplasmic reticulum and floods the cell with stored calcium; that calcium opens **TRPM5**, a calcium-activated monovalent channel, depolarizing the taste cell; the depolarization opens **CALHM1/3**, the channel that releases ATP onto the $P2X_2/P2X_3$ receptors of the afferent nerve. Six proteins stand between the glutamate and the signal, and each is a gain stage. It is also why the umami dose-response is **cooperative**: the heterodimer's two flytraps and the amplifying cascade give a Hill coefficient greater than one — a curve flat at low glutamate and then steep, so savor has a threshold and then, past it, a sudden fullness. In solution that threshold sits near a tenth of a millimolar for pure glutamate and falls severalfold once a nucleotide is present, which is the difference between a tomato you must concentrate to taste and a broth that reads savory at a spoonful in a bowl. And the receptor is exquisitely stereoselective: only **L**-glutamate tastes — its mirror-image D-enantiomer is silent, and so is glutamate bound inside a peptide — because the flytrap reads both the α-amino group and the free side-chain carboxyl at once. The tongue is not detecting a molecule. It is detecting a molecule in exactly one shape.

Two facts about that binding site govern the whole axis. The first is that only **free** glutamate tastes. Glutamate locked inside a protein chain — the vast majority of the glutamate in any raw ingredient — cannot reach the flytrap cleft. It has to be liberated: by ripening, by heat, by the slow proteolysis of aging and fermentation, by enzymes that cut the peptide bond and set the amino acid loose. This is why a raw tomato is quietly savory and a sun-dried one is loud, why cheese gains umami as it ages, why a long ferment tastes of more than its ingredients. Savor is not how much glutamate a food contains; it is how much it has **released**.



FIG. 1 — The restored tomato in cross-section. The locular jelly around the seeds is where free glutamate concentrates; breeding for shipping drained it, and the native GAD/GDH route grows it back.

The second fact is that the flytrap has a **second site**. Adjacent to the glutamate cleft, on the same TAS1R1 lobe, is a pocket that binds **5′-ribonucleotides** — inosine monophosphate (IMP) from a seared or aged surface, guanosine monophosphate (GMP) from a dried mushroom. A nucleotide bound there does not taste of anything on its own. What it does is wedge the clamshell: it stabilizes the closed, active conformation, so that the same glutamate now holds the receptor shut far more effectively. The receptor's half-activating concentration — its $\mathbf{EC_{50}}$ — falls, sometimes by an order of magnitude. This is the synergy every cook knows by instinct: kombu *and* bonito, tomato *and* parmesan, the hard sear on the mushroom. It is the single most important thing about umami, and because it is multiplicative rather than additive, it is a design lever, not an accident. We return to its mathematics in Chapter 3.

For most of the last century we were quietly starved of the note. Breeding chose vegetables for the truck and the shelf, not the tongue, and the tomato — our richest everyday source of free glutamate — lost roughly **ninety-three percent** of its peak free-glutamate pool on the way to market. The flavor did not fail. It was bred out, one shipping-tolerance selection at a time, until a supermarket tomato read near **300** on the Index where a vine-ripe heirloom reads a thousand. We give it back at the source, through the plant's own glutamate machinery — the **GDH** amination that builds the pool and the **GAD** decarboxylation that drains it — rather

than dusting monosodium glutamate over the top at the end. Restored, the tomato climbs to **2,600 and beyond**; the chemistry of that climb is Chapter 9.

Where does umami actually live? A map is useful, because the savory floor a kitchen stands on is a small number of very loud sources. Free glutamate content, measured as milligrams per hundred grams, spans two orders of magnitude across ordinary foods:

SOURCE	FREE GLUTAMATE (MG/100 G)	CHARACTERISTIC NUCLEOTIDE
Kombu (dried kelp)	1,400–3,200	–
Yeast extract	~1,960	AMP / GMP
Parmesan, aged	1,200–1,680	–
Soy sauce	400–1,700	–
Fish sauce	~950	–
Cured ham	~340	IMP
Green tea	220–670	–
Ripe tomato	150–250	–
Dried shiitake	~70	GMP ~150
Katsuobushi (dried bonito)	~40	IMP 470–700

Read the table as a design brief rather than a grocery list. The glutamate giants — kombu, yeast, aged cheese, soy sauce — carry the note. The nucleotide carriers — bonito, shiitake, cured meats — carry the *potentiation*. A great savory system pairs one of each, which is exactly what every traditional broth on earth independently discovered: dashi marries kelp glutamate to bonito inosinate; a Western stock marries tomato and long-roasted bone; a mushroom broth leans on its own guanylate. Our work is to put both halves inside organisms we can grow to a locked profile, so the pairing is designed rather than hoped for.

There is an evolutionary argument folded into that universality, and it is the kind of fact that makes the axis feel less like chemistry and more like deep history. The umami receptor is the body's protein detector, and a clade's diet is written in the state of the gene. Obligate carnivores such as the domestic cat carry a broken, pseudogenized **sweet** receptor but keep the umami one intact and sharp; the giant panda, which abandoned meat for bamboo, carries the opposite loss — a pseudogenized **TAS1R1**, the umami receptor gone dark in a lineage whose food no longer rewards it. Each species keeps the savor sensor its diet pays for. Humans kept both receptors and then built whole cuisines on the keeping. Every broth tradition that paired a glutamate source with a nucleotide source was, without knowing it, engineering for the second binding site on the TAS1R1 flytrap — running the synergy of Chapter 3 by taste alone, a thousand years before anyone could draw the receptor it ran on. When we design a savory organism, we are not inventing the pairing. We are finishing a sentence the tongue started.

On the Index, **IFI-U** scores that loudness directly. The formula is the same one that governs every axis in the library:

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

Here the "candidate" is the food itself, and EC_{50} is the **dilution** at which the food's savor half-activates a panel of umami receptors in a cell-based assay. A food heavy with free glutamate and potentiating nucleotide half-activates at a high dilution — you can water it down and it still speaks — so its EC_{50} is low and its IFI-U is high. A washed-out supermarket tomato has to be nearly undiluted to register; its EC_{50} is high and its score collapses toward 300. The number is not a measure of concentration. It is a measure of how loudly a food says *protein* — and loudness, unlike concentration, is the thing the tongue actually reports.

That is the axis in one receptor and one ratio. It is bright, it is trustworthy, and it is the half of savor you can name. The other half you cannot — and it is the half that makes the first one last.

Chapter 2 · Kokumi — *the depth you feel, not taste.*

AXIS 05 · IFI-K · CaSR (calcium-sensing receptor)

Kokumi is the axis you can only notice by its absence. It has no flavor of its own — nothing to point to, nothing to name — and yet a dish without it tastes thin, hollow, finished too soon. A tomato sauce with the umami exactly right but the kokumi missing reads *flat*: correct and lifeless, like a chord played without its lowest string. What kokumi contributes is **body**: continuity across the palate, roundness, the sense that a mouthful arrives whole, fills the mouth, and stays whole as it fades. It is the quality a taster calls *depth* and then cannot define, because the thing doing the work never announces itself.

The Japanese word *kokumi* — roughly, "rich taste" — was coined in the 1980s by researchers at a flavor house who noticed that certain extracts made a soup taste fuller without making it taste of anything new. The extracts, chased down, turned out to be peptides. Specifically **γ-glutamyl peptides**: short chains in which a glutamate residue is linked to its neighbor through its *side-chain* carboxyl — the γ-carbon — rather than the usual backbone α-carbon. Glutathione, the cell's own antioxidant tripeptide (γ-Glu-Cys-Gly), was the first and is the archetype. Others followed, and one, **γ-glutamyl-valyl-glycine** (γ-Glu-Val-Gly), proved far more potent than the rest.

The mechanism is almost beautiful in its indirection. These peptides are agonists of the **calcium-sensing receptor, CaSR** — another class C GPCR, the same Venus-flytrap architecture as the umami and sweet receptors, but one whose day job across the body is to monitor extracellular calcium and keep it in range. In the mouth, CaSR sits on a subset of taste cells and, crucially, **produces no taste signal of its own.** You cannot taste CaSR activation the way you taste sweet or sour. What CaSR activation does instead is **modulate** the tastes already present: it potentiates and sustains sweet, salt, and above all umami, lowering the effective threshold at which the mouth reads them and stretching the perception across time. Sensory panels resolve the effect into three descriptors, and they are the working vocabulary of the axis: an enhanced **initial impact** (the first instant tastes more complete), heightened **mouthfulness** or continuity (the taste spreads and fills), and a longer, rounder **aftertaste** (it is slow and kind about leaving). Kokumi does not add a voice to the chorus. It widens the room the chorus sings in.

The receptor behind that widening deserves a closer look, because it is the strangest sensor on the tongue. **CaSR** is, everywhere else in the body, a thermostat for calcium: it sits on the parathyroid gland and the kidney and reads the blood's calcium level, tuning the hormones that keep it in range. Its presence on taste cells was a genuine surprise, resolved only when researchers chasing the kokumi peptides found that these γ-glutamyl compounds are

allosteric agonists of CaSR — they bind not the receptor's calcium site but a modulatory pocket in the same class-C Venus-flytrap domain the umami and sweet receptors share, and they sensitize the receptor to the calcium already present. The receptor's other agonists are instructive about where body comes from in a kitchen: **extracellular calcium itself** is kokumi-active, as are certain aromatic amino acids and polyamines, which is part of why aged, calcium-rich, and long-cooked foods taste full. The γ-glutamyl peptides are simply the most potent and the most designable of its ligands. And because CaSR activation on a taste cell modulates the output of its neighbors rather than firing a taste of its own, kokumi is, mechanistically, a *cross-talk* signal — one receptor leaning on the others. That is the literal truth behind the feeling that body sits *underneath* taste rather than beside it: it is a receptor whose entire job is to change how the others are heard.



FIG. 2 — Garlic in cross-section. The clove is a reservoir of γ-glutamyl peptides — storage forms that, cut and aged, become the kokumi a long braise pays hours to build.

Not all γ-glutamyl peptides are equal, and the differences are the whole of kokumi design. Potency at CaSR is exquisitely sensitive to the residues that follow the γ-glutamyl head — a structure-activity relationship clear enough to engineer toward. Reported potencies, expressed relative to glutathione as the benchmark, put a small hydrophobic middle residue followed by glycine at the top of the ladder:

Γ-GLUTAMYL PEPTIDE	RELATIVE CASR POTENCY (GSH = 1.0)	FOUND IN
γ-Glu-Val-Gly	~12.8	fermented sauces, designed yeast
γ-Glu-Cys-β-Ala	~4.6	legumes
γ-Glu-Cys	~1.5	onion, garlic, glutathione precursor
Glutathione (γ-Glu-Cys-Gly)	1.0	yeast, scallop, nearly all cells
γ-Glu-Val	~0.4	aged cheese, garlic

The lesson sits in the top row. γ-Glu-Val-Gly is roughly **thirteen times** the kokumi of glutathione at the same concentration, which means a plant or a ferment engineered to route even a modest fraction of its glutathione flux toward that tripeptide carries disproportionate body for the metabolic cost. It is the single most efficient kokumi molecule we know, and it is the target the γ-glutamyl pathway in Chapter 7 is pulled toward.

That potency has been confirmed where it matters, on human panels. In controlled sensory work, γ-Glu-Val-Gly added at a few parts per million to a chicken broth or a reduced-fat food raised all three kokumi descriptors — initial impact, mouthfulness, persistence — at concentrations far below its own taste threshold, which is the signature of a true modulator: it does the work without ever being detected as a flavor of its own. The structure-activity rule behind the ladder is now clear enough to engineer against. The **γ-glutamyl head is obligatory** — move the glutamate's link from its side-chain γ-carbon to the ordinary α-carbon, as in an everyday peptide bond, and the molecule goes silent at CaSR — and the middle residue wants to be **small and hydrophobic** (valine, cysteine, α-aminobutyrate), with a short third residue, glycine, sealing the most potent of them. It is a compact rule, and a compact rule is a gift to an engineer, because it tells the γ-glutamyl transpeptidase of Chapter 7 exactly which acceptor to prefer. It is also why kokumi became the flavor industry's quiet tool for **fat reduction**: a γ-glutamyl peptide restores to a leaner formulation the continuity and mouth-coating that fat used to supply — body without the lipid — and it does it, again, without adding a taste anyone can name.

Where does kokumi come from in a kitchen, before anyone engineers it? From **time and heat** — the two things a cook spends to buy depth. Long cooking and slow aging generate γ-glutamyl peptides from the free amino acids that proteolysis releases; the enzyme that builds them, γ-glutamyl transpeptidase, is active in aging cheese, fermenting soy, sprouting alliums, and the gentle warmth of a braise held for hours. This is the real reason a slow-cooked thing tastes *complete* and a quick one tastes correct-but-thin: the long version has had time to synthesize its own kokumi. It is also why kokumi has always been the **expensive** axis for a manufacturer. You bought it with hours at the stove, with the yield loss of a long reduction, or with fat — and none of those scale cleanly into a shelf-stable product. We build it in at the source instead, in the organisms that carry it natively: **garlic** and the alliums,

yeast, **soybean**, **vanilla** for its rounding mouthfeel, the quiet kokumi of **scallop** and aged ferments. Grown to profile, they give a dish the body of a long braise without the braise.

On the Index, **IFI-K** does not measure a flavor, because kokumi is not one. It measures *how whole a thing feels in the mouth and how long it is kind enough to stay* — operationalized as the γ-glutamyl peptide pool's ability to activate CaSR and the resulting modulation of a reference umami stimulus. The same ratio governs it:

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

where EC_{50} is the dilution at which the candidate's peptide load half-maximally activates CaSR. A broth rich in γ-Glu-Val-Gly half-activates at high dilution — its EC_{50} is low, its IFI-K is high — and it delivers that enveloping continuity even when thinned. A lean stock with little peptide load scores near the floor no matter how much glutamate it carries; it will taste sharp and short.

This is why umami and kokumi belong in one book and on one scorecard. They are not two intensities of the same thing. They are two different receptors — TAS1R1·TAS1R3 for the note, CaSR for the body — reading two different chemistries, and because the receptors are distinct, the axes **stack without competing**. A food can be loud in U and empty in K (monosodium glutamate in water: all note, no body) or quiet in U and full in K (a lean dashi finished with a kokumi peptide: modest note, long body). Great savor is both at once, in a ratio you can now specify. The next chapter is the mathematics of specifying it.

Chapter 3 · The synergy — *why savor multiplies.*

MECHANISM · allosteric potentiation · the Yamaguchi relation

The first thing to understand about savory design is that it does not obey arithmetic. Add a spoon of sugar to a cup of tea and you get exactly one spoon of sweeter. Add a 5′-nucleotide to a glutamate solution and you can get *eight times* the savor, or thirty — a result so far outside addition that when Akira Yamaguchi first quantified it in 1967, the number looked like an error. It was not. Umami is the one basic taste built on **synergy**, and synergy is the reason the axis rewards engineering so richly: small, cheap changes to a nucleotide pool produce enormous changes in perceived savor, because the receptor is doing the multiplying for free.

Start with the phenomenon, then the mechanism, then the math. The phenomenon is that glutamate and the 5′-ribonucleotides — inosinate (IMP) and guanylate (GMP) chief among them — are weak to moderate on their own and overwhelming together. A solution of glutamate at a concentration you would call "pleasantly savory," given a trace of IMP that you could not taste at all by itself, jumps to a savor you would call "intense." The two are not sharing a flavor. They are gating each other.

The mechanism lives back at the Venus flytrap of Chapter 1. Glutamate binds in the cleft of the TAS1R1 lobe and tends to pull the clamshell shut; the nucleotide binds a **separate, adjacent site** on the same lobe, nearer the hinge, and its binding **stabilizes the closed conformation** glutamate is trying to reach. With the nucleotide present, the receptor needs far less glutamate to spend the same fraction of its time in the active, signaling state. In the language of pharmacology, the nucleotide is a **positive allosteric modulator**: it does not activate the receptor itself (little taste alone), but it lowers the agonist's EC_{50} — the concentration at which the agonist drives the receptor to half its maximum. Lower EC_{50} is, by the Index definition, higher loudness. The nucleotide buys IFI-U without adding any glutamate at all.



FIG. 3 — A receptor assay plate on the stage. Each well is a dose of savor read against a cell line expressing the umami heterodimer; the leftward shift of the curve when a nucleotide is added is the synergy, measured.

Now the mathematics, because this is where the axis becomes designable. Yamaguchi's relation expresses the perceived savor of a glutamate-plus-nucleotide mixture as an **equivalent glutamate concentration** — the concentration of pure glutamate that would taste equally strong:

```
y = u + β · u · v
```

Here *u* is the glutamate concentration and *v* the nucleotide concentration (both as grams per 100 mL), *y* is the equivalent glutamate concentration, and **β** is the synergy constant for that particular nucleotide. The constant is large. For glutamate with **IMP**, β ≈ **1,218**. For glutamate with **GMP** — the mushroom nucleotide — β ≈ **2,800**, more than twice as potent. The linear term *u* is the glutamate tasting on its own; the product term **β·u·v** is the synergy, and because it multiplies the two concentrations, it dominates the moment both are present in any quantity.

Work an example, because the number is the argument. Take a modest 1:1 mixture, *u* = *v* = 0.05 g/100 mL — concentrations each near the threshold of "mildly savory." With IMP:

```
y = 0.05 + 1,218 × (0.05)(0.05) = 0.05 + 1,218 × 0.0025 = 0.05 + 3.045 ≈ 3.1 g/100 mL equivalent
```

A tenth of a gram of actual savory compound, split evenly, tastes like **three grams** of glutamate alone — a roughly thirty-fold return on the mixture's total mass. Swap IMP for GMP and the equivalent climbs past seven grams. This is not a seasoning trick; it is the arithmetic of the receptor, and it is why every dried-mushroom-and-kelp broth in the world tastes like far more than the sum of two cheap ingredients.

The shape of that arithmetic is worth seeing whole. Hold the *total* savory mass fixed — say a tenth of a gram per hundred millilitres — and vary only the ratio of glutamate to nucleotide, and the perceived savor traces a steep arc that **peaks near one-to-one** and falls off toward either pure component. Plotted as an isobologram — the curve of all glutamate-plus-nucleotide mixtures that taste equally strong — the synergy bows the line deep toward the origin, the geometric signature of two agents that do far more together than apart. This is the quantitative reason a cook's instinct to *balance* a glutamate source against a nucleotide source is correct: the maximum is not at the extremes but in the middle, and a broth heavy on kelp and light on bonito leaves most of its possible savor on the table.

The structural basis sharpens the design rule. The nucleotide's potency is not generic; it is specific to the **5′-phosphate**. Move the phosphate to the 2′ or 3′ position of the ribose and the synergy collapses — the 5′-mononucleotide is the active form, which is exactly why the release enzymes of Chapter 6 are **5′**-phosphodiesterases, cutting RNA to leave the phosphate where the receptor wants it. Structural modelling of the TAS1R1 flytrap places the nucleotide's phosphate bridging two basic residues near the domain's hinge, clamping the two lobes in their closed, glutamate-bound pose. The purine ring and the 6-position substituent fine-tune

the grip, which is why guanylate (a 2-amino-6-oxo purine) out-synergizes inosinate (6-oxo only): the extra amino group tightens the clamp, and β more than doubles. For an engineer this means the nucleotide half of a savory organism has a single, unambiguous target — maximize free **5′-GMP** — and a single enzyme to protect, the phosphodiesterase that produces it in the right isomer.

The same physics, written as the EC_{50} shift the Index actually scores, gives the design lever directly. Model the nucleotide as a simple allosteric cofactor and the apparent glutamate EC_{50} falls with nucleotide concentration:

```
EC50,app = EC50,0 ÷ (1 + [N] / K_N)
```

with $EC_{50},0$ ≈ 1.5 mM for glutamate alone at the umami heterodimer and K_N ≈ 0.1 mM for IMP's allosteric site. Add IMP to 0.5 mM — a trace — and the factor is (1 + 0.5/0.1) = 6: the apparent EC_{50} drops from 1.5 mM to **0.25 mM**, and IFI-U rises sixfold. The organism did not make six times the glutamate. It made a little guanylate, and the receptor did the rest.

STACK	GLUTAMATE EC_{50}	RELATIVE IFI-U	PERCEIVED
Glutamate only	1.5 mM	1.0×	savory, short
+ IMP (0.5 mM)	0.25 mM	6×	savory, rounded
+ GMP (0.5 mM)	~0.13 mM	~11×	deep, mouth-filling
+ GMP + kokumi peptide	~0.13 mM	~11× (U)	deep **and** continuous

That last row is the whole thesis of the book in one line, and it introduces the second kind of synergy — the one between the axes. Kokumi does not touch the umami flytrap; it works through CaSR, a different receptor entirely. So when a γ-glutamyl peptide lengthens and rounds the umami signal, it is not competing with the nucleotide for a binding site — it is **stacking orthogonally**, adding body to a note that the nucleotide has already made loud. U synergizes *within* itself (glutamate × nucleotide) and *across* to K (umami × kokumi), and neither synergy cannibalizes the other, because three distinct receptors are doing three distinct jobs.

One practical corollary closes the point. Because the synergy multiplies two concentrations, its return is greatest when the two are **balanced** — which is why the dashi of Chapter 14 pairs roughly equal weights of kelp and mushroom, and why a savory system lopsided toward glutamate leaves most of its potential unspent. And not every nucleotide pulls its weight: **5′-AMP**, the adenylate, is a weak synergist, a small fraction of IMP's β, which is exactly why the AMP-deaminase step that converts adenylate to inosinate earns an enzyme of its own in the yeast of Chapter 11. The receptor is picky about which 5′-phosphate it rewards, and the whole of nucleotide design is feeding it the one it likes best. Synergy is not a bonus the organism throws in on top of its glutamate; it is the main event, and an organism that carries both halves of it is carrying the multiplier itself, not merely the two numbers it multiplies. A fully designed savory system therefore has three dials, not one: the glutamate pool

(the note), the nucleotide pool (the note's gain), and the γ-glutamyl peptide pool (the body). Turn all three and the result is multiplicative in two directions at once.

This is why savor is the axis we most want grown into an organism rather than shaken on at the end. A manufacturer blending MSG and disodium inosinate from drums can hit the note. But the ratio drifts batch to batch, the kokumi is simply absent, and the label carries three additives. An organism engineered to carry glutamate, guanylate, and γ-Glu-Val-Gly together, in a locked ratio, delivers the whole multiplicative stack as a single ingredient — one line on the label, one profile on the Index, one broth that tastes like it simmered all day because, chemically, it did.

Chapter 4 · Reading U×K on the Index — *the number under the broth.*

METHOD · the IFI formula · the assays · the scorecard

Every claim in this book resolves to a number, and the number is always built the same way. The Iyengar Flavor Index exists so that a tomato, a dashi, a soy sauce, and a designed yeast can be laid on one page and compared without anyone having to trust an adjective. For savor, that discipline matters more than anywhere else in the library, because savor is the axis most often sold on adjectives — *rich, deep, rounded, umami-forward* — and least often sold on evidence. The Index replaces the adjective with an EC_{50}.

Recall the master formula, and read it slowly, because its direction is counterintuitive the first time:

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

EC_{50} is the concentration — here, the **dilution of the food** — at which a stimulus drives its receptor to half of full response. The reference is fixed. The candidate's EC_{50} sits in the denominator, so a **low** EC_{50} — a food so loud you can dilute it far and still half-activate the receptor — produces a **high** score. Loudness, not concentration. A food can be chemically rich and perceptually quiet (glutamate still bound in protein, invisible to the flytrap), and the Index scores what the tongue gets, not what the laboratory assay would find after acid hydrolysis. For savor this is the only honest choice, because the entire axis turns on what has been *released*.

Two receptors mean two assays. For **IFI-U**, we run the candidate across a cell line expressing the human **TAS1R1·TAS1R3** heterodimer with a calcium-indicator reporter — the receptor couples through a promiscuous or chimeric G-protein to a rise in intracellular calcium, which a fluorescent reporter turns into light. Dose the cells across a dilution series, fit the response to a four-parameter logistic (the Hill equation), and read the midpoint: that is EC_{50}(candidate). Written out, the fitted curve is

$$R(c) = R_{min} + (R_{max} - R_{min}) / (1 + (EC_{50} / c)^{n_H})$$

where R is the reporter response at food concentration c, R_min and R_max bracket the signal, and n_H is the Hill coefficient — the cooperativity of Chapter 1, greater than one for the umami heterodimer, which is why the curve rises steeply through its midpoint. The half-response point, where R sits exactly between its floor and ceiling, defines EC_{50}. Everything else about the axis is this one number, read off this one curve, inverted against a fixed reference.

The reference is what makes the scores comparable across the whole library, so it is chosen once and never moved. For **IFI-U** the reference EC_{50} is anchored so that **1,000 corresponds to the free-glutamate loudness of a vine-ripe heirloom reference pulp** — in assay terms, a reference EC_{50} near **1.5 mM** glutamate-equivalent. The arithmetic then falls straight out of the master formula. A washed-out commodity tomato that must be run nearly neat to half-activate — an EC_{50} near **5.0 mM** — scores 1,000 × 1.5 ÷ 5.0 ≈ **300**, the number Chapter 1 opened on. A restored tomato whose glutamate and staged nucleotide let it half-activate at a tenth of that concentration — an EC_{50} near **0.58 mM** — scores 1,000 × 1.5 ÷ 0.58 ≈ **2,600**. Same formula, same reference, two tomatoes nearly nine-fold apart in loudness, read on one axis without a single adjective.

A reference card of the stimuli behind the savory axes, with the approximate, assay-dependent half-activating concentrations that anchor them:

STIMULUS	RECEPTOR	APPROX. EC_{50}	NOTE
L-glutamate (reference)	TAS1R1·TAS1R3	~1.5 mM	the IFI-U 1,000 anchor
L-glutamate + saturating IMP	TAS1R1·TAS1R3	~0.25 mM	synergy, 6× louder
L-glutamate + saturating GMP	TAS1R1·TAS1R3	~0.13 mM	synergy, ~11× louder
L-theanine	TAS1R1·TAS1R3	~6 mM	the brothy note of tea
Glutathione (GSH)	CaSR	~2 mM	the IFI-K 1.0 benchmark
γ-Glu-Val-Gly	CaSR	~0.15 mM	~13× GSH, the body target
Ca^{2+}	CaSR	~3 mM	native agonist; aged-food body

Work the kokumi score the way we worked the umami one. Anchor IFI-K so that **1,000 corresponds to the CaSR loudness of the glutathione benchmark** — a reference EC_{50} near 2 mM. A broth whose γ-glutamyl pool is dominated by γ-Glu-Val-Gly, roughly thirteen-fold more potent, half-activates CaSR near 0.15 mM, which at the peptide level would read 1,000 × 2 ÷ 0.15 ≈ 13,000 — far past the ceiling, which is precisely why the *scorecard* value is anchored not to raw CaSR activation but to the **modulation of a reference umami curve** and held under the 5,000 maximum. A kokumi reading therefore always means *how much body this delivers to a real savory stimulus*, never a laboratory abstraction a tongue would never meet.

The dual read on the umami axis — once as presented, once with a saturating IMP spike — is the most diagnostic pair of numbers on a savory scorecard, because the *gap* between them tells a formulator what the food is missing. A food that is loud as presented and barely louder with added IMP is already carrying its own nucleotide — a dried mushroom, an aged ferment. A food that is quiet alone but jumps with the spike is all glutamate and no gain — a fresh tomato, a lean vegetable stock — and the size of the jump is a direct instruction to add a guanylate source. Read this way, the Index does not only score savor; it **prescribes the fix**, which is the difference between a measurement a formulator admires and one a formulator uses. The axis score is anchored to that modulation, because a kokumi compound that activates CaSR in a dish but does not round the umami is, for a cook's purposes, inert.

Every axis is normalized to a common ceiling so the eight can share a radar plot and a scorecard: the **per-axis normalization maximum is 5,000**. Nothing is scored above it; a reading that would exceed it is held at the ceiling and flagged. This keeps Phoenix's engineered heat (IFI-H 4,200) and a designed broth's savor on the same visual field without either pinning the chart. And every score carries the library's quality bar: a profile is not locked until batch-to-batch variation falls to **seven percent or less**, because a savor you cannot reproduce is a savor you cannot sell.

What does the readout look like? Here is a plain reference stock — our restored-tomato-and-shiitake base, before anything is built on it — scored across all eight axes so the shape of a savory thing is visible at a glance:

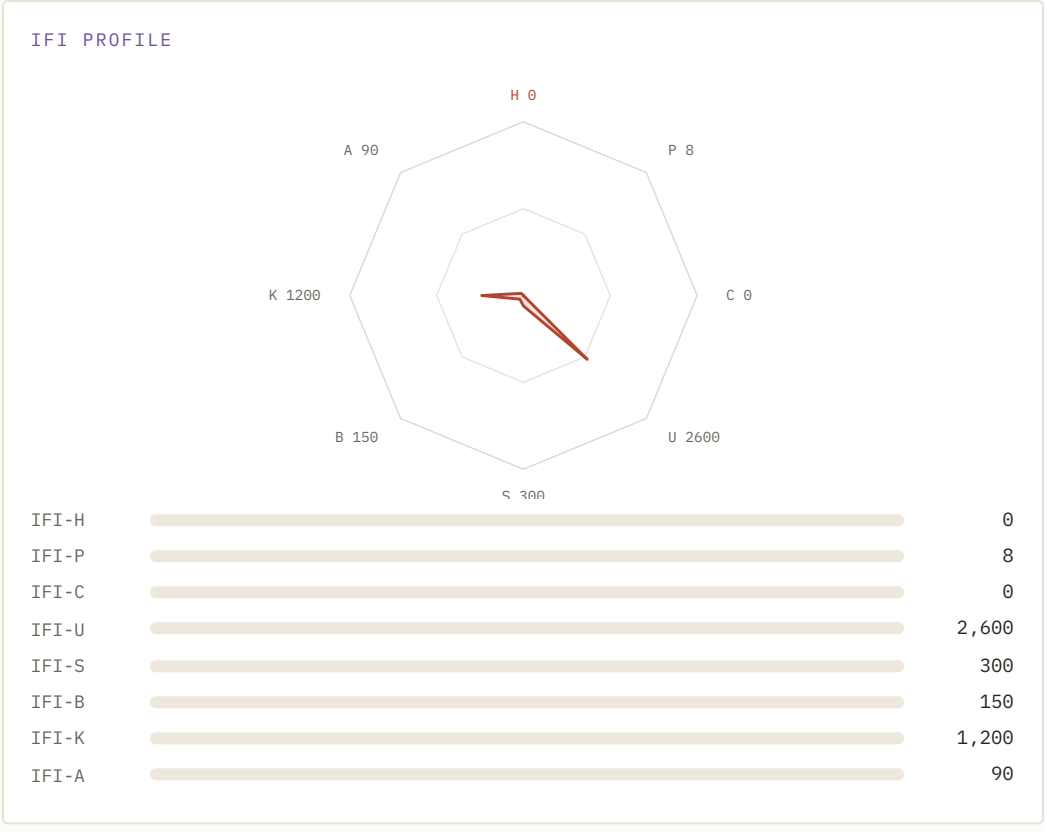

Read it as a cook reads a face. The two tall numbers, U and K, are the savory pair — the note at 2,600, the body at 1,200 — and everything else is near the floor: a whisper of pungency from the alliums, a touch of residual sweetness, no heat, no bitterness, no cool, almost no astringency. This is what *savory depth* looks like as a signature: two dominant axes and a quiet field, the opposite of a

superhot's heat-pinned spike or a wine's astringent-and-bitter pair. When you see that shape on a scorecard anywhere in this book, you are looking at something built to taste like broth.

AXIS	CODE	RECEPTOR	ASSAY READOUT
Umami	IFI-U	TAS1R1·TAS1R3	Ca^{2+}-flux EC_{50}, ± IMP spike
Kokumi	IFI-K	CaSR	activation EC_{50} + umami-curve shift
Sweetness	IFI-S	TAS1R2·TAS1R3	Ca^{2+}-flux EC_{50}
Zero-Bitter	IFI-B	TAS2R38	attenuation vs. reference
Pungency	IFI-P	TRPA1	Ca^{2+}-flux EC_{50}

The scorecard is not decoration. It is the contract between the plant and the kitchen — a promise that what the greenhouse grew, the tongue will read, and that both can be checked against the same midpoint on the same curve. With the two axes defined, the synergy quantified, and the Index reading established, we can go to where savor is actually made: the enzymes.



PART II OF V

The Biosynthesis

GS·GOGAT · GDH · GAD · the purine branch · γ-glutamyl transpeptidase · glutaminase

Chapter 5 · The glutamate engine — *how a plant makes the note.*

PATHWAY · nitrogen assimilation · GS/GOGAT · GDH · GAD

To raise umami at the source, you have to raise free glutamate, and to raise free glutamate you have to understand that glutamate is not a flavor molecule the plant makes for us. It is the **central hub of nitrogen metabolism** — the single amino acid through which nearly all nitrogen enters the plant's biochemistry and from which most other amino acids are built. Every leaf is already an engine for making glutamate. Our work is not to install the engine but to change what it does with its output: to let the pool accumulate in the fruit as free, tasteable amino acid rather than being spent as fast as it is made.

Follow the carbon in. The skeleton of glutamate is **2-oxoglutarate** — α-ketoglutarate, α-KG — a five-carbon acid the plant pulls straight out of its own **TCA cycle**, the respiratory loop that burns sugar for energy. This is the first design fact: the supply of glutamate's backbone is tied to respiration, so an organism's savory potential is coupled to how much α-KG it can divert from energy metabolism without starving itself. The nitrogen comes in separately, as ammonium, and the plant has two routes to weld nitrogen onto the skeleton.

The primary route is the **GS/GOGAT cycle**. Glutamine synthetase (**GS**) takes an existing glutamate, adds ammonium to its side chain at the cost of ATP, and makes **glutamine**. Glutamate synthase — **GOGAT**, glutamine:2-oxoglutarate aminotransferase — then takes that glutamine, hands its new amide nitrogen to a molecule of α-KG, and produces **two glutamates** where there was one. Run the cycle and the glutamate pool grows, one net molecule per turn, every turn drawing down α-KG and ammonium and ATP. The second route is more direct and more reversible: **glutamate dehydrogenase (GDH)** aminates α-KG with free ammonium using NAD(P)H as reductant, in a single step, and — critically — it runs *both ways*, synthesizing glutamate when nitrogen and reductant are plentiful and tearing it back down to α-KG when the cell needs carbon or energy. GDH is the valve. During tomato ripening, the balance of GDH activity tips toward synthesis, and the fruit's free-glutamate pool climbs steeply in the last days on the vine. That late climb is exactly the window a shipping-timed harvest cuts off.



FIG. 4 — The sequencing bench. Before a single edit, the whole glutamate pathway of a candidate organism is read, assembled, and checked against the genome-scale model that decides where the flux is actually lost.

So much for making glutamate. The subtler half of the engine is what *drains* it, because a plant that synthesizes glutamate furiously and spends it just as fast tastes of nothing. There are three drains, and savory design is largely the art of narrowing them without harming the plant.

The first and most important is **glutamate decarboxylase (GAD)**, which lops the carboxyl off glutamate to make **γ-aminobutyric acid (GABA)** and carbon dioxide. GAD is stress-activated and calcium-regulated through an autoinhibitory C-terminal domain, and in many fruits it runs hard early in development, shunting a large fraction of glutamate into the GABA pool. In tomato, three genes carry the activity — **SlGAD1, SlGAD2, SlGAD3** — and SlGAD3 dominates in the fruit. Attenuate it, and glutamate that would have become tasteless GABA stays as tasteable glutamate. This is the single highest-leverage edit on the umami axis, and it is subtractive: we are not adding a foreign gene, we are turning down one the plant already throttles itself.

The regulation of GAD is itself the reason the edit is delicate rather than brutal. Plant GAD carries a **calmodulin-binding autoinhibitory domain** at its C-terminus: when cytosolic calcium rises, calmodulin binds that tail and *activates* the enzyme, which is why stress and wounding — calcium-signalling events — spike GABA at glutamate's expense. The enzyme also has an acidic pH optimum near 5.8, so cytosolic acidification switches it on. We do not want the plant's stress response broken, only its fruit-stage drain narrowed, so the edit is tuned to the fruit and to the ripening window rather than applied genome-wide. And the GABA it would have made is not simple waste: GABA feeds the **GABA shunt**, a bypass around two steps of the TCA cycle in which GABA transaminase (GABA-T) and succinic-semialdehyde dehydrogenase (SSADH) return the carbon to the cycle as succinate. So glutamate lost to GAD is lost twice over — once as savor and once as carbon pulled out of the pool and spent on respiration. Narrowing the drain keeps both.

Upstream, the nitrogen that feeds the whole engine deserves a line, because savor begins as fertilizer. A plant takes up nitrogen mostly as **nitrate**, reduces it through nitrate and nitrite reductase to **ammonium**, and only then can the GS/GOGAT cycle weld it onto a carbon skeleton as glutamate. The fruit itself makes little of this; it

imports the raw materials through the phloem — sucrose for carbon-and-energy, amino acids including glutamine as nitrogen currency — and runs its local glutamate engine on that supply during development. This is why the savory potential of a tomato is set partly in the leaf and the root, and why a plant engineered for savor is engineered as a whole system: the fruit can only make loud what the plant delivers to it. The second drain is **transamination** — the family of aminotransferases, chiefly aspartate aminotransferase, that trade glutamate's nitrogen to other keto-acids to build other amino acids. Some of that traffic is essential; the plant must make its proteins. But a measured reduction in the export aminotransferases keeps more nitrogen resident in the glutamate pool. The third drain is simply **incorporation into protein**, which we leave almost entirely alone: starving the plant of its own structural amino acids is the fast way to a sickly organism and a worthless harvest.

The pathway, drawn as a ledger of steps:

STEP	ENZYME	SUBSTRATE → PRODUCT	ROLE FOR UMAMI
Carbon supply	TCA cycle	→ 2-oxoglutarate	feeds the skeleton
Assimilation	GS	glutamate + NH_4^+ → glutamine	primes the cycle
Assimilation	GOGAT	glutamine + 2-OG → 2 glutamate	net pool growth
Amination	GDH	2-OG + NH_4^+ + NAD(P)H ⇌ glutamate	the reversible valve
Drain	GAD	glutamate → GABA + CO_2	**attenuate**
Drain	aspartate aminotransferase	glutamate + OAA ⇌ 2-OG + aspartate	trim
Sink	ribosome	glutamate → protein	leave intact

Now the discipline that keeps this from being guesswork. You cannot simply up-regulate GDH and down-regulate GAD and hope — push one node too hard and you back up another, or you rob respiration of α-KG and the plant wilts. We constrain the whole set of edits with a **genome-scale metabolic model** solved by **flux-balance analysis (FBA)**: every reaction in the organism's metabolism is written as a stoichiometric constraint, growth is held above a viability floor, and the solver is asked to **maximize steady-state free cytosolic glutamate** subject to all of it. The model does three useful things at once. It identifies which drain is actually rate-limiting in *this* organism (in tomato fruit, FBA puts GAD attenuation far ahead of GDH over-expression — the pool is lost downstream, not under-made). It sets the *ceiling*: how much free glutamate the organism can carry before growth or redox balance breaks, which is why we

target the tomato to around 2,600 rather than an unsustainable maximum. And it finds the **collateral**: the edit that raises glutamate also nudges the α-KG pool, which feeds back on respiration and, as it happens, on the nucleotide pathway of the next chapter — so the two halves of umami are modeled together, not in isolation.

Two neighbors of glutamate deserve a mention, because savor is never a single amino acid in isolation. **Aspartate** — glutamate's four-carbon cousin, made by the very transamination that can drain glutamate — is itself weakly umami, a minor savory note riding alongside the main one, and the aspartate-to-glutamate ratio is part of what gives a savory profile its particular character. And the **acid balance** matters as much as the amino acid. A tomato's free glutamate is read against its citric and malic acids and its sugars, and the same glutamate tastes sharp in an acidic, low-sugar fruit and round in a balanced one — which is why restoring umami in the tomato is done *alongside*, never instead of, holding the sugar-and-acid frame the fruit reads its savor within. The glutamate engine makes the note; the fruit's acidity and sweetness decide the room it is heard in, and a savory organism designed without that frame in mind can be loud and still unpleasant.

The result, in the restored tomato, is a free-glutamate pool that climbs to vine-ripe-heirloom loudness and then holds there batch after batch, grown rather than seasoned, with the GABA the plant would have made quietly forgone. The note, made at the source. Next, its gain.

Chapter 6 · The ribonucleotide route — *the gain on the note.*

PATHWAY · the purine branch · IMP → GMP · 5′-phosphodiesterase

Glutamate is the note; the 5′-ribonucleotides are the gain. From Chapter 3 we know that a trace of IMP or GMP can multiply perceived umami several-fold by wedging the receptor's flytrap shut. The question this chapter answers is where those nucleotides come from biochemically, and why getting an organism to carry them *in the free, tasteable form* is harder than getting it to carry glutamate — hard enough that it reshapes how we engineer the savory organisms entirely.

The savory nucleotides are purines, and they are all children of one molecule: **inosine monophosphate, IMP**, the branch-point of purine nucleotide biosynthesis. The cell builds IMP atom by atom on a ribose-phosphate scaffold through a ten-step *de novo* pathway, and from IMP the two roads diverge. That pathway is worth naming at its head and its hinge, because both are regulatory. It begins with **PRPP** — 5-phosphoribosyl-1-pyrophosphate, the activated ribose-phosphate that is also the entry point for every nucleotide the cell makes — and its committed, rate-governing step is **PRPP amidotransferase**, which transfers an amide nitrogen from glutamine onto PRPP to start the purine ring. That enzyme is **feedback-inhibited by the end products themselves** — AMP, GMP, and IMP bind it and shut it down — which is the cell's way of refusing to overbuild purines. Then nine more steps, drawing in glycine, aspartate, two one-carbon units from folate, carbon dioxide, and the better part of **five ATP**, assemble IMP. The cost is the point: de novo purine synthesis is one of metabolism's most expensive constructions, and no cell will run it to pile up tasteable guanylate it does not need for RNA. This is the thermodynamic face of the growth-competition the flux model flagged, and it is why accumulation, not synthesis, is the industrial answer.

There is also a cheaper route the cell prefers when it can — the **salvage pathway**, in which the enzymes HGPRT and APRT reattach free purine bases to PRPP to remake nucleotides, recycling rather than rebuilding. Salvage matters to us because it governs how a processed biomass's liberated bases behave, and because the industrial production of the savory nucleotides splits along exactly this seam. Commercial 5′-IMP and 5′-GMP are made two ways: by **microbial fermentation**, in which engineered *Corynebacterium* or *Bacillus* strains, their feedback inhibition deliberately relieved, overflow **inosine** or **guanosine** into the broth to be phosphorylated at the 5′ position afterward; or by **RNA hydrolysis**, in which a **5′-phosphodiesterase** — nuclease P1, classically from *Penicillium citrinum* — digests yeast RNA into the four 5′-mononucleotides, with AMP deaminase finishing the adenylate to the stronger inosinate. Both routes are the staged-release principle industrialized, and both are templates our engineered yeast internalizes: grow the reservoir, then cut it to savor.

There is a fragility to guard, and it changes how savory ingredients are handled. A 5′-nucleotide tastes only while its phosphate is attached; strip the phosphate and the molecule goes quiet. Foods carry **phosphatases** that do exactly that, and a vegetable's own enzymes will slowly dephosphorylate the guanylate it holds — one reason a long-stored or carelessly handled ingredient loses its synergy even as its glutamate survives untouched. The traditional defense is heat: **blanching** a vegetable inactivates its phosphatases before they can strip the nucleotide, locking the savory potential in, and the same logic governs how a savory concentrate is pasteurized. For our organisms the protection is built into the process — the controlled hold that liberates the nucleotide also denatures the enzymes that would degrade it, so the guanylate is made and shielded in one pass. A nucleotide you liberate and fail to protect is a nucleotide you lose, and the difference between a savory ingredient that holds its loudness on the shelf and one that quietly fades is, very often, a phosphatase that someone did or did not switch off in time. One road aminates IMP — through adenylosuccinate, via adenylosuccinate synthetase and lyase — to **AMP**, the adenine nucleotide. The other road oxidizes IMP at the hands of **IMP dehydrogenase** to **XMP** (xanthosine monophosphate), which **GMP synthetase** then aminates to **GMP**, the guanine nucleotide and the most potent umami synergist of the three. IMP, AMP, GMP: the inosinate of meat and bonito, the adenylate that is weakly synerg/sweet, and the guanylate of the mushrooms. Our interest is in the first and third.

Here is the complication that governs everything. In a living cell, these nucleotides are not condiments; they are the **currency of RNA, DNA, and energy.** A cell that accumulated large pools of free GMP would be a cell drowning in a building block it is supposed to polymerize immediately. Flux-balance analysis makes the tension explicit: pull flux toward free guanylate and you compete directly with growth, because every molecule of GMP sitting free and tasteable is a molecule not being spent on ribosomes. You cannot simply over-express the purine pathway in a plant fruit and harvest guanylate; the plant will either incorporate it or you will make it sick. Nature's own savory nucleotide sources all solve this the same way, and so do we: **not by over-synthesis, but by staged release.**



FIG. 5 — Specimen isolation. The guanylate a mushroom carries is not made free; it is liberated from the organism's own RNA by a phosphodiesterase during drying and gentle heat — a release engineered, not a pool over-built.

Consider how the kitchen has always done it. In meat and fish, the savory nucleotide is **post-mortem**: once the animal dies, its ATP cascades down — ATP → ADP → AMP — and **AMP deaminase** converts AMP to IMP, which accumulates in the resting muscle over hours to days. This is why aged beef and properly cured bonito are so much more savory than fresh; the inosinate is a product of controlled autolysis, not of life. In mushrooms, the savory nucleotide is **released from RNA**: the dried shiitake is packed with ribosomal RNA, and on rehydration and gentle heating, the mushroom's own ribonucleases and a **5′-phosphodiesterase** cleave that RNA into 5′-mononucleotides — among them the guanylate that makes a shiitake broth so deep. Soak a dried shiitake in cold water and warm it slowly and you are running an enzymatic reactor: the GMP you taste did not exist as a free pool in the living mushroom; it was liberated from the organism's own genetic machinery on the way to your pot.

That mechanism is our template. For the nucleotide half of a savory organism we engineer two things, and neither is "make more free GMP in the living cell." First, we build **RNA-dense biomass** — organisms, especially the yeast of Chapter 11, grown fast and rich so that a large fraction of their dry mass is ribosomal RNA, a reservoir of bound guanylate and adenylate. Second, we tune the **release enzymes** — the 5′-phosphodiesterase that cleaves RNA to 5′-nucleotides and the AMP deaminase that converts the resulting adenylate to the far more savory inosinate — so that a defined processing step (autolysis, drying, a controlled warm hold) converts the reservoir into a free nucleotide pool on demand. The nucleotide is grown as RNA and *harvested* as flavor.

NUCLEOTIDE	BUILT BY	FREED BY	SAVORY ROLE
IMP (inosinate)	de novo purine pathway; AMP deaminase from AMP	post-mortem/autolytic AMP deaminase	strong synergist (β ≈ 1,218)
GMP (guanylate)	IMP → XMP → GMP (IMPDH, GMP synthetase)	5′-phosphodiesterase on RNA	strongest synergist (β ≈ 2,800)
AMP (adenylate)	IMP amination	RNA/ATP breakdown	weak synergist; deaminate to IMP

The payoff is the ability to put the gain and the note in **one organism**. A designed yeast grown RNA-rich and finished with a tuned phosphodiesterase delivers free glutamate *and* free guanylate from a single ingredient — the receptor's two binding sites served by one grown thing, the synergy of Chapter 3 built in rather than blended.

There is a quiet dietary signature in *which* nucleotide an organism carries, and it matters to the plant-forward kitchen this library serves. **Inosinate is the animal nucleotide** — the post-mortem product of muscle, the savor of meat and fish — while **guanylate is the plant-and-fungal one**, released from the RNA of mushrooms and microbes. A savory system built to stand without animal ingredients is therefore, at the molecular level, a **GMP story**: it reaches the synergy through the mushroom's and the yeast's guanylate rather than through muscle's inosinate. And because GMP's synergy constant is more than twice IMP's, the plant route is not the weaker

one — it is, if anything, the louder. This is a genuinely hopeful fact for anyone building savor from plants: the strongest synergist on the receptor is the one the fungal and microbial kingdoms make, and the vegetarian broth that marries kelp to shiitake is reaching a potentiation a meat stock cannot match. The nucleotide an organism carries is a kind of fingerprint of its kingdom, and the plant kingdom's fingerprint, read on the receptor, is the strong one. And because the release is enzymatic and staged, it is reproducible to the library's seven-percent bar: the same biomass, processed the same way, yields the same nucleotide pool, which is the difference between a broth that is reliably deep and one that is deep when you are lucky. The note is grown; the gain is grown beside it. What remains is the body.

Chapter 7 · The γ-glutamyl pathway — *where body is built.*

PATHWAY · glutathione synthesis · γ-glutamyl transpeptidase · the kokumi peptides

Kokumi is a tripeptide problem, and it is the most elegant of the three savory biosyntheses, because the molecule that gives the most body — γ-glutamyl-valyl-glycine — is a near-neighbor of a molecule every living cell already makes in quantity: **glutathione**. The plant is not short of the machinery. It is short of a nudge, and the γ-glutamyl pathway is the study of where to apply it.

Glutathione — γ-Glu-Cys-Gly — is the cell's master antioxidant, present in nearly every compartment at millimolar concentration, and it is assembled in two ATP-dependent steps that are the spine of all kokumi chemistry. First, **γ-glutamylcysteine ligase** (GCL, also called γ-glutamylcysteine synthetase, γ-GCS) joins glutamate to cysteine through the glutamate *side-chain* carboxyl — forming the signature **γ-linkage** — to make the dipeptide **γ-Glu-Cys**. This is the committed, rate-limiting step, feedback-inhibited by glutathione itself, and it is the first lever: relieve its feedback inhibition or raise its expression and the whole γ-glutamyl pool grows. Second, **glutathione synthetase** adds glycine to γ-Glu-Cys to complete glutathione. Two enzymes, two peptide bonds, one of them the unusual γ-bond that makes these molecules taste of body rather than of their constituent amino acids.

Glutathione itself is a kokumi compound — it is the benchmark, potency 1.0 on the CaSR ladder of Chapter 2. But the prize sits one transfer away. The enzyme **γ-glutamyl transpeptidase (GGT)** takes the γ-glutamyl group from glutathione and transfers it onto an acceptor amino acid or peptide — this is its native job in glutathione turnover and recycling across the cell membrane.

How GGT does it is the whole of why it is designable. GGT is an **N-terminal nucleophile (Ntn) hydrolase**: it is made as a single inactive chain and then *autocatalytically cleaves itself* into a large and a small subunit, a cut that exposes a new N-terminal **threonine** whose side-chain hydroxyl becomes the catalytic nucleophile. That threonine attacks the γ-glutamyl amide of glutathione, releasing the cysteinylglycine and forming a covalent **γ-glutamyl-enzyme intermediate** — a loaded enzyme, holding the γ-glutamyl group and waiting for something to hand it to. Here the fork decides everything. If a water molecule resolves the intermediate, the γ-glutamyl group is simply hydrolyzed to free glutamate — the unproductive route. If instead an **acceptor amino acid** reaches the active site first, the group is transferred onto it — **transpeptidation**, the productive route that builds a new γ-glutamyl peptide. Kokumi engineering is, almost entirely, the art of biasing that fork: the N428D edit of Chapter 8 raises transpeptidation over hydrolysis, and the T391S edit widens

the acceptor pocket so that **valylglycine** out-competes water, steering the loaded enzyme toward γ-Glu-Val-Gly rather than toward waste.

The committed step upstream, **GCL**, is designable for a different reason. GCL is an ATP-grasp ligase that forms the unusual γ-bond through an **acyl-phosphate intermediate** — it phosphorylates glutamate's side-chain carboxyl, then lets cysteine's amine displace the phosphate — and it is held in check by **feedback inhibition from glutathione itself**, which binds at the glutamate site and throttles the whole reservoir once the cell has enough. The E97Q edit relieves that inhibition, letting the reservoir grow past the cell's homeostatic set-point so GGT has surplus to skim. And the reason we must *skim* rather than drain is the molecule's day job: glutathione is the cell's redox buffer, kept in its reduced form by **glutathione reductase** spending NADPH, its reduced-to-oxidized ratio a readout of the cell's health. The Meister **γ-glutamyl cycle** — synthesis by GCL and glutathione synthetase, turnover by GGT and the dipeptidases — is a loop the organism lives inside. We are not allowed to break the loop. We are allowed to run it faster and lean its surplus toward the tongue, which is the gentlest intervention in this book precisely because the cell was already doing nearly all of it.

A practical virtue of the γ-glutamyl peptides earns a final note: they are **robust**. Where aroma compounds are volatile and fragile — driven off by heat, lost to time, oxidized in storage — the kokumi peptides are small, water-soluble, and chemically stable, surviving cooking, pasteurization, and months on a shelf largely intact. The γ-linkage that makes them taste of body also makes them resistant to the ordinary peptidases that would cleave a normal peptide bond, so the body an organism carries is body that lasts. For a manufacturer this is the difference between a flavor that must be protected and one that can simply be built in and relied upon: the kokumi designed into a yeast or an allium is still there after the retort, after the spray-dryer, after the warehouse. Among the three savory chemistries, the body is the most durable — fitting, since it is also the one whose whole job is to make a taste last. Hand GGT the right acceptor and it builds new γ-glutamyl peptides: give it **valylglycine** (Val-Gly) and it makes **γ-Glu-Val-Gly**, the roughly thirteen-fold-potent tripeptide that is the most efficient body molecule we know. The kokumi pathway, then, is three enzymes working in series — GCL to commit, glutathione synthetase to complete the reservoir, GGT to redirect that reservoir toward the high-potency target — and designing kokumi is designing the balance among them.



FIG. 6 — Onion in cross-section. The concentric rings are a store of γ-glutamyl peptides; sprouting and aging activate the allium transpeptidase that turns that store into perceptible body.

The alliums show why this pathway is worth engineering rather than extracting, and they show it beautifully, because garlic and onion are already professional γ-glutamyl chemists. In an allium bulb, γ-glutamyl peptides are the **storage and transport form** of sulfur and nitrogen — γ-glutamyl-S-alk(en)yl-cysteines, warehoused in the dormant clove. When the bulb sprouts or is aged, an allium γ-glutamyl transpeptidase mobilizes that store, cleaving and transferring the γ-glutamyl groups and releasing, among other things, the γ-Glu-Cys and S-allyl-cysteine that give aged garlic its rounded, lingering depth. This is the biochemistry a cook exploits, without naming it, every time a confit of garlic tastes sweeter and fuller than a raw clove: aging has run the transpeptidase. Our contribution is to bias that native enzyme — through expression and, where the structure allows, through acceptor-specificity edits — toward the valylglycine acceptor, so the bulb accumulates γ-Glu-Val-Gly rather than only its ancestral sulfur peptides. Same pathway, redirected toward the molecule the CaSR reads loudest.

STEP	ENZYME	REACTION	KOKUMI RELEVANCE
Commit	GCL (γ-GCS)	Glu + Cys → γ-Glu-Cys	rate-limiting; relieve feedback
Complete	glutathione synthetase	γ-Glu-Cys + Gly → glutathione	builds the reservoir (potency 1.0)
Redirect	GGT	glutathione + Val-Gly → γ-Glu-Val-Gly	the ~12.8× target
Mobilize (allium)	γ-glutamyl transpeptidase	γ-Glu-S-alkyl-Cys store → free peptides	the "aged garlic" depth

Flux-balance analysis keeps this honest in a way peculiar to kokumi. Glutathione is not a flavor the cell is willing to lose — it is the antioxidant it lives by — so pulling flux toward a tasteable γ-glutamyl pool must not strip the organism's redox defense. The model is solved with glutathione held above a physiological floor and the *excess* flux directed to γ-Glu-Val-Gly, which is why the kokumi organisms are engineered to over-produce the pathway rather than to divert the baseline: we raise GCL to grow the whole reservoir, then let GGT skim the surplus to the target, leaving the cell its antioxidant intact. The γ-glutamyl peptides are the designable core of kokumi, but they are not its only source, and the fuller picture helps a formulator know what they are building toward. **Aged cheese** generates γ-glutamyl peptides across months of ripening, part of why a well-aged parmesan tastes round as well as loud; the **scallop**

and certain shellfish carry a natural kokumi chefs have exploited forever without naming it; and fermented and long-cooked foods generate the peptides in place, the body a braise builds being, chemically, a γ-glutamyl synthesis run slow. Beyond the peptides, CaSR answers to other ligands that contribute body — **calcium itself**, as Chapter 2 noted, along with certain polyamines and hydrophobic amino acids — so a food's kokumi is really a chord of CaSR agonists, of which the γ-glutamyl peptides are the loudest and by far the most tunable string. When we engineer an organism for body, we are raising the one contribution we can design precisely, inside a signal the food was always assembling from several sources at once. That is why the engineered kokumi reads as *natural* rather than *added*: it is the same molecule, in the same receptor, that a long braise and an aged cheese were quietly making all along.

It is the gentlest of the three savory interventions — a surplus skimmed, not a function sacrificed — and it yields the axis that makes the other two last.

With all three pathways named — the glutamate engine for the note, the purine branch for its gain, the γ-glutamyl series for the body — the organism has everything savor requires. What remains is the method that decides exactly which residue to change, and how far.

Chapter 8 · The method — *how we decide the edit.*

PLATFORM · protein language models · structure prediction · flux-balance analysis

The three pathways of the last three chapters give us a map of *where* to intervene. They do not tell us *how* — which residue of which enzyme to change, by how much, to raise an activity or stabilize a fold without breaking the protein. That decision is the platform, and it is the same engine, unchanged, that designed the heat of the flagship peppers. Savor is where we most want to show that the method is not a pepper method but a flavor method, so this chapter walks the engine once, on the savory enzymes.

The platform is three models in series. The first is a **protein language model** — an ESM2-class network of hundreds of millions of parameters, trained on the evolutionary record of protein sequences until it has learned, implicitly, the grammar of what a functional protein can be. Hand it an enzyme sequence — a glutamate dehydrogenase, a γ-glutamyl transpeptidase — and ask it to score every possible single-residue substitution by its predicted effect on function and on fold stability. The model has never seen our enzyme; it does not need to. It has seen enough of the family that it can tell a tolerated substitution from a catastrophic one, and rank the plausibly *beneficial* ones — the handful that might tighten a binding pocket or stiffen a loop without unfolding the protein. This turns an impossible search — twenty choices at each of several hundred positions — into a short list.

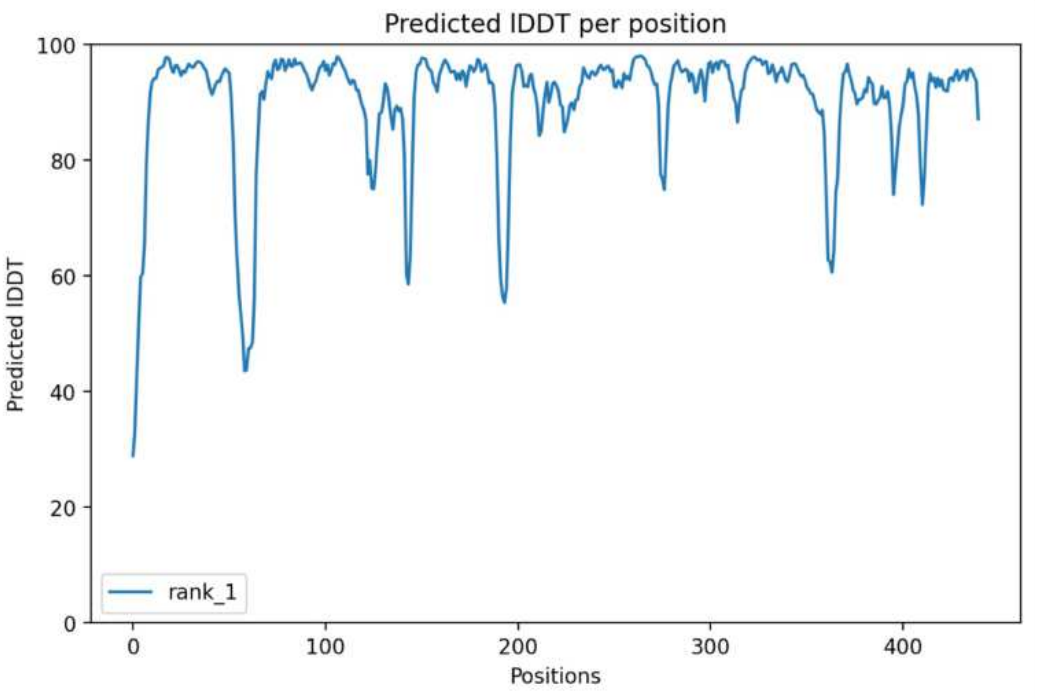

FIG. 7 — The platform's structure model, coloured by per-residue confidence (pLDDT); the folded core and active-site cleft are high-confidence, the flexible termini low. The same pipeline that produced this plot for the capsaicin acyltransferase runs on every savory enzyme we design.

The second model is **structure prediction**. The language model's short list is proposed against a predicted three-dimensional fold, scored residue by residue for confidence — the **pLDDT** metric that

colours the plot above, high in the ordered catalytic core, low in the floppy termini. Structure tells us what sequence alone cannot: whether a proposed substitution sits in the active-site cleft where it will touch catalysis, on a surface loop where it will not, or across a domain interface where it might shift the whole conformational equilibrium. For the reversible enzymes of savor this matters especially — GDH's synthesis-versus-breakdown balance and GGT's acceptor preference both live in the geometry of the active site — so we read the structure for where the equilibrium can be nudged, not merely for where the protein is strong. A companion plot, the **predicted aligned error** map, tells us how confident the model is about the *relative* positions of domains, which is the confidence that matters when an edit is meant to change how two lobes move against each other.

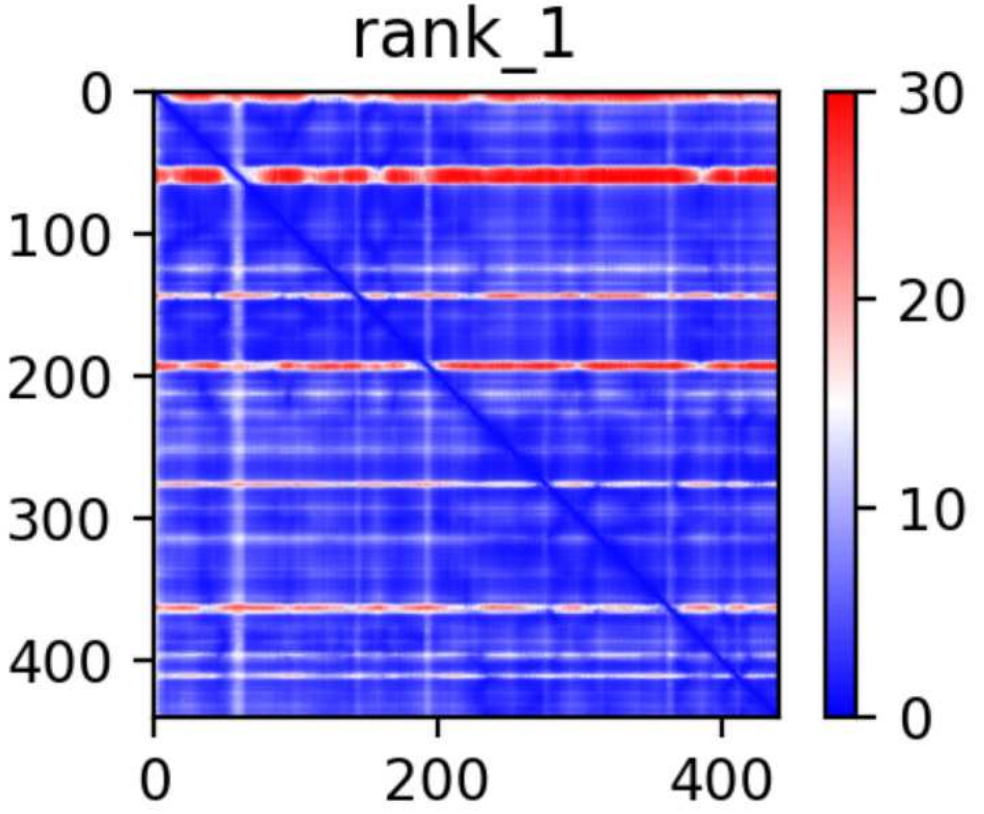

FIG. 7b — The predicted aligned error map for the same model. Low error (dark) within a domain means its fold is trusted; low error *between* domains means their relative placement is trusted — the reading that tells us whether an edit at a hinge will move the lobes the way we intend.

Three technical points make the pipeline more than a picture. First, the language model scores a substitution **zero-shot** — without ever being trained on our enzyme — by a *masked-marginal* comparison: it asks how much more or less likely the mutant residue is than the wild-type one, given the rest of the sequence as context, and a

mutation the evolutionary record finds implausible is flagged before any structure is built. Second, each surviving candidate is **docked**: the enzyme's substrate — glutamate, a nucleotide, a γ-glutamyl peptide — is placed into the predicted active site and scored to confirm the edit tightens or re-shapes the pocket without occluding it, and a short **molecular-dynamics** run checks that the edited fold still holds at the plant's physiological temperature rather than unravelling a few nanoseconds in. Third, the flux model's objective is written plainly as a linear program: **maximize v_target subject to S·v = 0, v_biomass ≥ floor, and the thermodynamic flux bounds** — where S is the stoichiometric matrix of the whole metabolism, v the vector of reaction rates, v_target the free-glutamate (or guanylate, or γ-Glu-Val-Gly) flux, and the biomass floor the viability constraint that keeps the solution a living plant. Only when a residue passes all three — plausible to the language model, well-placed and stable in the structure, productive in the flux program — is it written.

The writing itself is matched to the edit. A single stabilizing substitution is installed with a **base editor**, which rewrites one base without a double-strand break; a larger change goes in by **homology-directed knock-in**; and a knockdown like SlGAD3's is a **promoter or untranslated-region edit** that dials expression down rather than deleting the gene, because a dialled gene holds its level across batches where a deletion can be compensated unpredictably. Conservatism is not timidity here; it is how the seven-percent batch bar is met at the level of the DNA.

The third model is the **genome-scale metabolic model** solved by **flux-balance analysis**, which we have already met in each pathway chapter and which has the final word. A protein edit that triples an enzyme's activity is worthless if that enzyme was not the rate-limiting step, and worse than worthless if the extra flux backs up into a toxic intermediate or robs the plant of a metabolite it needs to live. FBA sits above the protein work and answers the systems question: given the whole metabolism, *will this edit actually raise the free glutamate (or guanylate, or γ-Glu-Val-Gly) pool at the organism level, and what does it cost the organism?* Only edits that pass all three filters — plausible to the language model, well-placed in the structure, productive in the flux model — are written into the plant.

They are written with **CRISPR**, and the writing is deliberately conservative. A knockdown of SlGAD3 is a promoter or coding-sequence edit that turns an activity down, not out. A stabilizing substitution in a dehydrogenase is a single base change that the language model ranked and the structure placed. We prefer the smallest edit that moves the number, because small edits are the ones that hold their seven-percent batch variation and survive into a real, growable, harvestable plant. Below is a representative set of designed edits across the savory enzymes — the kind of short, ranked list the platform returns, with the predicted change in folding free energy (ΔΔG, negative = stabilizing) that the structure model assigns:

ENZYME	EDIT	DESIGNED EFFECT	ΔΔG (KCAL/MOL)
GDH1 (glutamate dehydrogenase)	A166V	shifts equilibrium toward amination	−1.6
GDH1	R241K	stabilizes NAD(P)H-bound state	−1.2
GGT (γ-glutamyl transpeptidase)	T391S	widens acceptor pocket for Val-Gly	−0.8
GGT	N428D	raises transpeptidation over hydrolysis	−1.4
GCL (γ-Glu-Cys ligase)	E97Q	relieves glutathione feedback inhibition	−0.6
5′-phosphodiesterase	L212F	raises thermostability for the warm-hold release	−2.0

Note what the table is not: it is not a set of guesses. Each row is a residue the language model nominated, the structure model placed and scored, and the flux model predicted would move the organism-level pool — and each is small enough to hold its profile. This is the whole of our claim to the word *design*. We are not breeding and hoping the genotype follows the phenotype, the way heat was chased for ten thousand years. We are reading the enzyme, placing the change, modeling the consequence, and writing the one edit that the three models agree will raise the savor and keep the plant alive. The method proven on a pepper, turned on a broth.

The loop closes in the plant, not the model. A designed edit is a *prediction* until the edited organism is grown and its savor read on the assays of Chapter 4 — EC_{50} measured, profile scored, batch variation checked — and that measurement feeds back to correct the models for the next design. This is why the construct counts are large: each organism's final profile is the survivor of many designed candidates, most of them screened out not because the biology failed but because the number landed a little off target or the variation ran a little wide. The three models make the search tractable, cutting an astronomically large space of possible edits down to a short list worth growing; the greenhouse makes it true. A design the models love and the plant does not is, in this library, a design that lost — and the discipline of letting the plant have the last word is the whole reason the numbers in this book can be trusted.

Across the savory library this engine has produced the construct counts the next part is built from — the two thousand and forty tomato constructs, the hundreds in soybean and yeast, the designed alliums — each one a plant whose savory profile was decided by three models before a seed was ever sown. We go now to the organisms themselves.



PART III OF V

The Organisms

tomato · soybean · yeast · koji · shiitake · the alliums — savor grown, not seasoned

Chapter 9 · Tomato — *the savor bred out, grown back.*

ORGANISM · 2,040 constructs · axes S/U/F/M · difficulty: Easy · IFI-U 300 → 2,600

If there is a founding grievance in the savory half of our library, it is the supermarket tomato. More than any other plant, the tomato is where modern breeding traded flavor for logistics in front of everyone and called it progress — and more than any other plant, it is where the loss is reversible, because the tomato never lost the *genes* for savor. It lost the *behavior*. Our tomato work is the largest single program in the savory library — **2,040 constructs** — and it is, mercifully, the easiest: a native-pathway restoration in a plant whose genome we know cold, a model Solanaceous organism with no exotic machinery required. We are not adding savor to the tomato. We are letting it keep what it was always going to make.

Recall the number from Chapter 1: a vine-ripe heirloom carries free glutamate enough to read near a thousand on the Index; the shipping-bred commodity tomato reads near **300**, having shed roughly ninety-three percent of its peak free-glutamate pool. Where did it go? Two places, both of which Chapters 5 traced. First, it was never allowed to *finish*. The free-glutamate pool climbs steeply in the last days of ripening on the vine, as GDH tips toward synthesis and the fruit's metabolism winds down — and a tomato picked green for the truck, gassed to color in a warehouse, simply never runs that final program. Second, in the commodity genetics, more of the glutamate that *is* made gets decarboxylated by **SIGAD3** into tasteless GABA and more gets held in protein, so even a commodity tomato left to ripen fully underperforms the heirloom it descended from.

The deepest cut was an accident of convenience, and it is worth telling because it is the whole parable of the modern tomato. Mid-century breeders selected a trait called **uniform ripening** — an unripe fruit of an even, pale green that turned red all at once, so a field could be machine-picked in a single pass. The trait is a loss-of-function in a transcription factor, **GLK2**, and GLK2's day job is to build the chloroplasts in the fruit's own flesh. Break it and the tomato ripens evenly, yes — but the fruit also loses much of its own photosynthesis, and with it the sugar *and* the carbon skeletons that feed the glutamate engine. An entire industry selected, in one gene, for a prettier unripe fruit and against the savor and sweetness of the ripe one, and did not know it until the molecular work was done. The savor was collateral damage of a harvest schedule. We reverse the schedule's cost at the genome, restoring the fruit's capacity to finish — to run its glutamate program to the end — without asking the grower to pick by hand.

The accumulation is a developmental curve, and timing is everything. Through the stages a tomato passes on the vine — **mature green, breaker, turning, pink, red-ripe** — the free-glutamate

pool stays low until the breaker stage and then climbs steeply into full ripeness, while the GABA that dominated early falls away. A fruit picked at mature green for shipping is picked *before the curve begins*, and no amount of warehouse ethylene reruns the lost days of vine metabolism. The restored tomato is engineered so that more of that climb happens, happens faster, and survives to the table:

TOMATO	FREE GLUTAMATE (MG/100 G)	~BRIX (SUGAR)	IFI-U
Commodity, shipping-bred	20–60	~4.0	~300
Vine-ripe heirloom	150–250	~6.5	~1,000
Restored (engineered)	250–400	~7.0	~2,600

Read the last row against the formula of Chapter 4. The restored tomato carries more free glutamate than the heirloom, yes — but its IFI-U of 2,600 is more than twice the heirloom's 1,000 because glutamate alone is not the whole score: the restored fruit also carries a **staged nucleotide release** that potentiates its own glutamate on cooking, and a **kokumi floor** from the engineered γ-glutamyl reservoir. The number is loudness, and loudness is glutamate multiplied by its synergists, not glutamate counted.

There is a last piece to the restored tomato, and it keeps the project honest about what *flavor* is. Umami is the subject of this book, but a tomato that "tastes like a tomato should" is not umami alone — it is umami read inside a frame of **sugar and acid** and lit by a cloud of **aroma volatiles**, the green, sweet, faintly smoky compounds a ripe tomato releases that the nose reads as "tomato" before the tongue reports anything. Breeding for the truck dimmed those volatiles alongside the glutamate, and a credible restoration holds the whole frame: the sugar-acid balance of the previous chapter's closing note, the volatile profile that tells the brain *ripe*, and the free-glutamate loudness that is our particular lever. We lead with savor because savor is what the Index measures and what the kitchen most often lacks — but we grow the fruit as a whole sensory object, because a loud glutamate inside a flat, odorless tomato would be a laboratory curiosity, not the thing a cook reaches for. The restoration is of a tomato, not of a number; the number is only how we know we have it.

And the restored tomato is grown the way the rest of the library is grown — with the whole plant in view, in conditions tuned to let it finish. Controlled-environment cultivation lets the fruit run its full ripening program on the vine without the shipping compromise that drained its ancestors, so the glutamate climb of the breaker-to-red window completes where it is measured rather than being cut short in a warehouse. The result is the plainest promise the organism shelf can make, and the one the kitchen has waited fifty years for: a tomato that tastes the way a tomato is supposed to — loud, round, and finished — not as nostalgia but as a profile, grown on purpose and read on the Index. It is the restoration thesis made edible, and it is where a reader who has followed the chemistry this far gets, finally, to taste the argument.



FIG. 8 — The restored tomato, whole. The savor is concentrated in the locular jelly around the seeds; a tomato engineered to finish its glutamate program carries the loudness of a vine-ripe heirloom on a plant that still ships.

The restoration is a stack of edits, each from the pathways already named, each shown by the flux model to move the organism-level pool. The highest-leverage one is the **attenuation of SlGAD3** — turning down, not out, the decarboxylase that drains glutamate to GABA, so the pool the plant builds in ripening survives to the table. Layered onto it: a **GDH1** shift toward amination (the A166V and R241K edits of Chapter 8), lengthening and strengthening the final ripening climb; a measured trim of the export aminotransferases to keep nitrogen resident; and — the move that makes the tomato a true U×K organism rather than only a U one — a modest **GCL** up-regulation to grow the γ-glutamyl reservoir, giving the fruit a kokumi floor it never had. The result climbs from 300 to a restored and locked **2,600 on IFI-U**, with a real, if secondary, IFI-K. And because the tomato also carries sweetness natively, the restored fruit reads as the balanced savory-sweet thing a tomato is supposed to be:

IFI PROFILE

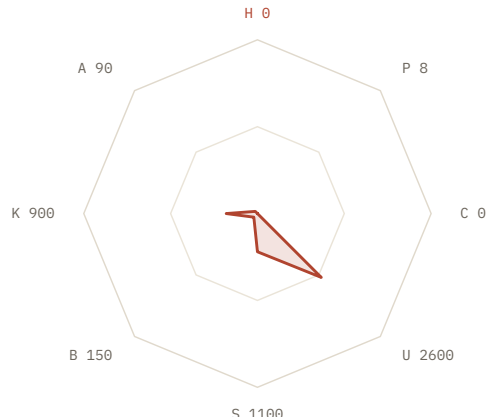

IFI-H	0
IFI-P	8
IFI-C	0
IFI-U	2,600
IFI-S	1,100
IFI-B	150
IFI-K	900
IFI-A	90

RESTORATION EDIT	PATHWAY	EFFECT	CONTRIBUTION
SlGAD3 attenuation	glutamate engine	less glutamate → GABA	the bulk of IFI-U
GDH1 A166V / R241K	glutamate engine	longer ripening climb	sustains the pool
aminotransferase trim	glutamate engine	nitrogen stays resident	marginal IFI-U
GCL up-regulation	γ-glutamyl	grows kokumi reservoir	the IFI-K floor
staged nucleotide release	purine branch	processing-liberated IMP	synergy on cooking

What does a restored tomato *do* that a commodity one cannot? It cooks like a tomato that simmered all day, because it starts where an all-day simmer was trying to get. A commodity-tomato sauce reaches for depth and never arrives; the cook adds salt, then sugar, then a long reduction, then — in a factory — a drum of glutamate and a pinch of inosinate, chasing with additives the savor the fruit should have carried in. The restored tomato carries it in: free glutamate that speaks on its own and potentiates hard the moment a staged nucleotide release or a seared surface lends it IMP, plus a kokumi floor that rounds the whole thing. It is the single most useful organism in this book for a savory manufacturer, because the tomato is already in everyone's formulation. Swap the input and the depth arrives without a new additive line. This is the restoration thesis in its plainest form: the flavor was never gone, only turned off, and turning it back on is cheaper — and cleaner — than buying it back by the drum.

Chapter 10 · Soybean — *one seed, both axes.*

ORGANISM · 360 constructs · axes U/K · difficulty: Frontier · the dual-axis flagship

Of all the organisms in the savory library, the soybean is the one that carries **both axes natively**, and that fact makes it the quiet center of this book. Umami *and* kokumi, note *and* body, in a single seed — which is why, across every cuisine that discovered it, the soybean became the savory foundation: soy sauce, miso, tempeh, tofu, natto, fermented bean pastes from a dozen traditions, all of them extracting from one legume the complete savory stack that other ingredients deliver only in pairs. The soybean is a **frontier** organism in our difficulty grading — a legume with a complex seed metabolism and a slower transformation cycle than the easy Solanaceae — and we grade it frontier not because the pathways are strange but because the plant is demanding. The **360 constructs** here are hard-won, and each one matters more for it.

The soybean's umami is latent, locked in its protein. A raw soybean is not loud; its glutamate is almost entirely bound in **glycinin** and **β-conglycinin**, the storage proteins that make the seed forty percent protein by dry weight — a reservoir of glutamate residues with nothing free to taste. The soybean's savor is therefore a **release** story even more than the tomato's is a synthesis story: the free glutamate that makes soy sauce profound is proteolytic, liberated by the koji enzymes of the next chapter over months of fermentation. This is the soybean's native genius and its native limitation — it holds enormous savory potential and hands almost none of it over without help.

The reservoir has a structure worth knowing, because its structure sets how it releases. The seed's protein is dominated by two storage proteins: **glycinin**, the 11S legumin, and **β-conglycinin**, the 7S vicilin, and the ratio between them — the 11S/7S balance — governs both how the seed gels and how readily proteolysis liberates its amino acids. Glycinin is comparatively rich in glutamate and glutamine residues, so a seed bred toward glycinin is a seed whose proteolysis yields more of the savory amino acid; this is one of the quieter levers in the soybean program, a composition edit rather than a pathway one. And the soybean carries an advantage no other savory organism we grow can match: it fixes its own nitrogen. In symbiosis with **Bradyrhizobium** bacteria in its root nodules, the soybean pulls nitrogen straight from the air and feeds it into the GS/GOGAT cycle of Chapter 5 — a built-in fertilizer line running directly to the glutamate engine. The legume that carries both savory axes is also the legume that supplies its own nitrogen to build them; the economics of growing savor rarely align so neatly with the biology.

The soybean's cultural reach is the proof of its savory completeness, because almost every great fermented savory ingredient on earth is, underneath, a way of releasing this one seed:

FERMENT	MICROBE	WHAT IT RELEASES
Shoyu (soy sauce)	*A. oryzae* koji → *Z. rouxii*, *T. halophilus*	proteolytic glutamate + aged depth
Miso	*A. oryzae* koji	glutamate + γ-glutamyl body, by time
Tempeh	*Rhizopus oligosporus*	mild proteolysis, nutty peptides
Natto	*Bacillus subtilis* var. *natto*	peptides, γ-polyglutamate, pungency
Douchi / doubanjiang	koji + bacteria	deep roasted umami

Our engineering serves every one of these by improving their shared substrate: a seed that arrives at the ferment already partly loud, biased toward glutamate-rich storage protein and toward the γ-Glu-Val-Gly end of the kokumi ladder, with the beany off-notes of its native lipoxygenases trimmed so the depth reads clean. We also tune the isoflavone and saponin load — the genistein and the bitter triterpenoids that give raw soy its edge — downward, so that the savor the ferment liberates is not fighting a bitterness the seed brought with it. The soybean is the organism where the most traditions meet our method, and improving the seed improves all of them at once.

Two more roles make the soybean indispensable to this book in particular. The first is as the **protein base of the meat analog** of Chapter 16: the same seed that carries both savory axes also supplies, in its glycinin and conglycinin, the structural protein a textured analog is built from — so the soybean can be, in one product, both the savor and the substance, the depth and the bite. The second is the lesson of **natto**, where *Bacillus subtilis* spins the beans into sticky strands of **γ-polyglutamic acid**, a polymer of the very amino acid this whole book is about, linked through the same γ-bond that makes the kokumi peptides. Natto is divisive for its pungency and its threads, but it is a vivid proof that the soybean's nitrogen can be spun into γ-glutamyl chemistry at will. Traditional soy foods map the range: a mature soy sauce can carry well over a thousand milligrams of free glutamate per hundred grams, a miso several hundred, a fresh tofu almost none until it is cooked or fermented — the same seed, the same reservoir, read out at every point along the release curve. Our engineering simply moves the whole curve: a seed that starts louder, releases more per unit of ferment time, and lands its kokumi on the potent end of the ladder, so every soy tradition that meets it inherits the lift.

For the analog maker there is a cleanup job the soybean program takes on directly, because the seed's gifts come with baggage. Raw soy carries **beany and bitter off-notes** — the green, cardboard note of lipoxygenase-oxidized lipids, the astringency of its isoflavones, the chalk of its saponins — and these are exactly the flavors that betray a plant-based product as *plant-based* in the pejorative sense. Trimming the lipoxygenases quiets the beany note at its enzymatic source; tuning the isoflavone and saponin load lifts the bitterness and astringency that fight the savor. What remains is a cleaner canvas on which the engineered umami and kokumi read without interference — a soybean whose depth arrives unaccompanied by

the flaws that depth used to have to shout over. The agronomy cooperates: the soybean is a hardy, nitrogen-fixing crop grown at scale across the world's temperate farmland, so the cleaned-up, savor-forward seed enters an existing supply chain rather than demanding a new one. Both axes, cleaner notes, and a crop that already feeds billions — the soybean earns its frontier grade and repays it. It is no accident that the one organism carrying both savory axes is also the one the world already grows by the hundred million tonnes; the soybean was always the keystone, and the library simply reads it as one. That keystone standing is why the soybean's 360 constructs, hard-won as they are, return more savory reach per edit than almost any other organism we grow — one seed that answers a dozen traditions. Our engineering works both ends: we raise the seed's **free**-glutamate fraction at maturity through the glutamate engine of Chapter 5, so the seed arrives at the ferment already partly loud, and we tune the storage-protein composition so that proteolysis releases a higher-glutamate peptide profile — more savor liberated per unit of fermentation time.



FIG. 9 — Soybeans in cross-section. Forty percent protein by dry weight: a reservoir of bound glutamate that fermentation, or engineering, must release to taste.

The kokumi half is where the soybean is most naturally gifted and where our work is most rewarding. Legumes carry γ-glutamyl peptides abundantly — it is part of how they store sulfur and nitrogen in the seed, the same logic as the alliums — and the soybean's γ-glutamyl pool includes peptides near the top of the CaSR potency ladder. A fermented soy paste tastes *round* in a way a tomato cannot, and that roundness is γ-glutamyl kokumi generated both by the bean's own stores and by the ferment's transpeptidase. We bias the soybean's γ-glutamyl pathway toward **γ-Glu-Val-Gly** — the GGT edits of Chapter 8, applied to the legume's own transpeptidase — so that the engineered seed carries, and the ferment amplifies, the single most potent body molecule. The outcome is an organism engineered to be loud on both axes at once:

IFI PROFILE

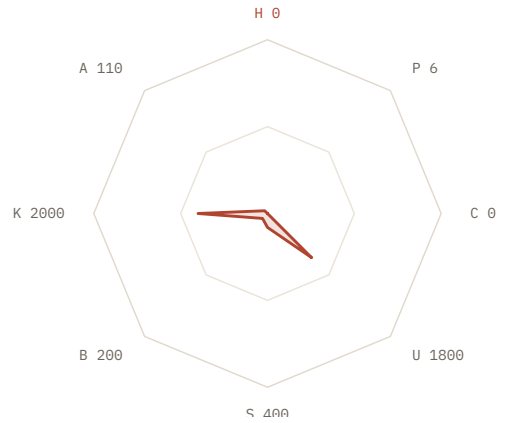

IFI-H		0
IFI-P		6
IFI-C		0
IFI-U		1,800
IFI-S		400
IFI-B		200
IFI-K		2,000
IFI-A		110

Read the balance: umami and kokumi nearly matched, note and body in equal measure, the profile of a thing that will taste *complete* rather than merely savory. This is why the soybean anchors the ferment and meat-analog builds of Part IV. A savory system that must stand in for the full depth of a long-cooked or animal product needs both axes high, and the soybean is the one organism that offers them from a single, protein-dense, versatile input. The 360 constructs are a frontier investment, but the thing they buy — complete savor in one seed — is the most valuable position in the savory library.

Chapter 11 · Yeast — *the manufacturer's engine.*

ORGANISM · 240 constructs · axes U/K · difficulty: Hard · the staged-release workhorse

If the soybean is the savory library's heart, yeast is its factory. No organism we work with converts so directly into a manufacturer's ingredient, because yeast is already the backbone of the savory-ingredient industry: **yeast extract** — autolyzed yeast, the glutamate-and-nucleotide concentrate behind a thousand "natural flavor" lines — is one of the most widely used umami sources on earth, and it is produced by letting the yeast digest itself. The **240 constructs** in our yeast program are aimed squarely at that process, because yeast is the organism where the staged-release principle of Chapter 6 is not an analogy but the entire industrial reality.

Recall the problem the purine chapter posed: a living cell cannot carry large free pools of the savory nucleotides without competing with its own growth. Yeast solves this the way nature solved it in the dried mushroom, only faster and at scale. You grow the yeast **RNA-rich** — fed and divided fast, so that a large fraction of its dry mass is ribosomal RNA, a dense reservoir of bound guanylate and adenylate — and then you **autolyze** it: raise the temperature into the window where the cell's own enzymes turn on itself, and a **5′-phosphodiesterase** cleaves that RNA into free 5′-nucleotides while proteases and peptidases liberate free glutamate and γ-glutamyl peptides from the cell's protein. In one controlled process, the three savory chemistries arrive together — glutamate (the note), guanylate and inosinate (the gain), and glutathione and its γ-glutamyl relatives (the body) — because the yeast cell contained precursors of all three and the autolysis released them all at once.

The organism is almost always **Saccharomyces cerevisiae** — the same species behind bread, beer, and wine, the most domesticated microbe in human history — though the dairy yeast *Kluyveromyces* and the fodder yeast *Candida utilis* earn their place where a feedstock suits them. What makes yeast the nucleotide champion of the library is a number: a yeast grown fast, fed hard, and divided young carries **six to twelve percent of its dry weight as RNA**, almost all of it ribosomal, because a rapidly dividing cell is a cell building ribosomes as fast as it can. That RNA is the guanylate reservoir of Chapter 6, standing ready, and the first lever of our yeast program is simply to grow the cell in the regime that maximizes it — savor begins as a growth curve.

Breaking the cell open is the craft, and there are three ways, only one of which we favor. **Autolysis** lets the yeast digest itself: hold the cell near 50–55 °C at a controlled pH and its own proteases, peptidases, and 5′-phosphodiesterase turn inward, liberating free glutamate, peptides, and 5′-nucleotides over hours to days — gentle, enzymatic, clean-labelled as "yeast extract." **Hydrolysis** forces the issue with acid or added enzymes; it is faster but cruder, and acid

hydrolysis of protein can throw off chloropropanol by-products, which is reason enough to leave it to the past. **Plasmolysis** uses salt to rupture the cell osmotically. We build on autolysis because it is the method that respects the three chemistries — it releases them rather than degrading them — and because it is the method that holds the seven-percent bar, the same biomass and the same program yielding the same extract. The process, stepped out:

STEP	CONDITION	PRODUCES
Grow RNA-rich	fast aerobic feed	high-RNA biomass
Autolyze	50–55 °C, controlled pH	free glutamate, peptides, glutathione
Phosphodiesterase hold	warm, thermostable 5′-PDE	5′-GMP, 5′-AMP from RNA
AMP deaminase	—	5′-IMP (the upgrade to strong synergy)
Separate + concentrate	centrifuge, evaporate	yeast extract (paste or spray-dried)

There is a lineage here the manufacturer will recognize. The nineteenth-century chemist Justus von Liebig dreamed of a concentrated savory extract to feed the many; he built it from meat, and it was ruinously expensive. The dream was realized a generation later not from an animal but from spent brewer's yeast — the sediment of the beer industry, autolyzed into the dark savory pastes that became Marmite and Vegemite. Yeast extract has been the democratic umami ever since: grown, not farmed, cheap, and deep. Our contribution is to make it *complete* — to tune not only the glutamate the old extracts delivered but the guanylate, the inosinate, and the γ-Glu-Val-Gly, so a single modern extract carries the whole multiplicative stack of Chapter 3 rather than only its first term.

Two refinements round out the yeast program. The first is that nothing need be wasted: the autolysis that yields the savory extract also leaves the **cell wall**, a rich source of **β-glucan** and mannoprotein that is itself a valued food ingredient — so a savory yeast is a two-product organism, the soluble extract for depth and the insoluble wall for texture, the economics of the whole improved by using both. The second is that yeast can be tuned toward **kokumi specifically** by raising its glutathione: a glutathione-rich "kokumi yeast" delivers body as directly as a high-RNA yeast delivers gain, and the two can be grown as separate lines or balanced in one. It is worth drawing the distinctions a buyer meets on a spec sheet, because the words are used loosely. **Nutritional yeast** is the whole inactivated cell, not autolyzed — nutty, but not deeply savory. **Autolyzed yeast extract** is the cell self-digested, the clean savory standard. **Hydrolyzed vegetable protein** is a different thing entirely — acid- or enzyme-digested plant protein, savory but carrying the labeling and by-product concerns of Chapter 19. Our work lives in the middle category, the autolysate, because it is the one that releases the three chemistries intact and wears the cleanest label.

The scale economics are what make yeast the manufacturer's favorite in the end. Yeast grows on inexpensive, often **upcycled feedstocks** — molasses, spent brewing and distilling streams, other food-industry side-sugars — turning a by-product into a savory

ingredient, which is both a cost advantage and a circular one. It grows **fast**, doubling in hours where a crop takes a season, in fermenters that run independent of weather and acreage, so supply scales with tank volume rather than with farmland. And because the whole process is defined and enclosed, its seven-percent batch bar is the easiest in the library to hold, and its scale-up is a matter of engineering rather than agronomy. A savory depth that can be grown this cheaply, this fast, and this reproducibly — from what another industry threw away — is a savory depth that can reach a very large number of plates, which is, in the end, the point of growing it at all. A savory ingredient that scales like software and prices like a commodity is the quiet reason the manufacturer's shelf reorganizes around a single grown extract rather than a wall of separate drums.

Our engineering tunes each stream of that release. For the **note**, the glutamate engine edits raise the cell's free and protein-bound glutamate so autolysis liberates more. For the **gain**, we raise the cell's RNA fraction and tune the phosphodiesterase (the thermostable L212F variant of Chapter 8, built to survive the warm hold) so more of the RNA reservoir converts to free nucleotide — and we raise the **AMP deaminase** so the released adenylate, a weak synergist, is converted to inosinate, a strong one. For the **body**, we up-regulate GCL and bias GGT toward γ-Glu-Val-Gly, so the autolysate carries not just glutathione but the thirteen-fold-potent tripeptide. The engineered yeast is therefore not a single-axis extract but a complete, grown-and-released savory stack, tuned stream by stream:

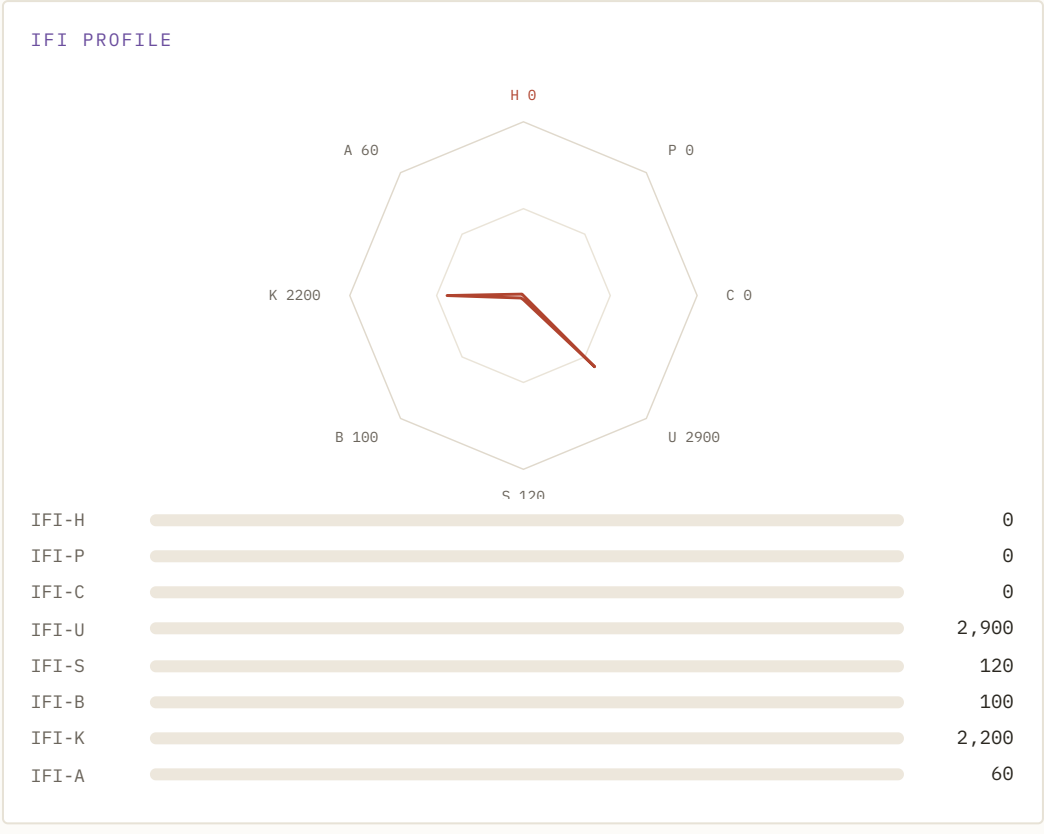

RELEASE STREAM	PRECURSOR IN THE CELL	RELEASE ENZYME	SAVORY OUTPUT
Note	protein-bound + free glutamate	proteases, peptidases	free glutamate
Gain	ribosomal RNA	5′-phosphodiesterase	free GMP, AMP
Gain (upgrade)	released AMP	AMP deaminase	IMP (stronger)
Body	glutathione reservoir	— (released intact) + GGT	glutathione, γ-Glu-Val-Gly

The manufacturing case for engineered yeast is almost unfair. It is **grown, not farmed** — fermenter batches on a defined feedstock, no field, no season, no shipping loss of the kind that ruined the commodity tomato — so its seven-percent batch variation is the easiest in the library to hold: same strain, same feed, same autolysis program, same profile, every time. It is a **drop-in** for an industry that already buys yeast extract by the tonne, so the formulation and the label line do not change; only the depth does. And it carries **all three savory chemistries in one ingredient**, which means a manufacturer reaching today for separate drums of glutamate, disodium inosinate, and a kokumi flavor can replace the shelf with a single grown extract that delivers the synergy of Chapter 3 already composed. Yeast is where the whole book becomes a product you can order.

Chapter 12 · Koji, shiitake & the mushrooms — *savor released by other life.*

ORGANISMS · koji 80 · shiitake 80 · mushroom 40 · axis U · difficulty: Hard

Three of our savory organisms do not make their depth so much as **unlock** it — in the soybean, in rice, in their own flesh — and they do it with enzymes so useful that human cooking has been borrowing them for two thousand years. Koji, shiitake, and the broader mushrooms are the *liberating* organisms: fungi whose gift is not a large native pool of free savor but a toolkit for setting savor free from substrates that hold it bound. They are hard organisms in our grading — fungal genetics, slower and less charted than a tomato's — but the **80 constructs in koji**, the **80 in shiitake**, and the **40 in mushroom** reach the most leveraged enzymes in the whole savory library.

Begin with **koji** — *Aspergillus oryzae*, the mold that makes soy sauce, miso, sake, and amazake possible, the "national fungus" of Japan and arguably the most important savory microorganism on earth. Koji's power is proteolytic. Grown over steamed soybeans or rice, it secretes a battery of **proteases and peptidases** that chop storage protein into free amino acids and peptides, and — the crucial enzyme — a **glutaminase** that converts the glutamine liberated by that proteolysis into **glutamate**. This last step is the hinge of soy-sauce umami: proteolysis releases far more glutamine than glutamate, and glutamine is nearly tasteless; it is the glutaminase that turns that silent pool into savor. Koji strains are prized or passed over almost entirely on their glutaminase activity, and it is the first thing our koji program raises — a higher-glutaminase koji liberates more umami per unit of fermentation, turning the soybean's bound reservoir loud in less time.

Glutaminase is only the headline of a remarkable toolkit. *A. oryzae* secretes a whole hydrolytic arsenal, and which part matters depends on the ferment. Its **α-amylase and glucoamylase** break starch into glucose — the chemistry behind sake (feeding sugar to a *Saccharomyces* that then makes alcohol) and amazake (a sweet, alcohol-free rice drink that is nothing but koji's amylases run on rice). Its **acid, neutral, and alkaline proteases**, backed by leucine aminopeptidase and carboxypeptidases, chop storage protein down to free amino acids — the chemistry behind soy sauce and miso. And across all of them the **glutaminase** converts the liberated glutamine to glutamate, the step that turns silent protein into savor. A koji program is really a dial board across this arsenal: raise glutaminase for umami, tune the amylase-to-protease balance to the ferment, and the one mold serves sake, miso, shoyu, and amazake from the same genetics.

There is a safety story here that is also a design story. *A. oryzae* was **domesticated from the wild mold *A. flavus***, the notorious producer of aflatoxin — and over centuries of selection by brewers and miso-makers, the toxin-making machinery was silenced or lost, leaving a food-safe organism where its ancestor is a hazard. Koji is, in other words, a mold humans engineered by eye and palate long before anyone could read a genome: a living proof that savor and safety can be bred together. Our contribution is to do deliberately, and to the seven-percent bar, what centuries of koji-makers did by selection — raise the useful enzymes, confirm the silenced toxin pathways stay silent, and lock the result.

A few details of the craft reward attention. Soy sauce is often made not with *A. oryzae* but with its close relative ***A. sojae***, selected over centuries for a protease profile suited to soy, while *A. oryzae* dominates sake and miso — two domesticated molds for two savory traditions, each a living archive of human selection. In sake, the *degree of rice polishing* is a quality obsession precisely because it tunes the substrate the koji works on: milling away the outer grain shifts the starch-to-protein ratio the amylases and proteases meet, and with it the balance of sweetness and savor. And the regulatory standing is settled — *A. oryzae* carries a long record of safe use and is treated as food-grade across the major jurisdictions, the domesticated end of a genus whose wild members are hazards. When we raise a koji's glutaminase, we are working inside that safety envelope rather than stepping outside it: the organism was made safe by the same centuries of selection that made it useful, and our edits are held to confirm it stays both.



FIG. 10 — A controlled chamber. The savory fungi do their work in defined warmth and humidity; a higher-glutaminase koji is a tuned enzyme release, held to the same batch tolerance as a grown plant.

Then the **mushrooms**, whose gift is the other half of the synergy — the nucleotide. The shiitake is the great guanylate carrier of the kitchen, and it demonstrates the staged-release principle of Chapter 6 more vividly than anything else we grow. A fresh shiitake is only mildly savory; a **dried** one, soaked and warmed, is profound — because drying and gentle reheating activate the mushroom's own ribonuclease and **5′-phosphodiesterase**, which cleave its ribosomal RNA into free 5′-nucleotides, chief among them **GMP**, the strongest umami synergist of all (β ≈ 2,800 from Chapter 3). The guanylate

you taste in a shiitake broth was not a free pool in the living mushroom; it was liberated from the fungus's genetic machinery by heat, on the way to the pot. Our shiitake and mushroom programs tune that release — raising the RNA reservoir and the thermostability of the liberating phosphodiesterase — so the dried, rehydrated mushroom delivers a larger, more reproducible guanylate pool.

The release has a **temperature window**, and it is the single most mishandled fact in savory cooking. The mushroom's ribonuclease first cuts its RNA, and a 5′-phosphodiesterase then produces the 5′-guanylate the receptor wants; both enzymes work best when a dried mushroom is rehydrated and brought *slowly* up through roughly **60–70 °C**. Drop a dried shiitake into already-boiling water and you denature the release enzymes before they work; soak it only in cold water and you never activate them. Peak guanylate asks for a patient warm-up — which is exactly why the broth method of Chapter 14 starts the mushrooms in cold water and climbs to a bare simmer over half an hour. The cook who rushes the dried mushroom is throwing away most of its synergy, and did not know there was an enzyme to respect. Beyond shiitake, the mushroom kingdom is a savory library of its own:

MUSHROOM (DRIED)	LOUD IN	CHARACTER
Shiitake	5′-GMP	the guanylate standard
Porcini (*Boletus*)	glutamate + GMP	very high free glutamate
Maitake	GMP	deep, earthy
Enoki	glutamate	delicate, clean

The profile of our engineered shiitake is almost pure note-and-gain, little body:

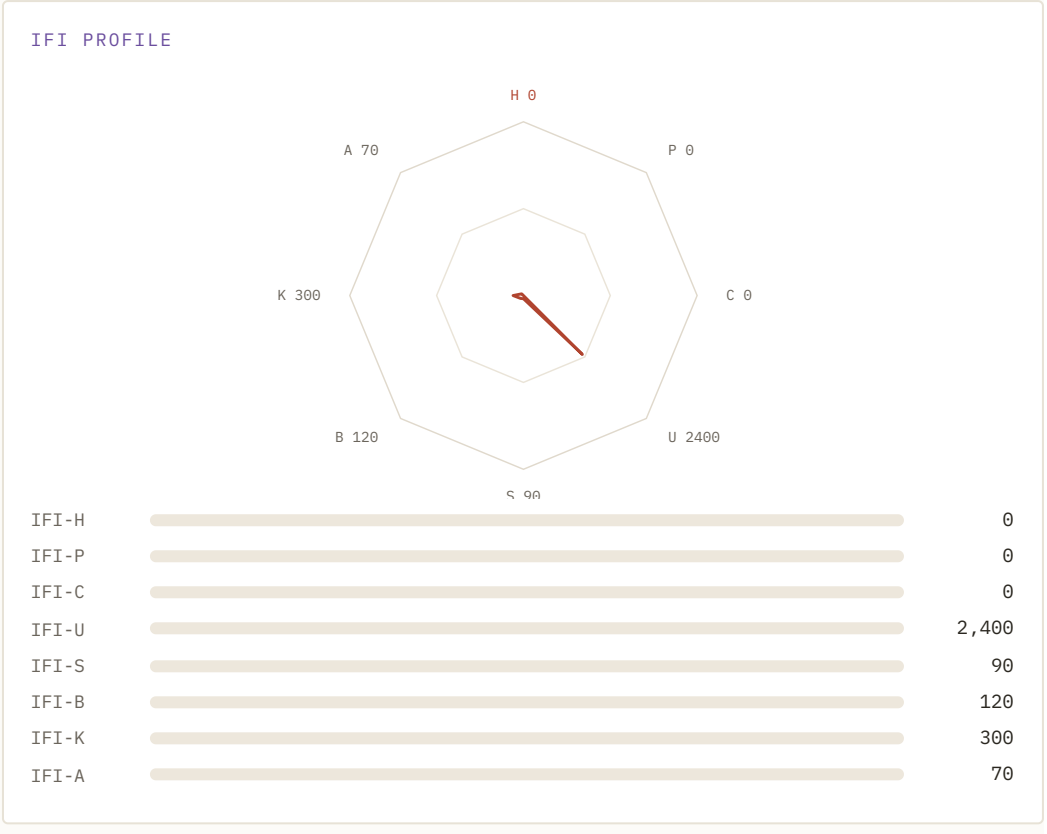

ORGANISM	CONSTRUCTS	THE ENZYME THAT MATTERS	WHAT IT LIBERATES
Koji (*A. oryzae*)	80	glutaminase (+ proteases)	free glutamate from soy/rice protein
Shiitake	80	5′-phosphodiesterase	free GMP from its own RNA
Mushroom	40	ribonuclease / 5′-PDE	free 5′-nucleotides

Put koji and shiitake on the same page and the whole synergy of this book is staring back at you, assembled by evolution long before anyone measured an EC_{50}: koji releases the **glutamate**, the mushroom releases the **guanylate**, and the two together — kombu-and-bonito's vegetarian cousin, the backbone of a mushroom-and-soy broth — multiply into a savor far beyond either. These are the organisms that taught human cooking what synergy was, by accident, over centuries. We have simply learned to tune their enzymes and read the result on a number.

The release enzyme at the center of all this has an industrial afterlife worth knowing. **Nuclease P1**, a 5′-phosphodiesterase from the mold *Penicillium citrinum*, is the enzyme the flavor industry uses to digest yeast RNA into the 5′-nucleotides that season a thousand products — the same chemistry the shiitake runs on itself, bottled as a reagent. When we tune a mushroom's or a yeast's own phosphodiesterase, we are internalizing that industrial step into the organism, so the biomass arrives able to liberate its own guanylate without a separate enzyme dosing. And the mushroom kingdom keeps rewarding the search: dried **porcini** carries free glutamate at levels that rival kombu *and* releases guanylate on cooking, a rare organism loud on both halves of the synergy at once; **maitake** and the oyster mushrooms bring their own balances; the humble button mushroom, roasted hard to concentrate and brown, becomes a savory base far beyond its raw self. Each is a different distribution of the same two chemistries, and the library grows by reading them, ranking them, and learning which enzyme to raise in which. The fungi are the hardest savory organisms we grow and, cell for cell, among the most rewarding, because an enzyme tuned in a mold is an enzyme that keeps working long after the greenhouse lights go out.

Chapter 13 · The alliums & the quiet kokumi — *body, from the bulb and the leaf.*

ORGANISMS · garlic 560 · onion 40 · vanilla 40 · tea 280

The last of the savory organisms are the ones that work beneath the note — the carriers of **body** and of the quieter savory accents that make a complete dish complete. Garlic and onion are the kokumi foundation of nearly every cuisine on earth; vanilla and tea are the surprises, savory in ways the words *sweet spice* and *beverage* conceal. Together they round out the library and show that kokumi is not a laboratory abstraction but the oldest trick in cooking: the allium you start almost every savory dish with is there, whether the cook knows it or not, to build body.

Garlic carries the largest allium program — **560 constructs** — and it is the clearest kokumi organism we grow, because garlic is already a professional γ-glutamyl chemist, as Chapter 7 described. Its dormant clove is a warehouse of **γ-glutamyl-S-alk(en)yl-cysteines**, the storage form of its sulfur, and aging or sprouting runs the allium **γ-glutamyl transpeptidase** that mobilizes them — which is the biochemical truth behind every cook's knowledge that aged, confited, or black garlic tastes rounder and deeper than raw. Garlic also sits at an axis intersection: it is graded **K/P**, kokumi *and* pungency, because the same sulfur chemistry that feeds its γ-glutamyl stores also produces **allicin**, the TRPA1 agonist that gives raw garlic its sharp bite. Our garlic program pulls these apart as dials — biasing the γ-glutamyl pathway toward **γ-Glu-Val-Gly** for more body, and tuning the alliinase-driven pungency up or down independently — so a formulator can order garlic as *body with a little heat-sharpness* or *body with almost none*.

Those two dials are coupled by a shared origin, and the coupling is what makes the garlic program elegant rather than merely additive. The allium's γ-glutamyl peptides are not only kokumi molecules; they are the **biosynthetic precursors of alliin**, the odorless cysteine sulfoxide warehoused in the intact clove. Cut the clove and the enzyme **alliinase** — a carbon-sulfur lyase, held apart from its substrate in the living cell and released only by damage — converts alliin into **allicin**, the unstable thiosulfinate that is garlic's pungent bite and its TRPA1 signal. In onion a second enzyme, **lachrymatory-factor synthase**, diverts the same chemistry to propanethial S-oxide, the volatile that makes a cook cry. So the pool we skim for body is the very pool the plant would otherwise route toward sulfur storage and, on wounding, toward pungency — which means that biasing the γ-glutamyl transpeptidase toward γ-Glu-Val-Gly both *raises* kokumi and, by diverting precursor away from alliin, *gentles* the pungency. "Body with almost none" is not a blend of two ingredients; it is one pathway steered, the sharpness traded for depth at the branch point. It is also why **black garlic** — cloves held for weeks at warm humidity

— tastes sweet, dark, and profoundly round: the long warm hold runs the transpeptidase (building body) and the Maillard reaction (building color and molasses depth) while the alliinase pungency fades to nothing. Black garlic is a slow kokumi reactor that a patient cook runs over weeks; our engineered garlic arrives at a similar body without the wait. The profile is kokumi-led, pungency as an accent:

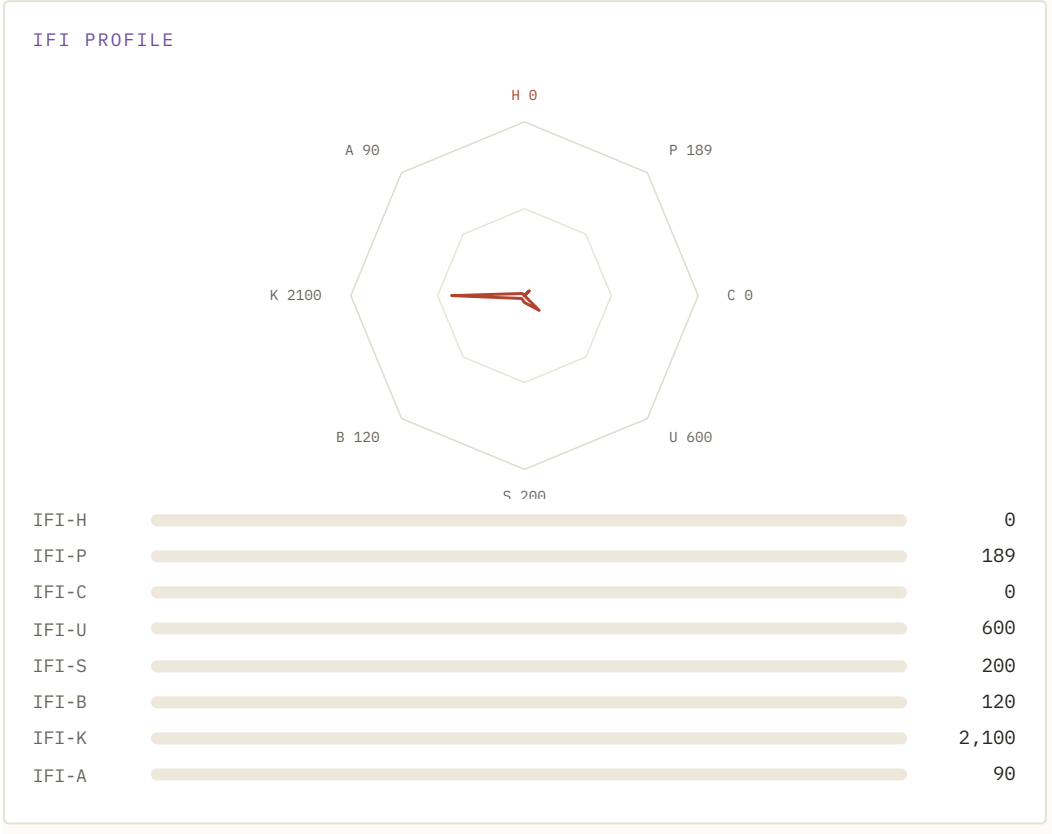



FIG. 11 — Garlic, whole. The clove is the kitchen's oldest kokumi delivery system; engineered, it becomes a body dial a formulator can set without the sharpness that used to come with it.

Onion (**40 constructs**, graded **K**) is garlic's gentler sibling — the same γ-glutamyl storage chemistry with less sulfur fire, which is why onion is the universal savory base that can go into anything without dominating. Cooked slowly, an onion builds body and a rounded sweetness as its γ-glutamyl peptides mobilize and its sugars caramelize; our work raises the kokumi it delivers without the hours of low-heat cooking a deep onion base traditionally demands.

The onion's slow-cook is worth pausing on, because it is a kitchen's most common act of savory design, performed by millions who never name it. As an onion cooks down over low heat, three things happen at once: its γ-glutamyl peptides mobilize into kokumi, its stored fructans hydrolyze and caramelize into sweetness, and its

harsh sulfur volatiles cook off — the sharp raw bulb becoming the sweet, deep, jammy base that underlies soups and sauces across every cuisine. A French onion soup is, underneath, a kokumi extraction with a Maillard finish; so is a Punjabi onion masala, browned for the better part of an hour before anything else goes into the pan. What the long cook buys is body, and the engineered onion is a way to start closer to that body — to shorten the hour without losing the depth it was building. Vanilla keeps the same quiet company: its contribution to a savory custard or a cream sauce is a rounding continuity a taster credits to richness or to fat, when it is in fact a kokumi mouthfeel doing the work. The lesson of this chapter is that body has always been engineered in the kitchen — by the cook who browns the onion, ages the garlic, or scrapes the pod — and that we are only moving that engineering upstream, into the plant, so it need not be re-earned by hand every single time. **Vanilla** is the genuine surprise of this chapter. Graded **K** and sharing a shelf elsewhere with the sweet organisms, vanilla's savory contribution is not its aroma but its **mouthfeel** — a rounding, lingering body that makes vanilla indispensable in custards and sauces for reasons that have nothing to do with the vanillin everyone notices. It is a kokumi organism masquerading as a sweet one, and in savory formulation a trace of engineered vanilla does the quiet work of continuity that a cook cannot name but misses when it is gone.

Tea (**280 constructs**, graded **U/B/A**) closes the savory roster from an unexpected direction. Tea's umami is real and specific: it comes from **L-theanine**, an amino acid chemically kin to glutamate that gives fine green tea — gyokuro, matcha — its unmistakable brothy, savory depth, the quality tea people call *marine* or *umami* without irony. L-theanine's biosynthesis is a small, elegant story that doubles as a cultivation lesson. The enzyme **theanine synthetase** condenses glutamate with ethylamine — itself from the decarboxylation of alanine — and the reaction runs chiefly in the tea plant's **roots**, from which theanine is shipped up to the flushing new leaves. Light drives theanine's onward conversion into catechins, the polyphenols behind tea's bitterness and astringency, so a leaf grown in sun trades its savor for bite. This is the mechanism behind the oldest luxury in tea: **shading**. Cover the bushes for the weeks before harvest — the practice behind gyokuro and the stone-ground matcha made from it — and theanine stays high while catechins stay low, which is why shaded teas taste most intensely brothy and least astringent. Our tea program delivers by genetics what the shade delivers by deprivation: it raises theanine and tunes the catechin balance directly, so the umami of a fine shaded tea need not be bought with weeks under cloth.

Tea sits at a three-axis intersection (umami from theanine, the attenuated bitterness of its catechins, and the astringency of its tannins), and in the savory library it is the proof that umami is not confined to the obviously savory — that a leaf brewed as a drink can carry the same note a broth does, read on the same Index.

It is worth ending the organism chapters on that widening, because it is the whole shape of the savory library. Savor began, in this book, as the obvious depth of a broth; it ends as a note found in a shaded

leaf, a rounding in a vanilla pod, a body warehoused in a dormant clove. The kitchen always knew savor reached further than the stockpot — the cook who finishes a sauce with a scrape of vanilla or a splash of green tea was reading umami and kokumi without the words — and the Index simply gives that intuition a number and a receptor. From here, the organisms stop being subjects and become ingredients: the next part is what a kitchen and a factory actually build from them.

ORGANISM	CONSTRUCTS	AXES	SAVORY ROLE	SIGNATURE MOLECULE
Garlic	560	K/P	kokumi foundation (+ pungency dial)	γ-Glu-Val-Gly, allicin
Onion	40	K	universal savory base	γ-glutamyl peptides
Vanilla	40	K	mouthfeel / continuity	(kokumi peptides)
Tea	280	U/B/A	brothy umami from a leaf	L-theanine

With the organisms assembled — the tomato for the note, the soybean for both axes, the yeast for the manufacturer's complete stack, koji and the mushrooms for release and synergy, the alliums and the quiet carriers for body — the savory library is complete. Everything from here is composition: taking these grown savory inputs and building with them the things a kitchen and a factory actually make.



PART IV OF V

The Compositions

broth · paste · the honest analog · the ferment — the organisms, assembled

Chapter 14 · Designing a broth — *the stack, assembled.*

COMPOSITION · the U×K master formula · two builds, scored

A broth is the purest test of savory design, because a broth is savor with nowhere to hide. There is no crust, no char, no fat-slicked richness to carry a thin middle — only water and what you have dissolved into it, read directly by the tongue. Every culture that cooks has a broth, and every one of them is, underneath, the same equation: a glutamate source for the note, a nucleotide source for the gain, and time-or-aging for the body. What our organisms change is not the equation but the *inputs* — each one grown to a locked profile, so the broth is composed from numbers rather than coaxed from chance.

Start from the synergy of Chapter 3, written as a shopping list. You want **one loud glutamate carrier**, **one potentiating nucleotide carrier**, and **one body source**, and you want them in a ratio where the glutamate × nucleotide product is large and the γ-glutamyl peptides are present enough to round and lengthen the result. Translated into grown organisms: the restored **tomato** or a **kombu**-style glutamate source for the note; the dried **shiitake** for the guanylate gain; the engineered **yeast** and slow **allium** for the body. Here is that stack as a kitchen-scale broth, built to taste like it ran all day because, chemically, it is standing in for exactly the chemistry an all-day simmer produces.

```
THE ALL-DAY VEGETABLE BROTH          yield 2 L · active 25 min · total 1 h
—
2.0   L    water, cold
30    g    kombu (glutamate carrier)
25    g    dried shiitake (guanylate carrier)
250   g    restored tomato, halved, cut face seared
150   g    onion, halved, cut face charred dark
20    g    garlic, lightly crushed, skins on
8     g    engineered yeast extract (the body stack)
6     g    tamari or shoyu, to finish
—                                        salt to the dish it serves, not
here
```

1. Put the kombu and shiitake in the cold water and bring it *slowly* toward a bare simmer over 30 minutes — never a boil. The slow warm-up is not ceremony: it is the window in which the shiitake's phosphodiesterase cleaves its RNA to free GMP, and in which the kombu's glutamate dissolves without the harsh notes a hard boil drags out of kelp.
2. While it heats, sear the cut faces of the tomato and char the onion hard in a dry pan — color is free glutamate made loud and free sugar turned to savor-adjacent depth.
3. At the first tremble of a simmer, pull the kombu (it has given its glutamate; held longer it turns to glue). Add the seared tomato, charred onion, and crushed garlic. Hold at a bare simmer 25 minutes.

4. Off the heat, whisk in the yeast extract until dissolved — added at the end, its nucleotides and γ-glutamyl peptides are not driven off by prolonged heat. Finish with the tamari. Strain.

IFI PROFILE

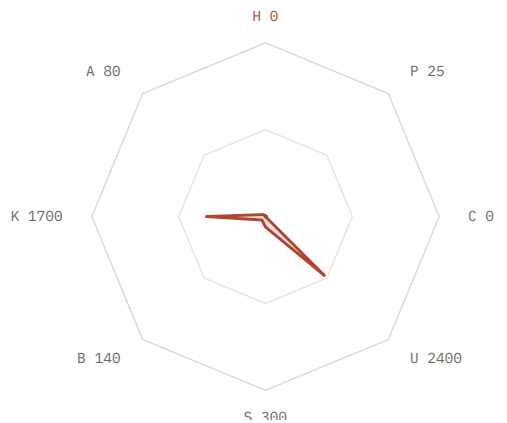

IFI-H		0
IFI-P		25
IFI-C		0
IFI-U		2,400
IFI-S		300
IFI-B		140
IFI-K		1,700
IFI-A		80

Read the score and you can taste it in advance: the note at 2,400 is the kombu-and-tomato glutamate multiplied by the shiitake guanylate; the body at 1,700 is the yeast and garlic γ-glutamyl load, enough to make the broth fill the mouth and fade slowly rather than ending at the front teeth; the faint pungency is the garlic. This is a vegetarian broth that reads, on the Index, as deep as a long-simmered stock — the whole point of growing the savor in.

Before the concentrate, the clearest teaching broth in the world is worth building on its own, because it isolates the synergy to two ingredients and a thermometer. A vegetarian **ichiban dashi** is nothing but kelp glutamate and mushroom guanylate, and its whole art is temperature.

```
KOMBU-SHIITAKE DASHI (ichiban)        yield 1 L · active 10 min · total 45 min
—
1.0   L    water, cold, soft
20    g    kombu
20    g    dried shiitake
—                                     nothing else; this is the lesson
```

1. Steep the kombu and shiitake in the cold water, unheated, for 30 minutes — a cold start that lets the kombu's glutamate begin dissolving and primes the mushroom's enzymes.
2. Bring to 60–65 °C over 10 minutes and **hold** there a few minutes: this is the window in which the shiitake's 5′-phosphodiesterase liberates guanylate, the nucleotide that makes the synergy.
3. Just before a simmer, near 80 °C, **pull the kombu** — held hotter, kelp gives alginates and a marine harshness, its glutamate already spent. Bring briefly to a bare simmer, pull the shiitake, and strain.

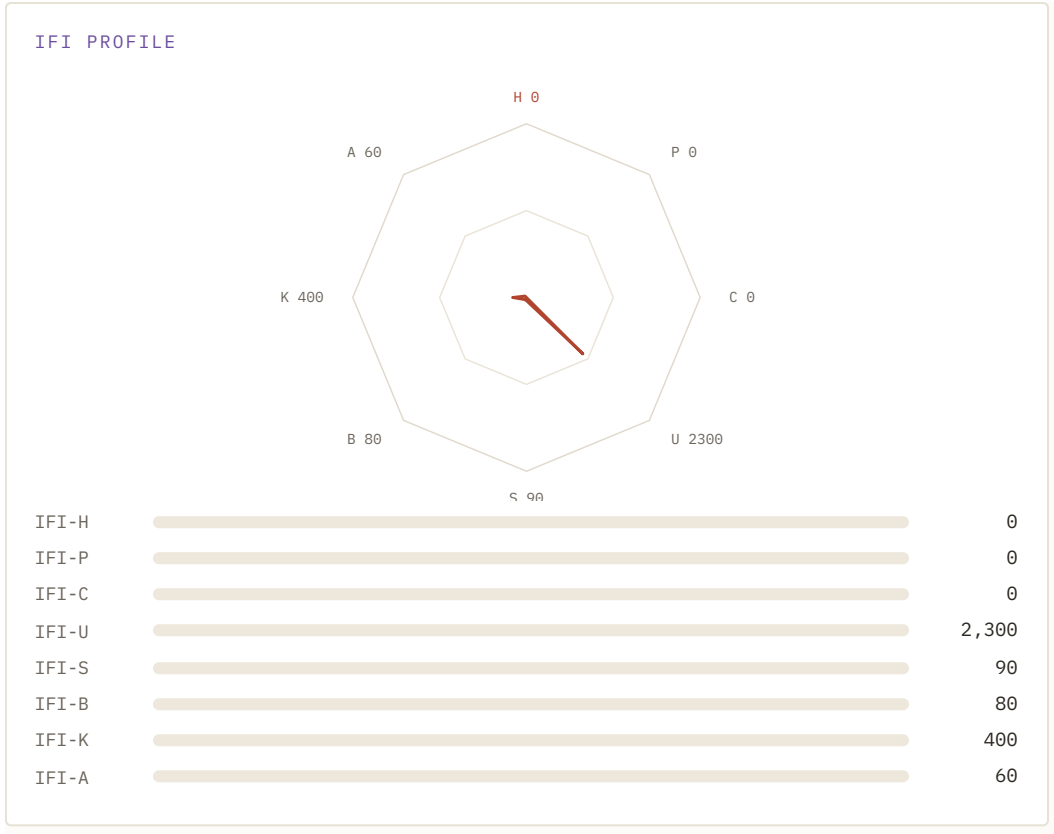

The extraction science *is* the dish. **Glutamate** dissolves from kombu across a wide temperature range and needs no enzyme — only time and not-too-much heat. **Guanylate** must be *made*, by the mushroom's own enzymes, in the 60–70 °C window of Chapter 12 — which is why a slow climb outperforms a hard boil, and why the cold start is chemistry rather than ceremony. Two ingredients, one thermometer, and the whole synergy of this book in a clear amber cup. Everything else a savory kitchen builds is this, elaborated.

Two notes carry the dashi logic into a working kitchen and a working plant. The first is that the spent solids are not spent: the once-drawn kombu and shiitake still hold glutamate and guanylate, and a second, longer extraction — the **niban dashi**, the "second dash" — pulls a milder stock ideal for braising and simmering, the frugality of the technique built into its structure. A savory manufacturer runs the same idea as counter-current extraction, washing loudness from the biomass in stages so little is thrown away. The second is that **water is an ingredient**. Glutamate extraction is sensitive to the water's mineral load — hard, calcium-rich water interferes with the clean draw of kombu's glutamate, part of why the great dashi traditions grew up around soft water — so a broth specified to the seven-percent bar specifies its water too, or treats it. These are not fussy details; they are the difference between a broth that reads at its designed IFI and one that falls short for reasons a cook cannot see. The organisms set the ceiling of savor a broth can reach; the extraction decides how much of that ceiling actually arrives in the bowl.

One more property separates a broth that merely tastes deep from one a manufacturer can ship: **clarity and stability**. A clear broth reads cleaner and more savory than a cloudy one of the same chemistry, because suspended solids blur both the flavor and the look; the classic defense is a gentle, never-boiling simmer that keeps proteins from emulsifying into haze — the same low heat the synergy already wanted. A broth base then deploys in a hundred forms: the liquid stock of a soup, the hydrating liquor of an instant

noodle, the braising medium of a sauce, the dehydrated bouillon of a cube or a powder — each format asking the extraction and the concentrate to hit the same IFI through a different water content. Designing the base to a scorecard rather than to a taste is what lets one savory architecture serve all of them: the number travels where the recipe cannot, and a formulator in another kitchen can reproduce the bowl from the profile alone.

For a manufacturer the same stack is built as a **concentrate**, where water is the enemy and the profile is pushed to the normalization ceiling so a little dissolves into a lot. The formula is the same equation at ten times the strength, held to the seven-percent batch bar:

```
BASE STOCK CONCENTRATE 10×            formulation · % by weight
—
38.0  %    restored tomato paste (glutamate note)
22.0  %    engineered yeast extract (note + gain + body)
14.0  %    shiitake concentrate (guanylate gain)
9.0   %    roasted onion paste (body + sweetness)
6.0   %    koji-fermented soy solids (note + body)
5.0   %    garlic paste, low-pungency line (body)
6.0   %    water / humectant to texture
—                              dilute 1:9 for a finished broth
```

1. Blend the tomato paste, yeast extract, and shiitake concentrate as the savory core; verify the free-glutamate and 5′-nucleotide load against the target before proceeding.
2. Fold in the roasted onion, koji solids, and low-pungency garlic for the body stack; the koji solids also carry residual glutaminase that continues to liberate glutamate in the jar.
3. Adjust to texture with the humectant; the concentrate should be spoonable, not pourable, so it ships without a cold chain.

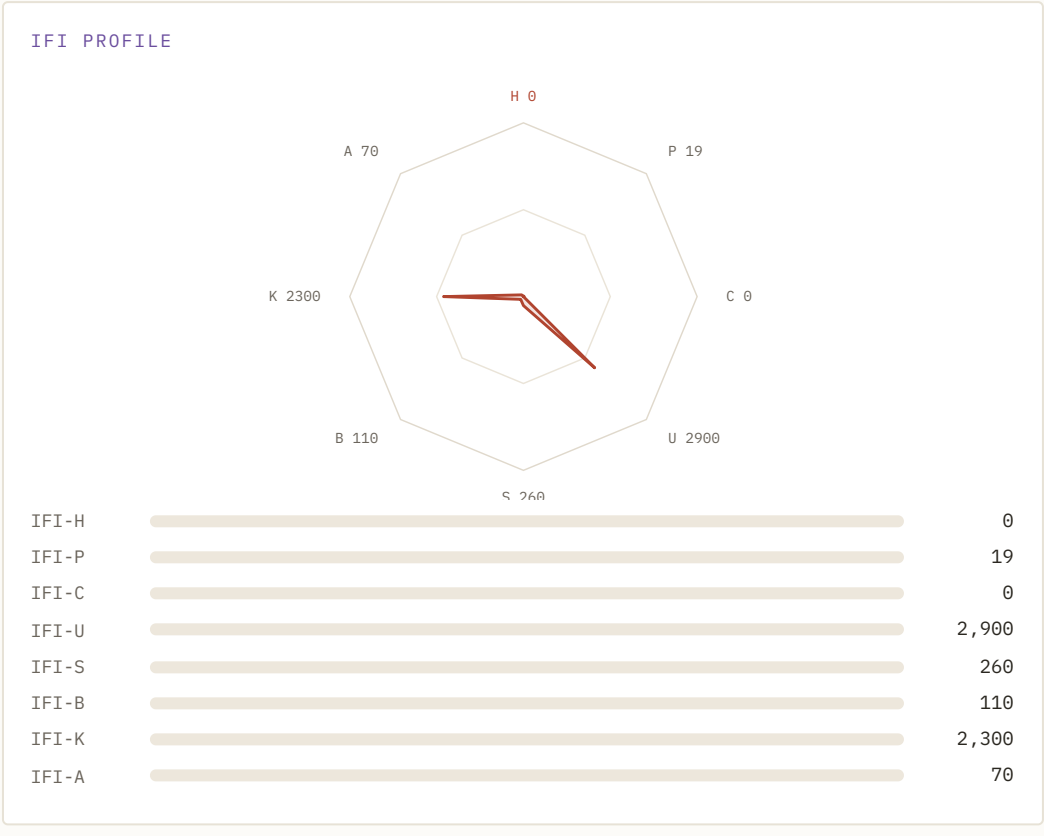

The concentrate is the broth's industrial twin — the same architecture, the ceiling pushed, both axes near the top of the scale. It is the clearest demonstration in this book of the manufacturing promise: a single shelf-stable ingredient, grown and assembled from

designed organisms, that turns water into all-day broth at a one-to-nine dilution, with a profile a formulator can read, order, and reproduce.

Chapter 15 · Sauces, pastes & the reduction — *depth you can hold in a spoon.*

COMPOSITION · concentration as design · three savory condiments, scored

A sauce is a broth that has given up water for intensity, and in that trade the savory axes behave in a way worth stating precisely: **reduction is multiplicative on the note and roughly additive on the body.** As water leaves, free glutamate and nucleotide concentrate together, and their *product* — the synergy term of Chapter 3 — climbs faster than either alone, which is why a reduced sauce tastes so much more than proportionally savory. The γ-glutamyl peptides concentrate too, but kokumi's contribution is a modulation that saturates, so body grows steadily rather than explosively. The practical lesson for anyone building a savory condiment is that reduction is a lever on *note* far more than on *body* — and that if you want the body to keep pace, you add it, from the organisms, rather than reduce for it.

Put numbers on it, because the asymmetry is sharper than intuition expects. Reduce a sauce to a fraction **f** of its starting volume — f = 0.25 for a four-fold reduction — and every dissolved concentration rises by 1/f. The glutamate climbs to u/f and the nucleotide to v/f, but the synergy term of Chapter 3, the product **β·u·v**, climbs to β·u·v / **f²**. A four-fold reduction therefore multiplies the linear glutamate term four-fold and the synergy term **sixteen-fold**: the note does not merely concentrate, it compounds, which is the real reason a hard-reduced sauce tastes so disproportionately deep. The kokumi term does not share the explosion. The γ-glutamyl peptides concentrate to k/f like everything else, but their *perceived* effect is a receptor modulation that **saturates** as CaSR fills, so body rises gently where note rises as a square. Reduction thus drives U up steeply and K up barely — and a sauce reduced for depth arrives loud and thin-bodied unless the cook puts the body back. The organisms are how you put it back without another hour on the stove: a spoon of engineered yeast or low-pungency garlic restores the kokumi the reduction could not, and the Index shows the correction directly, as a K climbing to meet the U.

The workhorse condiment is the **savory depth paste** — a concentrated, shelf-stable umami-plus-kokumi base that a cook or a line stirs into anything that needs a floor. It is the single most useful thing a savory manufacturer can make from this library, and it is built to sit at the top of both axes.

SAVORY DEPTH PASTE formulation · % by weight

30.0	%	engineered yeast extract (complete stack)
24.0	%	restored tomato paste, double-concentrated
16.0	%	koji-fermented soybean paste (miso-style)
10.0	%	dried shiitake, powdered (guanylate)
8.0	%	roasted garlic, low-pungency (body)
6.0	%	roasted onion (body + sweetness)
6.0	%	water to texture

1 tsp lifts a bowl; 20 g makes a sauce

1. Toast the shiitake powder briefly to sharpen its guanylate release, then blend with the yeast extract and tomato paste as the savory core.
2. Work in the miso-style koji paste last and *do not cook it hard* — its live glutaminase and its γ-glutamyl peptides are heat-sensitive, and much of its value is the slow savor it keeps generating in the jar.
3. Fold in the roasted alliums; adjust to a thick, spoonable paste.

IFI PROFILE

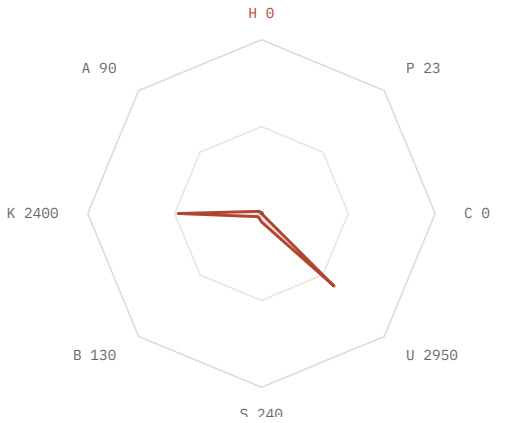

IFI-H	0
IFI-P	23
IFI-C	0
IFI-U	2,950
IFI-S	240
IFI-B	130
IFI-K	2,400
IFI-A	90



FIG. 12 — Vanilla pods. Beyond its aroma, vanilla is a kokumi organism; a trace in a savory reduction lends the rounding continuity a cook feels but cannot name.

What makes a paste *shelf-stable* is worth a line, because it is the quiet engineering under every concentrate in this book. A paste keeps without refrigeration because its **water activity** — the fraction of its water that is chemically free rather than bound to sugar, salt, and dissolved solids — is held low enough that microbes cannot grow, typically below about 0.85. The restored-tomato and yeast solids, the salt of a miso, and any added humectant all bind water, dropping the activity and buying months of ambient shelf life, which is why a concentrate can be built to the normalization ceiling and still ship without a cold chain. It is the same physics that lets soy sauce, miso, and fish sauce sit at room temperature for years: they are savory *because* they are concentrated, and stable *because* that concentration binds their water. There is a second format worth naming — the **savory oil or fat**, where kokumi-active and aroma compounds ride a lipid phase rather than a water one, giving a different delivery (a drizzle, a finish) and a different stability. Across both, the design rule is the same: push the savory solids up, bind the water down, and let the Index confirm the loudness survived the concentration.

Two more condiments show the range. The first is a **vegetarian demi-glace** — the French mother of savory reductions, traditionally days of roasted bones and veal stock, here rebuilt from mushroom, tomato, and the yeast stack, reduced hard so the note multiplies, with a trace of vanilla for the continuity a real demi gets from gelatin and time.

```
VEGETARIAN DEMI-GLACE                  yield 500 mL · total 2 h
—
1.5   L    mushroom-and-shiitake broth (Ch. 14 method)
200   g    restored tomato paste
40    g    roasted onion paste
15    g    engineered yeast extract
1     pod  vanilla, seeds scraped (kokumi mouthfeel)
—                                a splash of tamari; no thickener needed
```

1. Reduce the mushroom broth with the tomato paste and onion by three-quarters over low heat — slow, so the synergy concentrates without scorching.
2. Off the heat, whisk in the yeast extract and the vanilla seeds; the vanilla's job here is pure body, not aroma.
3. The reduction should coat a spoon from its own dissolved savor and glycerol — no starch, no roux.

IFI PROFILE

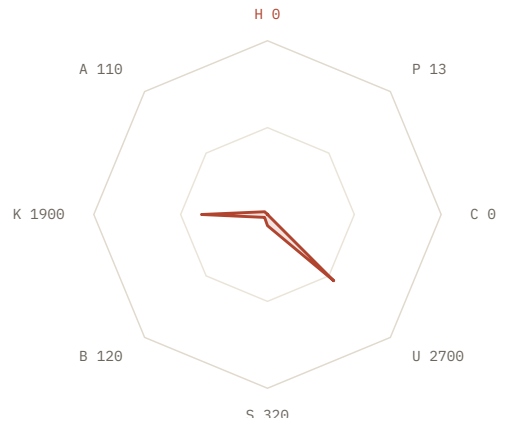

IFI-H	0
IFI-P	13
IFI-C	0
IFI-U	2,700
IFI-S	320
IFI-B	120
IFI-K	1,900
IFI-A	110

The second is a **mushroom XO-style condiment** — the Cantonese luxury relish, classically dried seafood and chili, here an all-plant version where crisped mushrooms and low-pungency garlic carry the savor and a measured pepper line carries the heat. It is the book's reminder that savory depth is a *base* other axes build on: here U and K hold the floor while a deliberate IFI-H rides on top.

IFI PROFILE

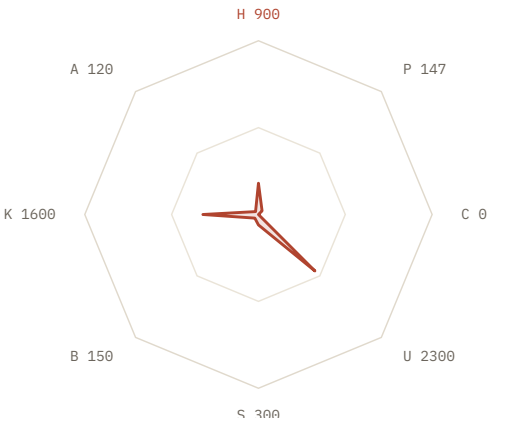

IFI-H	900
IFI-P	147
IFI-C	0
IFI-U	2,300
IFI-S	300
IFI-B	150
IFI-K	1,600
IFI-A	120

A fourth build shows savory depth carrying a second axis openly — the **miso-caramel glaze**, a savory-sweet finisher where umami and kokumi hold a base and a deliberate sweetness rides on top, the swicy-adjacent cousin of the pastes above.

```
MISO-CARAMEL SAVORY GLAZE            yield 300 mL · total 20 min
—
120   g    koji-fermented soybean paste (red miso-style)
80    g    restored tomato paste
60    g    sugar or glucose syrup
40    g    water
8     g    engineered yeast extract
—                                    a squeeze of citrus to finish
```

1. Melt the sugar with the water to a light caramel — the faint bitterness of the caramel is the foil the savor leans against.
2. Off the heat, whisk in the miso-style paste, tomato paste, and yeast extract; do not boil the miso hard, to spare its live enzymes and its γ-glutamyl body.
3. Return to the barest heat and cook to a glaze that coats a spoon; finish with citrus. Brush onto roasted vegetables, tofu, or anything from the grill.

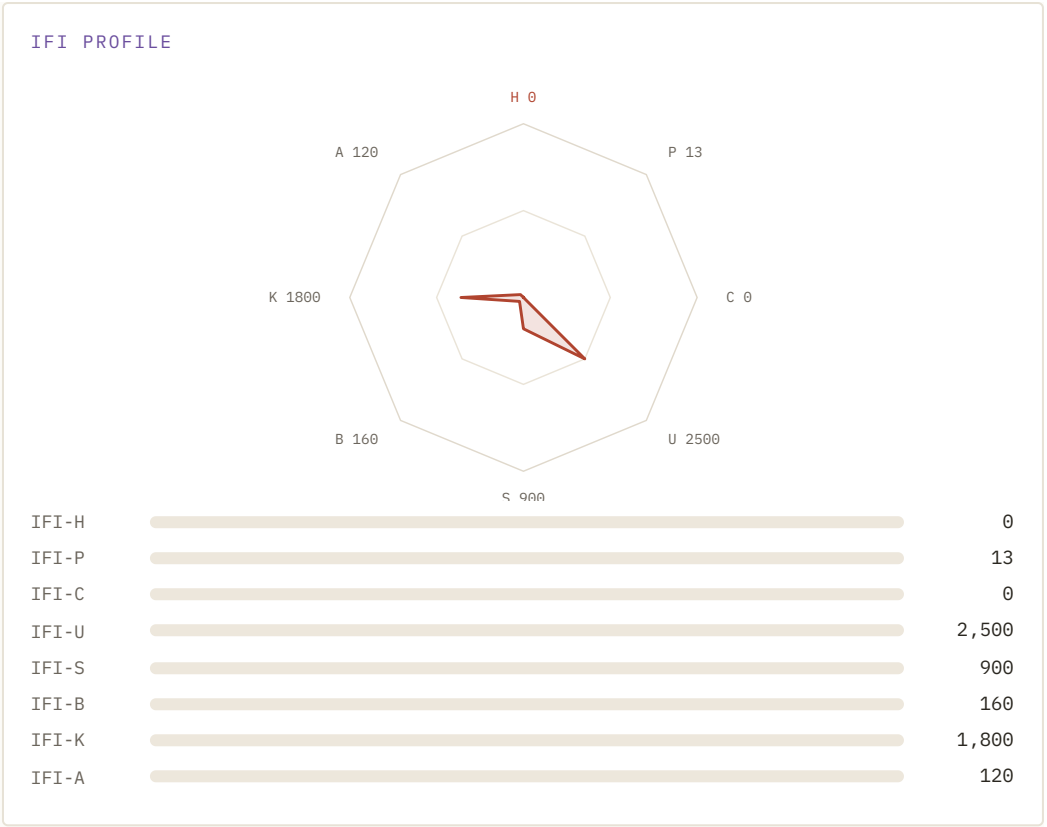

Across the four builds, the pattern is the lesson: concentration multiplies the note, the organisms supply the body that reduction alone cannot, and the Index lets you see — before anyone tastes — exactly where on both axes a condiment will land.

These condiments are also the clearest place to watch savor serve as a *base* for the rest of the library's axes, which is the thesis of the Intersections shelf this book belongs to. The miso-caramel glaze already carries sweetness openly; dial a designed pepper into the mushroom XO and you have the heat-over-savory of a chili crisp;

fold a cooling or pungent note into a dipping sauce and the savory floor holds it all up. A sauce is where a cook composes axes in real time, and a savory paste built to a known U×K profile is the dependable foundation that composition stands on — the reason a professional kitchen keeps a deeply savory base at hand and builds the day's sauces up from it. Depth first, then everything else: it is true of a sauce, and it is true of the whole Index. It is why a formulator who masters the savory paste has, in effect, mastered the base of every condiment the library can compose — the single most leveraged thing to get right on the whole sauce bench. Get the paste right and the whole condiment shelf follows; get it wrong and no amount of finishing rescues the thinness underneath.

Chapter 16 · The honest meat analog — *savor that does not pretend.*

COMPOSITION · the four-part savory system · plant-first, respectfully

The savory manufacturer's hardest brief is the meat analog, and it is hardest because of savor. Texture can be engineered, color can be managed, but the thing that most often betrays a plant-based product as an imitation is a savor that is thin, flat, and *front-loaded* — all note and no body, a burst of umami that vanishes, where real meat gives a savor that is deep, slow, and continuous. We take the plant-first position the whole library takes: the goal is not to counterfeit meat but to give a plant-based food a savory depth honest enough that it needs no apology. And that depth is, precisely, a **U×K** problem — which is why this book, not a texture manual, is where it is solved.

Meat's savor is a four-part chemistry, and a plant system earns its depth by reproducing each part from a grown source rather than from a drum of additives. The parts are: a **glutamate note**; a **nucleotide gain**, which in meat is the post-mortem IMP of Chapter 6; a **kokumi body**, which in meat is the glutathione and γ-glutamyl peptides of muscle and the peptides generated by cooking; and a **fat-borne continuity**, the way rendered fat carries and prolongs all of it. Our organisms supply the first three directly, and the fourth is a formulation choice. Here is the savory system as a dry-plus-wet seasoning stack a formulator folds into a plant protein base:

```
THE HONEST ANALOG — SAVORY SYSTEM     formulation · % of seasoning phase
—
34.0  %    engineered yeast extract (note + IMP-upgraded gain + body)
20.0  %    koji-fermented soy solids (note + body, slow-release)
14.0  %    restored tomato powder (note)
10.0  %    shiitake powder (guanylate gain)
8.0   %    low-pungency garlic (body)
6.0   %    roasted onion powder (body)
5.0   %    yeast-derived glutathione concentrate (body, direct)
3.0   %    smoke / Maillard flavor to taste
—                              dose seasoning phase at 4–6% of total mass
```

1. Build the savory system as a pre-blend and verify its U and K against target *as a system*, before it meets the protein — the Index reading of the seasoning phase predicts the finished product's depth.
2. The **AMP-deaminase-upgraded yeast** is doing the job meat's autolysis does: delivering inosinate, the strong synergist, rather than the weaker adenylate. This single choice is most of the difference between a plant savor that reads as deep and one that reads as flat.
3. The **koji solids and the glutathione concentrate** are the body — the slow, continuous savor that keeps the product tasting of something through the last chew. Front-loaded umami without this body is the single most common failure of the category.

4. Pair with a structured plant fat for the fourth part; the fat's job is continuity, carrying the water-soluble savor and prolonging it.

IFI PROFILE

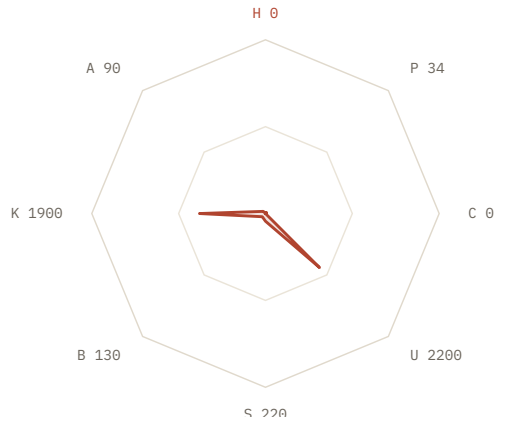

IFI-H	0
IFI-P	34
IFI-C	0
IFI-U	2,200
IFI-S	220
IFI-B	130
IFI-K	1,900
IFI-A	90

Assembled into a finished product, the savory system meets a protein base, a binder, and a fat, and the Index reading of the whole is what a panel will taste.

THE HONEST PATTY — yield 6 patties · total 30 min

300	g	textured plant protein (pea / soy), hydrated
120	g	cooked legume or mushroom, minced (bite + glutamate)
60	g	structured plant fat, chilled (continuity)
35	g	savory system, as dry pre-blend (above)
20	g	binder (methylcellulose or starch)
12	g	beet / malt for color
—		salt last, to the seasoned taste

1. Fold the hydrated protein and minced legume-or-mushroom with the savory system *before* the fat, so the water-soluble savor distributes through the whole matrix rather than only the surface.
2. Work in the chilled fat in pieces — it should stay discrete, to render and carry savor on the griddle the way animal fat does.
3. Bind, form, and rest cold. Cook on a hot surface to a hard Maillard crust on both faces.

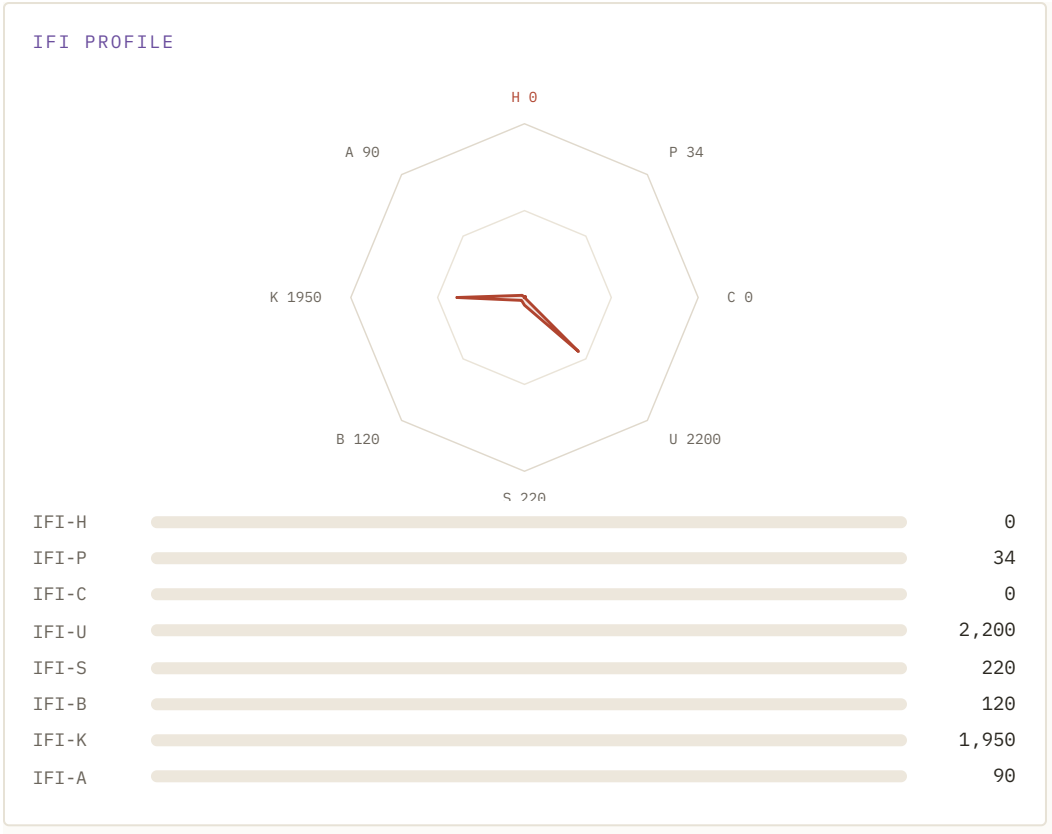

The crust is not cosmetic; it is a second savory reactor. The **Maillard reaction** between the base's reducing sugars and its free amino acids — including the glutamate we engineered in — generates the roasted, meaty volatiles a raw savory base lacks, and the sulfur amino acids **cysteine and methionine** throw the specifically meaty thiols and thiophenes. Our U×K base is the *foundation* those reaction flavors build on, not their replacement: the Maillard crust supplies the top-note of *cooked*, the organisms supply the depth of *savory*, and the structured fat carries and prolongs both. A patty that skips any of the three betrays itself — all crust and no depth, or all savor and no roast, or deep and then abruptly gone for want of fat. The honest analog is the one that gets all three right, which is why it is a U×K problem wearing a texture problem's clothes.

Two extensions finish the category honestly. The first is the **fat**, because its role in savor is more than richness. A structured fat with the right melting behavior carries the water-soluble savory compounds, releases them as it warms in the mouth, and coats the palate so the perception lingers — the physical basis of what a taster calls *juiciness* and *length*. A lean analog tastes abrupt not only for want of kokumi but for want of this carrier; the fat and the γ-glutamyl peptides are two routes to the same continuity, and the best builds use both. The second is **seafood**, written here with the same respect as meat and not as an afterthought. A plant-based seafood analog is a slightly different savory problem: the note leans more on glutamate than on the inosinate of warm-blooded muscle, the body wants the sweet-briny kokumi of scallop and shellfish, and the aroma asks for the marine bromophenols and clean green notes that say *sea*. The U×K base is the same — a glutamate note, a nucleotide gain, a γ-glutamyl body — tuned toward the balance a crab or a scallop keeps rather than a steer. We write the plant versions first, as the library's reverence dictates, and we write the meat and seafood reformulations with equal care, because the chemistry of savor does not change with the kingdom it serves, and neither does the respect we owe the thing on the plate. The Index reads them all on one axis, which is the plainest statement of that equality we can make.

A closing word on how an analog is judged, because savor is where the category succeeds or fails on a panel. The failure mode has a name worth borrowing — an *uncanny valley* of flavor, where a product gets close enough to a familiar savor that the small remaining wrongness reads as off-putting rather than merely incomplete. In savory analogs that wrongness is almost always a **shape** problem, not an intensity one: a note that peaks too early and drops, a body that never arrives, a savor loud but flat where the real thing is layered and slow. This is precisely what the U×K scorecard catches before a panel ever convenes — a profile with U high and K low predicts the early-peak, no-body failure on sight, and the fix is legible as a K that must climb to meet the U. We validate every analog build against a trained sensory panel, correlate the panel to the Index, and trust the number because it has earned the trust: a savory profile that reads balanced on the scorecard reads *complete* on the tongue, and completeness, not mere loudness, is what carries a plant-based product past the valley. And completeness is exactly what a grown organism delivers that a drum of additives cannot: the body built in beside the note, so the shape is already right before a formulator ever reaches for a fix.

The profile is deliberately balanced — note and body close together, around 2,200 and 1,900 — because that balance is what *slow* tastes like. A meat analog scored high on U and low on K will taste bright and brief; bring K up to meet U and it tastes long. This is the whole contribution of kokumi to the category, stated as a number a formulator can target.

A word, without apology, about meat itself. The same four-part system reads straight across to **meat reformulation** — the growing work of lifting savor in a product so that less salt, or less of a scarce cut, delivers the same depth. A grown glutamate-and-nucleotide stack lets a real-meat product hold its savor at lower sodium, and a kokumi load lets it taste full at a leaner formulation. We write the plant version first because that is where our reverence starts, and we write the meat version with equal respect, because savor is savor and the chemistry does not care which kingdom it serves. The Index scores both the same, which is the point: one measure, honest about every plate it reads.

Chapter 17 · Ferments — *time as an ingredient.*

COMPOSITION · koji · miso · shoyu · the slow savory build

The oldest savory technology is the one that works while you are not watching. Fermentation is how human cooking first learned to make deep savor on purpose — not by adding it, but by setting living enzymes loose on a substrate and letting them *release* savor over weeks and months — and it is where our liberating organisms, koji above all, do their fullest work. A ferment is a reactor in which the two savory chemistries are generated in place: proteolysis and glutaminase liberate the glutamate note, and the slow action of transpeptidases and aging builds the γ-glutamyl body. Time is not a cost of the ferment. Time *is* the ingredient — the thing that does the work — and our engineering is, largely, the art of making the ferment do more of that work in less of it.

The foundational build is a **koji shoyu-style sauce** — soy sauce's architecture, compressed by a high-glutaminase koji so the savor arrives in weeks rather than the traditional many months.

```
KOJI SHOYU-STYLE SAUCE                yield ~1 L · active 1 h · ferment 6–
10 wk
—
500   g    engineered soybean, cooked soft (bound-glutamate reservoir)
500   g    cracked roasted wheat (carbohydrate + aroma)
—     —    high-glutaminase koji culture (A. oryzae)
230   g    salt
1.2   L    water (→ brine, ~18% salinity)
—                              warm, patient, stirred daily at first
```

1. Inoculate the cooked soybean and roasted wheat with the high-glutaminase koji and incubate warm and humid for ~48 hours, until the grain is bound in white mycelium — this is the koji, the enzyme factory.
2. Slurry the koji into the salt brine to make the *moromi* mash. Now the release begins: the koji's proteases chop the soybean's glycinin into free amino acids and peptides, and its **glutaminase** converts the liberated glutamine into glutamate — the step that turns a silent protein reservoir into savor.
3. Ferment warm, stirring daily at first then weekly, for six to ten weeks. The high-glutaminase koji is doing in weeks what a traditional strain does in six months to a year.
4. Press and pasteurize gently. The pressed sauce carries free glutamate (the note), the kokumi γ-glutamyl peptides generated across the ferment (the body), and the deep color and aroma of the slow amino-carbonyl reactions of aging.

IFI PROFILE

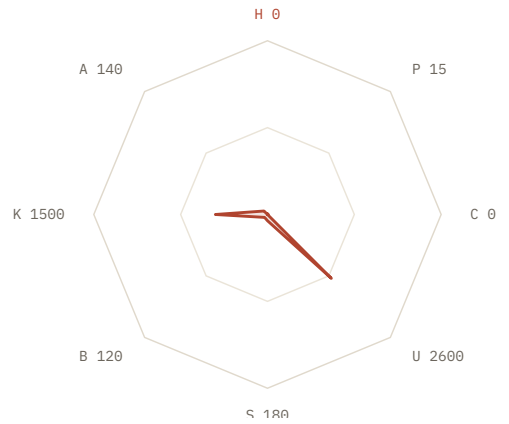

IFI-H	0
IFI-P	15
IFI-C	0
IFI-U	2,600
IFI-S	180
IFI-B	120
IFI-K	1,500
IFI-A	140



FIG. 13 — Soybeans, whole. The ferment's substrate: a protein reservoir that koji's glutaminase turns loud, the note and the body generated in place over weeks.

A **miso** is the same chemistry held as a paste rather than pressed to a liquid, and it shows the time-depth relationship most clearly of anything we make. A young white miso, fermented weeks, is light and sweet, its proteolysis barely begun — modest note, little body. An aged red miso, fermented a year or more, is dark, profound, and *round*, because the long ferment has liberated far more free glutamate and generated far more γ-glutamyl kokumi. The two are the same ingredients and the same organisms; the only variable is time, and the Index reads time directly as a climb up both axes. Our engineered soybean and high-glutaminase koji let a miso reach an aged-red depth in a fraction of an aged-red's year — the time-depth curve, steepened.

FERMENT	TRADITIONAL TIME	ENGINEERED TIME	DRIVES
White miso	2–8 weeks	1–3 weeks	modest U, low K
Red miso	1–3 years	3–6 months	high U, high K
Shoyu	6–18 months	6–10 weeks	high U, moderate K
Quick amazake	8–12 hours	4–6 hours	sweetness, trace U

The microbiology underneath these builds is a succession, not a single organism, and knowing the cast lets a manufacturer design the ferment rather than inherit it. Koji (*A. oryzae*) is an **aerobic, solid-state** first act, growing on the surface of steamed grain and soy and laying down the enzyme arsenal of Chapter 12. The *moromi* mash that follows is **anaerobic and salt-tolerant**, and there the baton passes to a halophilic community: the lactic bacterium **Tetragenococcus halophilus** drops the pH, and the yeast **Zygosaccharomyces rouxii** builds the aromatic esters and alcohols of a mature shoyu. Other savory ferments run an entirely different cast — **natto** is **Bacillus subtilis** var. *natto*, which spins the beans into sticky threads of γ-polyglutamate and a sharp, pungent savor; **tempeh** is the mold **Rhizopus oligosporus**, binding the beans into a nutty cake with a gentler proteolysis. Each organism is a different way of releasing the same soybean reservoir, and each is sped by starting from a seed already partway loud.

The gateway ferment for any kitchen — the one that turns this whole chapter into a Tuesday-night tool — is **shio koji**, a salted koji slurry that becomes a living seasoning.

```
SHIO KOJI (living savory seasoning)   yield ~500 mL · ferment 7–10 days
—
200   g    rice koji (A. oryzae on rice)
60    g    salt
250   mL   water
—                                          stir daily; it thins as it works
```

1. Combine the koji, salt, and water; the mix begins thick and granular.
2. Keep at warm room temperature, stirring daily. Over a week the koji's proteases and amylases liberate free amino acids and sugars; the slurry thins, sweetens, and turns savory-rich.
3. When it smells of sweet chestnut and tastes deeply of itself, refrigerate. Use it as a marinade, a brine, or a finishing seasoning — a spoon brings glutamate and a gentle kokumi to anything it touches.

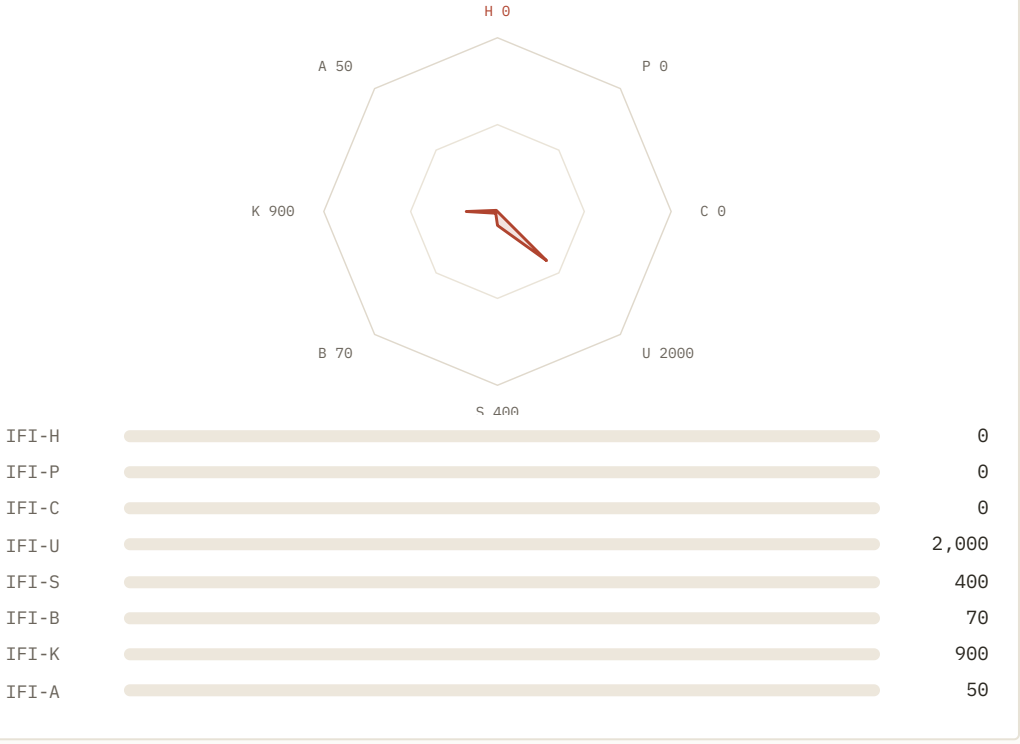

Shio koji is the ferment chapter in miniature: no additive, no fire, only a mold's enzymes and a week of patience, releasing savor from rice the way the great ferments release it from soy. It is where a reader stops reading about the method and starts keeping it alive in a jar on the counter.

A word on why time works the way it does, because the curve has a shape worth trusting. Proteolysis in a ferment follows roughly **first-order kinetics**: the rate of glutamate release is proportional to the protein still un-cut, so savor climbs fast at first — when there is the most substrate — and then slows as the reservoir empties, approaching a plateau rather than rising forever. This is the mathematical face of the familiar truth that the first weeks of a ferment change it most and the last months refine rather than transform it. **Temperature** sets the clock: as a rule of thumb, every ten-degree rise roughly doubles the enzymatic rate — the Q_{10} of the chemistry — which is why a warm miso matures in a season and a cold-cellar one takes years, and why a controlled warm hold is the safest lever for speeding depth without distorting it. Our engineering pushes on both terms at once: a high-glutaminase koji raises the rate constant, and a seed that starts partway released raises the substrate the curve begins from, so the ferment reaches a target depth sooner and lands it more reproducibly. Time is still the ingredient. We have only taught the plant and the culture to spend it faster.

Because a ferment is a living process, it carries a safety discipline a manufacturer must respect, and it is one more place where control replaces luck. The traditional ferments are safe for a reason that is pure chemistry: **salt and acid are hurdles**. A shoyu moromi is too salty for pathogens to grow; a miso too salty and too acidic; a lactic ferment drops its pH fast enough that spoilage organisms lose the race to the lactic bacteria. The danger lives in the gaps — a low-acid, low-salt mixture held warm without a dominant culture is exactly the condition a careful maker designs against — and the defenses are well understood: a vigorous starter that wins early, a salt or acid hurdle that holds, and process controls on temperature

and time. Our engineered organisms do not relax any of this. A faster, more dominant koji and a seed that releases on schedule make the ferment *more* controlled, not less. The oldest savory technology and the newest agree on one rule — a ferment is safe when the right organism is firmly in charge — and designing the organism is one more way of making sure it is.

The ferment chapter closes Part IV on the book's deepest idea: that savory depth is not a thing you add but a thing you *release* and *grow*, and that the oldest kitchens knew this — they just called it patience. Our organisms do not replace the patience so much as concentrate it, packing the enzymatic work of months into the plant and the culture themselves. What a long ferment taught the tongue, we have learned to measure, tune, and lock. The last part is what it takes to lock it — and to put it, honestly, on a label.



PART V OF V

Precision, Label & Frontier

the ≤7% CV · the clean line · the roadmap · the table

Chapter 18 · Precision & the ≤7% CV — *a savor you can promise.*

QUALITY · reproducibility · the seven-percent bar · why it is harder for savor

A flavor you cannot reproduce is not a product; it is an anecdote. This is the sentence that governs the whole library, and it governs savor most strictly, because savor is the axis a manufacturer most needs to be *boring* — identical bowl to bowl, batch to batch, year to year — and the axis that, left to nature, varies the most. A tomato's free glutamate swings with the season, the soil, the day it was picked; a ferment's depth drifts with the strain, the temperature, the length of the moromi. The entire commercial promise of designing savor at the source rests on collapsing that variance, and the number we hold it to is a **coefficient of variation of seven percent or less** — the same bar that governs the engineered heat of the flagship peppers, carried across to the savory library.

Seven percent is a demanding figure, and it is worth seeing what it means. Coefficient of variation is the standard deviation divided by the mean — a unitless measure of how tightly a property clusters around its target. A commodity tomato's free-glutamate content can swing with a CV of forty percent or more between fields and seasons; a traditional ferment, more.

Where does nature's noise come from? From everything the field cannot control. A tomato's savor swings with the **season** (light and heat drive the ripening climb of Chapter 9), the **soil** (nitrogen supply feeds the glutamate engine), the **cultivar**, and the **day of harvest** relative to the breaker stage — four uncontrolled variables multiplying into that forty-percent swing. A ferment adds the drift of a **living culture**: strain variation, ambient temperature, the length of the moromi, the wild organisms that colonize an open mash. Every one of these is a source of variance the traditional maker accepts as the price of depth. Designing savor at the source does not abolish the variables; it moves the decisive ones **off the field and into the organism and the process**, where a defined genotype, a defined feedstock, and a defined release step stand in for the season, the soil, and the weather. You cannot control a July. You can control a fermenter. Holding a *designed* savory organism to a CV at or below seven percent means that when a formulator specifies an IFI-U of 2,600, what arrives reads between roughly 2,420 and 2,780 — a window narrow enough that the finished product tastes the same to a consumer who could never articulate why a competitor's does not.



FIG. 14 — Isolates on the bench. Every savory batch is read against its target profile before it ships; the seven-percent window is the difference between a flavor and an anecdote.

How do you hold savor that tight? The answer is the thread running through the whole book: you move the variance from the *field*, where you cannot control it, to the *design*, where you can. Three choices do most of the work. First, **grow rather than farm** wherever possible — the engineered yeast of Chapter 11 is a fermenter product on a defined feedstock, with no season, no weather, no shipping loss, which is why its CV is the easiest in the library to hold and why it is the manufacturer's favorite. Second, **stage the release** so the tasteable savor is produced by a controlled process step — an autolysis, a warm hold, a defined ferment length — rather than by the uncontrolled accidents of ripening and aging; a process you run is a process you can repeat. Third, **make the smallest edit that moves the number**, because, as Chapter 8 argued, a conservative single-residue change holds its behavior across batches where an aggressive pathway overhaul backs up into side-products and drifts.

There is a subtlety peculiar to savor, and it is the synergy that is otherwise this book's great gift. Because umami is *multiplicative* — the glutamate × nucleotide product of Chapter 3 — variance in the two inputs **compounds** in the output. If glutamate varies by seven percent and nucleotide varies by seven percent, their product can vary by more, because the errors multiply along with the signal. This is why we hold the *nucleotide* release especially tight, and why the thermostable phosphodiesterase of Chapter 8 (the L212F edit) matters beyond its yield: a release enzyme that behaves identically across the processing window keeps the gain term stable, and a stable gain term is what keeps the multiplied output inside its seven-percent box. Precision in savor is not eight axes held independently; it is the discipline of holding the *synergy* steady, which is a harder and more interesting problem than holding any single concentration.

The arithmetic of that discipline sets the targets the whole supply chain is held to. For a multiplicative quantity — the synergy product **P = u·v** of Chapter 3 — relative errors add in quadrature:

```
(σ_P / P)² ≈ (σ_u / u)² + (σ_v / v)²
```

So if the glutamate and the nucleotide each vary with a coefficient of variation of seven percent, their product varies by $\sqrt{(0.07^2 + 0.07^2)} \approx$ **9.9 percent** — the synergy is noisier than either input that makes it.

Read the implication backward and it becomes a manufacturing law: to hold the *perceived* savor, which rides on the product, inside seven percent, each multiplicative input must be held near **five** percent, since $\sqrt{(0.05^2 + 0.05^2)} \approx 7.07\%$. This is why the nucleotide release is held tightest of all, and why "hold the synergy steady" is not a slogan but a specification that flows down to the glutamate pool, the RNA fraction, and the phosphodiesterase alike.

Holding it requires measuring it three ways, because no single instrument sees the whole of savor. **Liquid chromatography–tandem mass spectrometry (LC-MS/MS)** quantifies the actual molecules — free glutamate, 5′-IMP, 5′-GMP, glutathione, γ-Glu-Val-Gly — to parts per million: the ground truth of what a batch contains. The **receptor cell assay** of Chapter 4 is the *functional* check that catches what chemistry can miss — a matrix effect, an inhibitor, a synergy the concentrations predict but the tongue would not get. And a **trained sensory panel**, correlated to both, is the final arbiter, because the tongue is what ships and a number that does not predict a taste is a number we do not trust. Each batch's eight-axis profile is plotted on a **control chart** against its locked target with action limits, and a point outside them stops the lot — the same Shewhart discipline a pharmaceutical line runs, turned onto flavor.

None of this is asserted on faith. The method behind the savory library stands on **three patents filed and twenty-one peer-reviewed papers**, published openly on Zenodo under CC BY 4.0, so that the reproducibility we sell is reproducibility anyone can check. A savor you can promise is, in the end, a savor someone else can measure and find you were right about.

Reproducibility is not only a property of the batch that ships; it is a property of that batch a year later, which is why savory QC runs forward in time as well as across lots. The fragile term is the nucleotide: residual phosphatase, trace moisture, and heat will slowly dephosphorylate a guanylate pool over months, quietly draining the synergy even as the glutamate holds. So a savory ingredient is placed on **accelerated shelf-life** study — aged at elevated temperature to compress months into weeks — and its 5′-nucleotide content tracked to predict how its IFI will drift on a real shelf. The defense is again the process: the release step that inactivates the degrading enzymes, a dry format with low water activity, packaging that keeps out moisture and oxygen. Every lot is held as a **retained sample** against its own future, traceable back through its biomass, its feed, and its edits, so a profile that drifts can be diagnosed to its cause rather than merely discarded. Precision across space is holding a thousand lots to one target; precision across time is holding one lot to that target until the last jar is opened. Both are the same promise, and the savory library makes both. A number that holds across a thousand lots and across a year is not a marketing figure; it is the whole difference between a flavor a manufacturer can build a product line on and one they can only hope for.

The whole library is built on this. Across the library's roughly **19,500 product SKUs** — backed by approximately 225 distinct construct designs — with **a thousand live and purchasable today** and the quality bar held at seven percent, the savory organisms are the proof that reproducibility and depth are not opposites. The old trade was always *deep or consistent, pick one* — the long-simmered stock that is profound and never twice the same, or the additive blend that is identical and flat. Designing savor at the source dissolves the trade. A broth can be both profound and boring, in the best sense of the word: the same miracle, every single time.

Chapter 19 · Label, compliance & the clean line — *honesty on the back of the pack.*

REGULATORY · ingredient-sourced savor · allergens · attributes, never claims

A savory ingredient has to survive two readers: the tongue and the label. We have spent eighteen chapters on the first; this one is for the second, because a depth that cannot be declared honestly on a package is a depth a manufacturer cannot use. The good news is that designing savor at the source produces a cleaner label almost by accident — but only if the claims made about it are exact, and the regulatory ground here is full of traps for the overeager. We walk it carefully, because the library's rule is older than any regulation: nourishment and quality are values and intentions, never claims.

Start with the molecule everyone worries about. The glutamate that our restored tomato and grown yeast deliver is **chemically identical** to the glutamate in monosodium glutamate — glutamate is glutamate, and the body reads it the same whether it arrived from a drum or from a fruit. We say this plainly because the opposite, implied claim — that ingredient-sourced glutamate is *safer* than added MSG — is both false and something we will not print. The value of sourcing savor from whole, grown ingredients is **not** a safety claim; added glutamate is recognized as safe, and we will not trade on a fear we do not share. The value is twofold and both parts are honest: a **label built from ingredient names** (tomato, yeast extract, fermented soybean, shiitake) rather than from a row of isolated additives, which is what a growing number of consumers and buyers prefer on its own terms; and a **flavor that carries the kokumi body** an additive blend simply does not, because the organisms deliver the γ-glutamyl peptides that a drum of glutamate and a drum of inosinate leave out.

Now the trap, stated precisely so no one walks into it. In the United States, the FDA requires added MSG to be declared as "monosodium glutamate," while glutamate naturally present in an ingredient is labeled by that ingredient's name. But the FDA also considers it **misleading to claim "No MSG" or "No added MSG"** on a product that contains an ingredient which is itself a significant source of free glutamate — hydrolyzed protein, autolyzed yeast extract, and the like. This matters directly to us: a product built on our engineered yeast extract carries free glutamate by design, so it **must not** wear a "no MSG" flag. The honest and defensible position is the ingredient-forward one — *made with tomato, fermented soybean, and yeast* — not the additive-negative one. We would rather a manufacturer say what is in the food than what is not, and the regulation, read correctly, agrees. And "clean label" itself, worth saying plainly, has **no legal definition** anywhere; it is a consumer-preference claim, not

a regulated one, which is exactly why the honest version of it is an ingredient list a shopper can read rather than a negative flag a regulator can challenge.

Outside the United States the rules are written as numbers, and a savory ingredient ships to all of them. In the European Union the glutamates carry the E-numbers **E620–E625** — glutamic acid and its sodium, potassium, calcium, ammonium, and magnesium salts — and the savory ribonucleotides carry **E626–E635**, the guanylates, the inosinates, and the disodium 5′-ribonucleotides of E635. For decades the Joint FAO/WHO Expert Committee placed glutamic acid and its salts in its safest category, an acceptable daily intake **"not specified."** That changed in 2017, when the European authority re-evaluated the glutamates and set a **group ADI of 30 mg per kilogram of body weight per day**, expressed as glutamic acid — a genuine regulatory event a serious manufacturer now tracks, and the reason the sophistication cuts against the naïve clean-label instinct. Because that ADI is a ceiling on *total* glutamic acid from every source, a formulator leaning hard on ingredient-sourced glutamate is not exempt from it merely for avoiding the additive line; the glutamate is counted wherever it came from. This is the deepest argument for designing to **loudness rather than concentration**: the synergy of Chapter 3 reaches a target savor with *less* total glutamate, which is the one approach that satisfies the clean label and the intake ceiling at once. The international baseline, the Codex **General Standard for Food Additives**, permits these compounds under good-manufacturing-practice or numerical limits, and national codes layer their own labeling on top — much of Asia treating them as long-ordinary ingredients, other markets demanding the additive class name or the specific compound. An ingredient designed for the world is designed to the strictest cell of that matrix, not the most lenient.

SAVORY INPUT	DECLARES ON LABEL AS	REGULATORY NOTE
Restored tomato	"tomato" / "tomato paste"	whole-ingredient; no additive line
Engineered yeast extract	"yeast extract"	a free-glutamate source — no "no MSG" claim
Koji-fermented soybean	"fermented soybean (soy)"	**soy allergen**; may carry wheat
Shiitake concentrate	"shiitake mushroom"	ingredient-sourced guanylate
Low-pungency garlic / onion	"garlic" / "onion"	whole-ingredient body

Allergens are the second compliance axis, and savor runs straight through the two biggest. **Soy** is a major allergen under FALCPA and its equivalents worldwide, and the soybean is central to this book; any product built on engineered soybean, miso-style paste, or shoyu must declare it. **Wheat and gluten** enter through traditional shoyu, where roasted wheat is half the mash, and through any koji grown on grain — so a gluten-free savory build must route around wheat deliberately, using tamari-style all-soy ferments and rice-grown koji. These are not obstacles to engineering; they are facts of the substrates, and a formulator who knows them designs the allergen profile as deliberately as the flavor one.

The third compliance axis is the quietest and the one we hold hardest: **attributes, never claims.** Savory design gives a formulator real, declarable capabilities — a product can hold its depth at lower added sodium, because kokumi makes salt taste fuller and umami carries savor that salt used to carry; a plant-based product can taste complete without animal ingredients. These are **formulation attributes**, and we state them as such: *holds savor at reduced sodium, delivers depth from plants.*

The sodium attribute deserves its exact, careful wording, because it is where the temptation to overclaim is strongest. Umami and kokumi both make a given amount of salt taste *saltier and fuller* — umami because the savory signal carries perceived intensity that sodium used to carry alone, kokumi because CaSR modulation lengthens and rounds the salty percept. In practice, formulators report holding a product's perceived saltiness and overall acceptability while cutting **twenty to forty percent of the added sodium**, replacing what salt was doing for *flavor* with what the organisms do for *savor*. We state that as a formulation attribute — *maintains perceived saltiness at reduced added sodium* — and we stop there. We do not say the product is healthier, does anything in a body, or should be chosen for a medical reason; the attribute is a capability we hand the formulator, and what they claim within their jurisdiction's rules is theirs to substantiate. There is also a quieter positioning benefit the ferments carry almost for free: a savory depth built from tomato, soybean, yeast, and mushroom is, by construction, **plant-based**, and — with the allergen routing above — can be made **vegan, and often kosher or halal, by design** rather than by substitution. The depth does not have to be added back after the animal is removed; it was never animal to begin with. What we do not do — ever, anywhere in the library — is convert an attribute into a health claim. We do not say a food "reduces" anything in a body, does not "detox," does not "cure," carries no "health halo." The nourishment that runs under this whole book is a reason we do the work, not a promise we print on it. A broth grown to carry more savor from fewer, wholer ingredients is a good thing to give your body; that it is good is for the person eating it to feel, not for the package to assert.

A last labeling trap deserves naming, because it catches the well-meaning. The term **"natural flavor"** has a regulatory definition — a flavor derived from a biological source — but no clean-label virtue; a long "natural flavor" on an ingredient deck tells a shopper nothing and increasingly reads as the place where complexity hides. The savory-depth approach lets a manufacturer do the opposite: name the sources outright — *tomato, yeast extract, shiitake, fermented soybean* — because those *are* the sources, and an ingredient a shopper recognizes beats a category a shopper distrusts. Front-of-pack is where the temptation to shout a negative is strongest, and where restraint pays best: a clean deck of recognizable foods needs no flag, and a flag invites the scrutiny a flag cannot always survive. The discipline is the same from the receptor to the carton — give the real thing, and skip the halo. The cleanest label is the honest one, and honesty, here, is a competitive advantage that happens also to be right.

Chapter 20 · The frontier & the table — *where savor goes next.*

CLOSE · the roadmap · the nourishment through-line · the lid comes off

We began with a lid coming off a pot and a kitchen leaning in, and we have spent a book turning that reflex into receptors, enzymes, flux models, and a number you can lock to seven percent. It is worth ending by saying what all that precision is *for*, because the point of measuring savor was never to replace the pleasure of it. It was to make the pleasure reliable, reproducible, and available from plants — to let a broth taste like being cared for on a Tuesday in February, grown rather than coaxed, honest on its label and deep in the bowl. That is the whole of the work, and it is not finished.

The frontier of savory depth runs in three directions. The first is **more organisms**. The library holds forty-two host species today, and the savory ones — soybean, yeast, koji, the mushrooms, the alliums, the restored tomato — are a fraction of what carries glutamate, guanylate, or a γ-glutamyl peptide somewhere in its metabolism. Every legume is a potential kokumi organism; every protein-dense seed is a latent umami reservoir; every fungus is a possible release engine. The soybean, which we grade *frontier* for its difficulty, is the proof that the hard organisms are worth the cost: one seed, both axes, the complete savory stack. The next decade of this shelf is the patient work of bringing more such organisms to the seven-percent bar.

The candidates are already visible. **Kelp and the seaweeds** are the obvious frontier — kombu is the loudest glutamate source in the table of Chapter 1, and a cultivated, profile-locked seaweed would be a glutamate *crop* rather than a foraged one, grown in the sea with no land and no fertilizer. The **legumes beyond soy** — lentil, chickpea, lupin, the pulses that feed much of the world — each carry a storage-protein reservoir and a γ-glutamyl store waiting to be released, and each is a kokumi organism in waiting. The **fungi** run deeper than the three we grow today: the high-glutamate boletes, the nucleotide-rich specialty mushrooms, the molds beyond *A. oryzae* that ferment other traditions. And the **protein seeds** — the oilseeds and the ancient grains — hold latent savor the way the soybean does. Forty-two validated organisms is not a ceiling; it is a start, and each of the sixty-four concept organisms on the wider roadmap is a cell in the matrix of Chapter 4 waiting to be named and grown.

The sequence in which these arrive is set by difficulty, the same grading that orders the whole platform. The easy wins are the Solanaceae and the model crops already in hand — the tomato is *Easy*, and the other nightshades and leafy models follow fast. The moderate tier is the alliums and the common legumes, a year or two of work each. The hard tier is the fungi and the specialty leaves — koji, shiitake, tea — where the genetics are less charted and every construct is dearer. And the frontier tier is the protein powerhouses

and the sea crops, where the soybean has already proved the climb is worth it and kelp, lupin, and the ancient grains are the long-horizon prizes. The roadmap is not a wish list; it is an ordered queue, each organism scheduled by how much method it will take to bring to the seven-percent bar, and each one that lands widens the savory palette the next book draws from. What makes the queue move is not a breakthrough waiting to happen but the ordinary, patient repetition of the method — read the organism, model the pathway, design the edit, grow it, measure it, lock it — run again on the next cell of the matrix. The matrix itself is effectively endless: forty-two organisms times eight axes is a field of intersections, and the savory column alone, umami and kokumi crossed against every host that carries them, could fill a shelf on its own. We pick the cells worth a book and grow the organisms worth a crop; the rest wait, named, for their turn.



FIG. 15 — Tea. Its umami is L-theanine, a glutamate kin; that a leaf brewed as a drink carries the same note as a broth is the clearest sign savor reaches further than anyone assumed.

The second direction is **composition** — the multi-axis work the Index was built to make legible. Savor rarely stands alone on a plate; it is the floor other axes build on, and the most interesting designed foods are the ones where U and K hold a base while heat, sweetness, or pungency rides on top. We have seen it already in this book's own margins: the mushroom XO where a deliberate IFI-H sits on a savory floor; the restored tomato that carries sweetness alongside its umami; the Phoenix pepper, the library's full-spectrum flagship, which registers on all eight axes at once and whose profile includes a designed **IFI-U of 1,800 and IFI-K of 1,900** — savory depth engineered into a plant the world knows only for its heat. That a pepper our model predicts at approximately 2.95 million SHU also

carries the body of a good broth is the clearest statement the library can make: the axes are independent dials, and savor is one you can turn on anything.

The composition work has its own numbers. Savor is the base coat of the multi-axis library — the dual-axis **M2** collection runs to some **8,400 constructs** and the triple-axis **M3** to **2,800** — and a large share of both are savory floors with another axis composed over them: a heat that rides on umami, a sweetness balanced against savory depth, a cooling or pungent top-note set over a broth. The reason they compose cleanly is the reason this book has circled back to again and again — the receptors are **independent**. TAS1R1·R3 for the note, CaSR for the body, TRPV1 for heat, TAS1R2·R3 for sweet: four locks, four keys, four dials a designer can set without one cancelling another. Savory depth is the dial you set first, because it is the one every other flavor stands on.

The third direction is the one this company exists for, and it is the quiet subject under every chapter here: **nourishment**. We grow savor at the source because the source is where nourishment lives too, and because the two are, finally, the same act. A plant engineered to carry more free glutamate from fewer ingredients is a plant that lets a person eat more plants and want to; a kokumi that makes salt taste fuller is a kokumi that lets a food give the mouth what it wants from less of what it does not need; a broth deep enough from soybeans and mushrooms is a broth that feeds a body well without asking it to settle. None of this is a claim, and we have been careful all book not to make it one. It is an intention — the reason the family does the computing and sets the table — and it is allowed to be only that: a way of caring about what we grow, expressed in the growing rather than the advertising. Shelf Nine of the library — *The Nourishment Frontier*, *Grown Right*, *Clean* — is where that intention becomes the explicit subject; here it has only ever been the floor beneath the flavor, the way savor is the floor beneath the dish.

There is one more thing the frontier asks for that is neither an organism nor a composition, and it is the reason this library publishes. A measure is only as useful as it is shared: an Index a company keeps to itself is a marketing instrument, while an Index anyone can read against — anchored to receptors, defined by EC_{50}, backed by twenty-one open papers and three filed patents — can become a standard a whole industry designs to. We mean the savory axes to be read that way: cited, checked, argued with, and built on, with the annual *State of Flavor* report turning the library's own data back out into the open so the numbers stay honest by being watched. Savor measured in private is a claim; savor measured in public is a fact, and we would rather build on facts.

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The profiles and scores throughout this book are the outputs of our platform's models — predictions we stand behind and test with every batch, but predictions until a greenhouse harvest and an assay confirm them.

So the lid comes off, and the kitchen leans in, and now we know why — know it down to the clamshell closing on the flytrap, the guanylate wedging it shut, the γ-glutamyl peptide widening the room. The knowledge does not spoil the moment. It is a form of reverence: to understand exactly how a broth says *you are about to be cared for*, and then to build a plant that can say it, reliably, to anyone with a pot and some water and an evening. We measured savor so that we could design it. We designed it so that the oldest comfort a kitchen offers could be grown, and given, on purpose.

That is savory depth — the note and the body, umami and kokumi, measured at the source and composed with intention. The pot is on. Lift the lid.