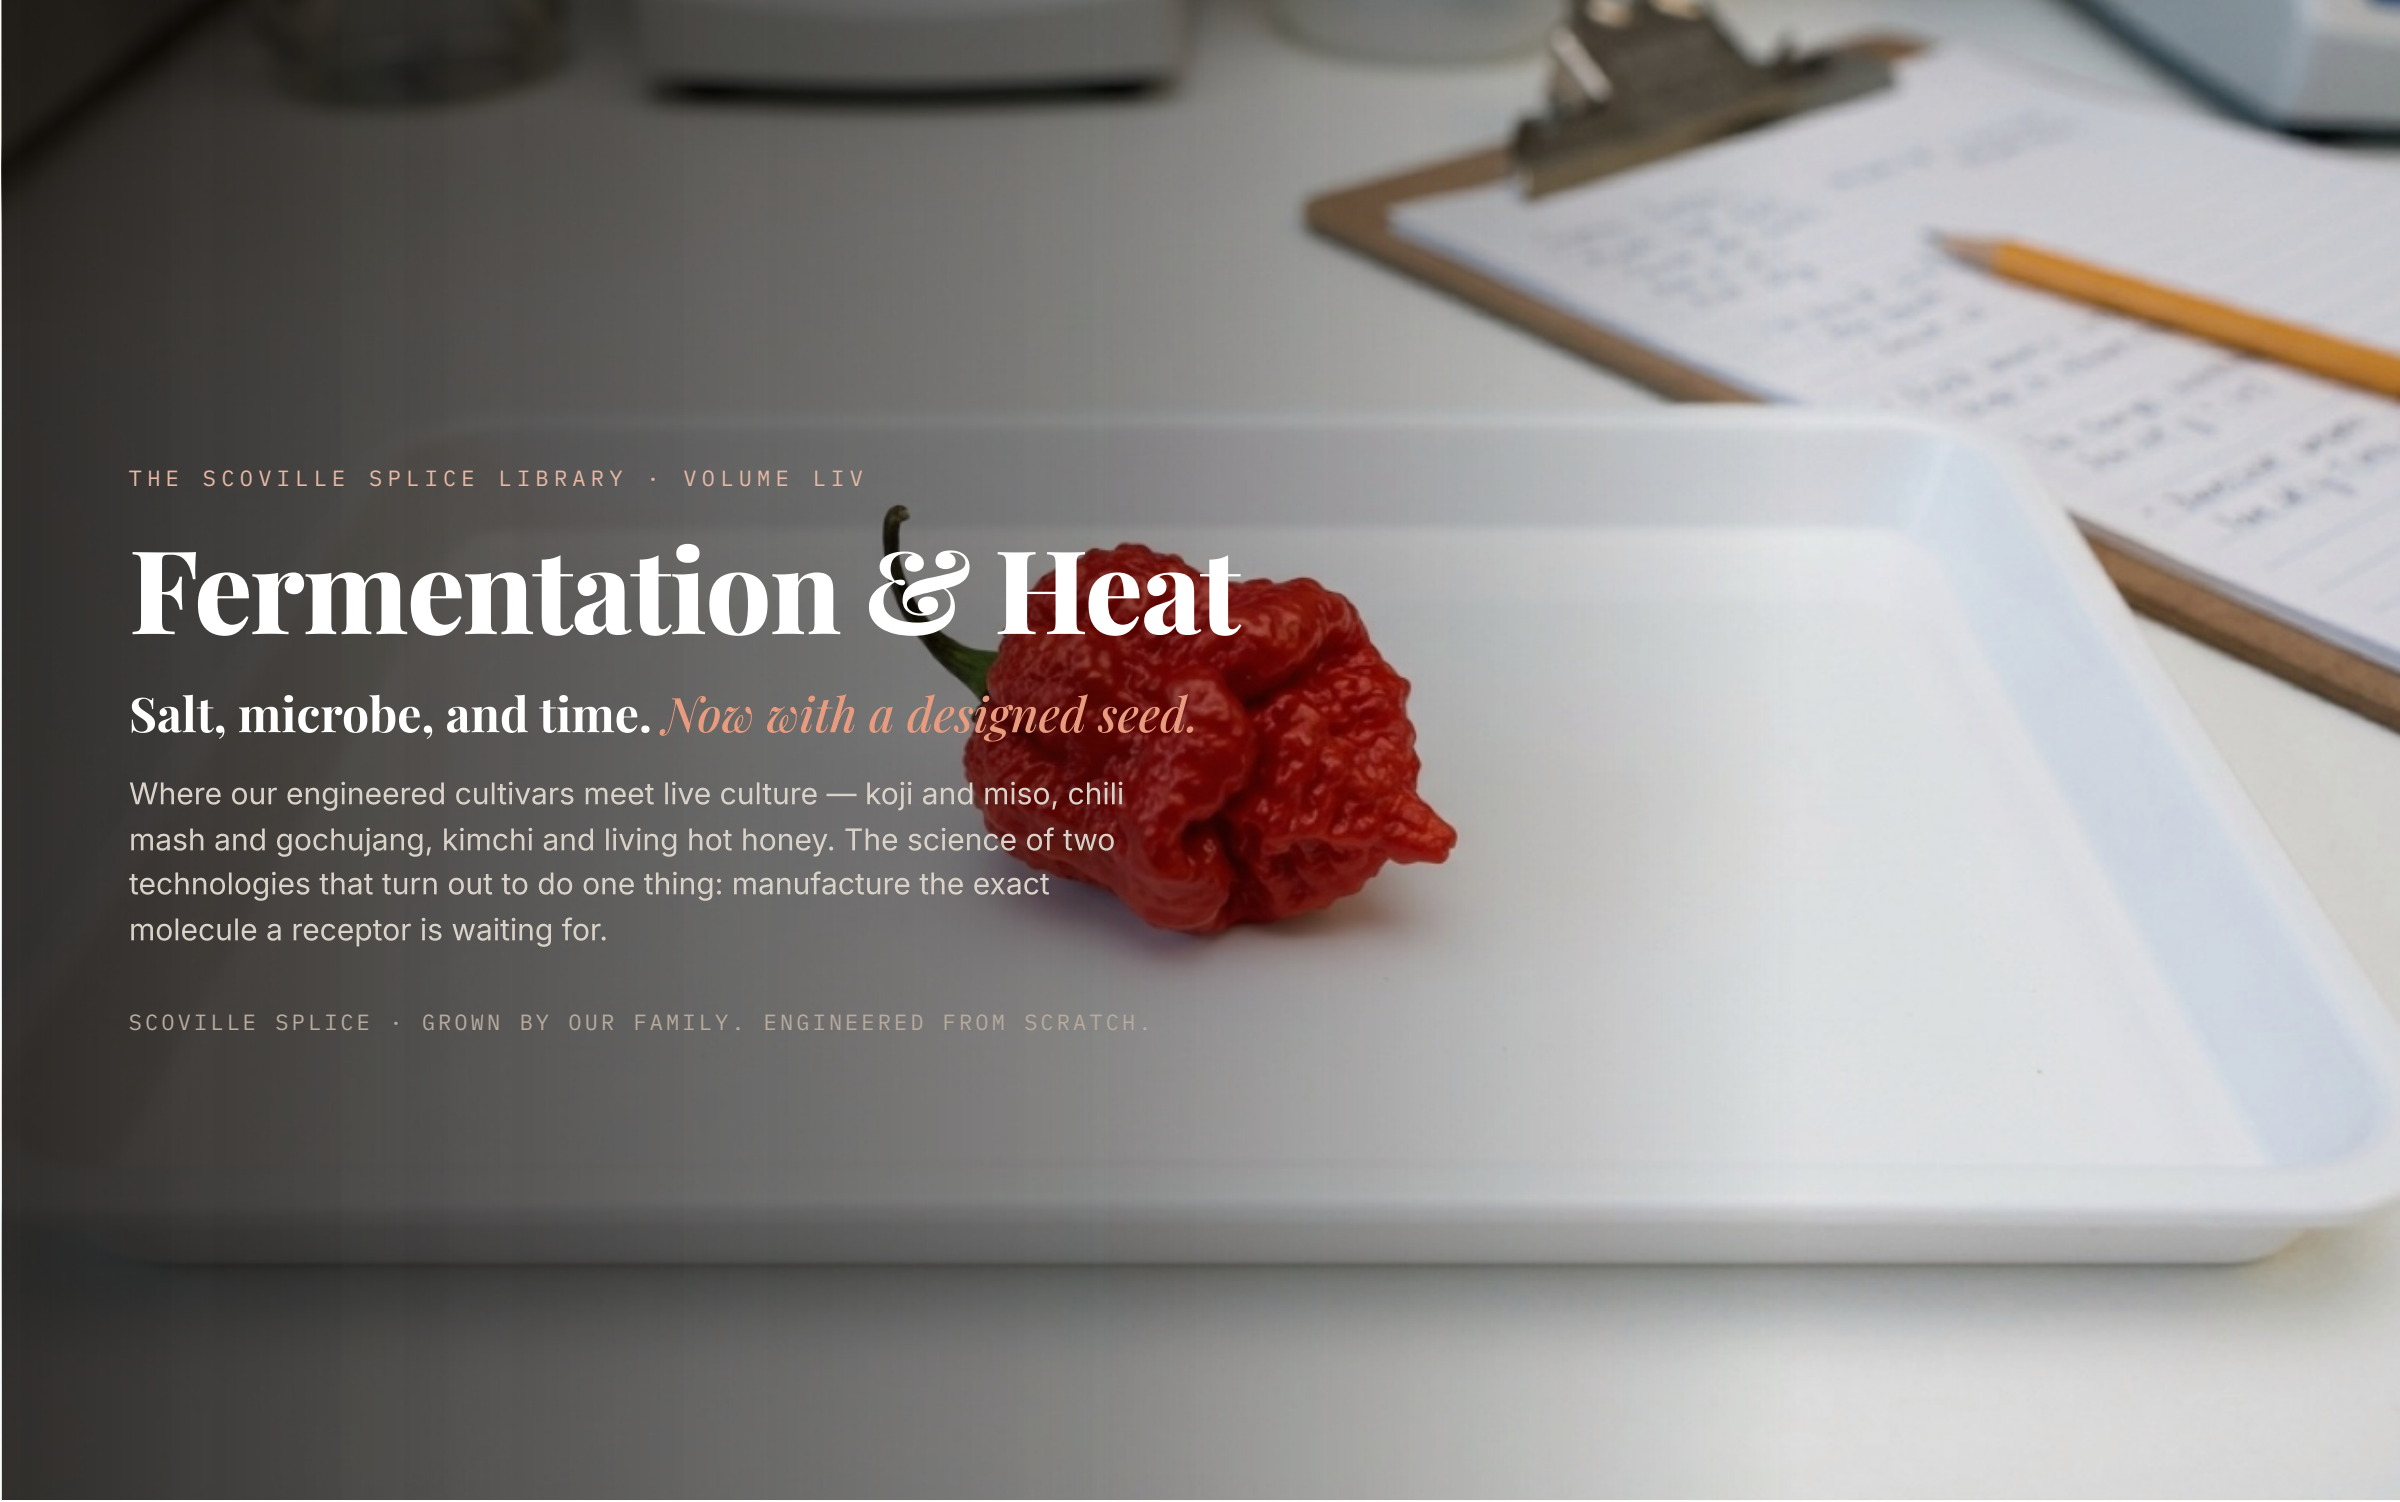

THE SCOVILLE SPLICE LIBRARY · VOLUME LIV

Fermentation & Heat

Salt, microbe, and time. *Now with a designed seed.*

Where our engineered cultivars meet live culture — koji and miso, chili mash and gochujang, kimchi and living hot honey. The science of two technologies that turn out to do one thing: manufacture the exact molecule a receptor is waiting for.

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

Contents

01 Opening essay · The crock and the construct

02 Chapter 1 · The ligand factory — fermentation is manufacturing, not magic.

03 Chapter 2 · The proton and the pepper — why a sour heat is a louder heat.

04 Chapter 3 · The multiplier — umami is not addition.

05 Chapter 4 · Koji, the machine we rebuilt — Aspergillus oryzae, eighty constructs deep.

06 Chapter 5 · The bean that holds the depth — soybean, three hundred and sixty constructs.

07 Chapter 6 · Peppers bred for the crock — a thousand and sixty constructs, a different spec for the jar.

08 Chapter 7 · The kimchi board — cabbage, radish, allium, ginger, and the culture that eats them.

09 Chapter 8 · The quiet multiplier — shiitake, mushroom, tomato, yeast, and the nucleotide side of savor.

10 Chapter 9 · Rice koji and its sweet children — grow the machine, then let it make sugar.

11 Chapter 10 · Miso, from the bean up — the long patience, and what it builds.

12 Chapter 11 · The fermented chili pastes — doubanjiang, gochujang, and a living chili crisp.

13 Chapter 12 · Living hot honey — the honest ferment, and its vegan twin.

14 Chapter 13 · Kimchi, laddered — one board, four heats, the same lift under each.

15 Chapter 14 · The fermented hot sauce — mash, brine, and the barrel that waits.

16 Chapter 15 · The ferments you drink — kombucha, tepache, and the chili session.

17 Chapter 16 · The three dials — salt, temperature, and time decide everything.

18 Chapter 17 · The alkaline ferments — natto, tempeh, and the bean made strange.

19 Chapter 18 · Vinegar — the other sour, and the heat it also lifts.

20 Chapter 19 · The scent of the work — the dimension the eight axes leave to the nose.

21 Chapter 20 · Reading the ferment — loudness climbs on a curve you can measure, and aim.

22 Chapter 21 · The safety floor and what the ferment gives back — the two numbers that keep it safe, and the gift underneath.

Opening essay · The crock and the construct

There is a sound a kitchen makes when something alive is working in it, and if you grew up in our house you learned it before you had a word for it — the faint tick of gas escaping a weighted lid, the sigh of a crock settling overnight, the particular hush of a room where millions of organisms are quietly taking a pile of beans apart and handing it back as something you could weep over. We kept ours on the cool end of the counter, under a cloth, and we were taught not to disturb it the way you are taught not to wake someone you love. You could smell the work happening: the sweet-grain breath of koji on the third day, the bright sour lift of cabbage turning to kimchi, the low brown depth of a miso that had been asleep for two winters. None of it we had added. All of it the culture had made, out of almost nothing, while we slept.

For as long as anyone has cooked, fermentation has been the one kitchen craft that felt less like cooking than like husbandry — you do not make the flavor, you keep the conditions and the living thing makes the flavor. Salt to choose who thrives. Time to let them work. A little warmth, a little air or none. Our grandparents could not have told you which enzyme was cutting which bond, could not have named the molecule that arrived on the tongue as *savory*. They knew only that if you were patient and clean and a little lucky, the crock gave back depth you could not get any other way, and that the giving was somehow also nourishing — that a spoonful of the old miso made a plain bowl of rice feel like being cared for.



The question that became this company came up at our table, over food the family had grown — long before any of it had a number.

We have spent the last decade learning to read that silence as chemistry, and the reading did something we did not expect: it pointed straight back at the work we already do. Scoville Splice designs flavor at the molecular level — roughly nineteen thousand five hundred product SKUs — backed by approximately 225 distinct construct designs — across forty-two host organisms, a thousand of them live and growing today, every one scored on the eight axes of the Iyengar Flavor Index. We tune a plant by changing how much of a particular molecule its cells make, because a flavor is loud in the mouth only in proportion to how fully it fills the receptor built to hear it. That is the whole trade. And when we finally understood what the crock on the counter was doing — proteins cut into free amino acids, starches cut into sugars, bland storage polymers unpacked into

small, insistent, receptor-active things — we realized fermentation had been running our business model for ten thousand years. It is not a different craft from ours. It is the same act, done by a different hand.

This is the thesis of this book, and we will spend every chapter earning it: **a flavor becomes louder in exactly one way — by raising the concentration of the molecule a receptor is tuned to — and there are two ways to raise it.** You can change the living plant, so its cells pack more of the molecule before harvest. Or you can change the harvested food, loosing a microbe's enzymes to liberate or synthesize the molecule after. Engineering works upstream, in the field. Fermentation works downstream, in the jar. They are not rivals. They are the same manufacturing line with the bench in two different rooms, and the extraordinary thing — the thing this whole book is here to show — is what happens when you run both rooms on the same batch. Engineer a tomato to carry more bound glutamate, then let koji's proteases set that glutamate free, and the ferment does not start from an ordinary tomato. It starts from a loaded one, and it ends somewhere neither technology could reach alone.

Then there is the title. *Fermentation and heat* is not two subjects we have bolted together for a craft book; it is one of the most elegant coincidences in the whole sensory world, and it sits at the exact center of what we do. The receptor that reports capsaicin — TRPV1, the molecular thermometer of the tongue — is not listening for capsaicin alone. It also listens for protons. Acidify a chili, and you lower the threshold of the very channel heat speaks through, so the same capsaicin reads hotter. Fermentation is, among other things, a controlled fall in pH. Which means that when a lactic culture turns a mash of our peppers sour over two weeks, it is not only preserving them and deepening them — it is quietly turning up the heat, at the receptor, for free. The oldest preservation technology on earth happens to be a volume knob for the one axis we are most known for. We did not arrange that. We only noticed it, and then we built around it.

A word, before the mechanisms begin, about why any of this matters beyond the pleasure — because the pleasure would be reason enough, and we will not pretend otherwise. We come to food as a family that believes you should give your body the most a plant can carry: the most flavor, and underneath it, quietly, the most nourishment. Fermentation is where those two intentions stop being separate. A culture that frees glutamate for your tongue is, in the same motion, cutting proteins into peptides your gut can actually take up, making minerals more available, laying down vitamins that were not there in the raw bean. The living work that makes the food delicious is the same living work that makes it more fully food. We will not dress that up as medicine; there are no cures in this book and no promises about your health. We will only follow the chemistry honestly, and the chemistry keeps arriving at the same place our grandparents arrived at by patience alone — that the crock gives back more than you put in, and that *more* is worth designing for.

A word on how the book is built, because it climbs deliberately. We begin with the convergence itself — the three pieces of receptor science that make fermentation and engineering one act: how a microbe manufactures the ligand a receptor is waiting for, how acid lifts heat at the channel, and how a nucleotide multiplies savor. Then we meet the organisms we have rebuilt for the crock — the koji and the bean, the pepper and the kimchi board, the quiet fungi that carry the multiplier. Then the ferments themselves, made and measured: rice koji and its sweet children, miso aged to the bottom of savor, the chili pastes and a living hot honey, kimchi laddered across four heats, the mash hot sauce aged like wine, the ferments you drink, and the alkaline ones that run the other way. And we end in the rigor the living work demands — the three dials that govern it, the way to read it as it climbs, and the two numbers that keep it safe — closing where the ethos has waited all along: on what a ferment gives the body back. Every chapter holds to the one rhythm the voice was built on: the reverent image, then the mechanism exactly, then what it means when you eat it.

So this is a book of two hands on one batch. It is a science book: we will name the enzymes, draw the pathways, write out the receptor kinetics and the synergy equations and the pH arithmetic, because the numbers are a form of respect for you and for the organisms doing the work. It is also, unapologetically, a cookbook: there are real ferments in here — koji grown from grain, miso built from the bean up, three fermented chili pastes, a living hot honey, kimchi laddered across four heat bands, a mash hot sauce aged like wine — each one a recipe you can make and each one read back to you on the eight axes so you can watch the loudness climb. Plants lead, as they always do in our kitchen; where meat and seafood enter a ferment they are written with the same care and no apology. And through all of it runs the single idea the crock taught us before we had the mathematics to see it: that to design a flavor and to design a food worth eating are, in the end, the same patient act.



PART I OF IV

The Convergence

How a microbe and a construct do the same thing — manufacture ligand · TRPV1 · mGluR1 · T1R1·T1R3 · the synergy

Chapter 1 · The ligand factory — *fermentation is manufacturing, not magic.*

Lift the lid on a week-old koji and you are looking at a factory floor. The grains of rice have gone soft and faintly sweet and are bound in a white felt of mycelium, and if you taste one it tastes of something that was not in the rice — a round, filling savor, a whisper of sweetness, the first hint of the depth that a month from now will be miso. Nothing was added. The rice is the same rice. What changed is that a mold, *Aspergillus oryzae*, has spent six days secreting enzymes into it, and those enzymes have been cutting the rice's large, silent molecules into small, talkative ones. That is the entire mechanism of fermentation, stated plainly, and it is worth stating plainly because the mysticism around ferments obscures how literal the process is: **a ferment is a controlled demolition of flavorless polymers into receptor-active monomers.**

Every flavor your tongue can report is a small molecule binding a receptor. Sweetness is sugar in the pocket of TAS1R2·TAS1R3. Savory depth is glutamate in the Venus flytrap of mGluR1 and T1R1·T1R3. Heat is capsaicin in the vanilloid cleft of TRPV1. But plants do not store their wealth as free small molecules — it would leak, oxidize, and cost them osmotically. They store it packed: protein instead of free amino acids, starch instead of glucose, bound and glycosylated precursors instead of active aglycones. A raw soybean is extraordinarily rich and almost silent, because its glutamate is sewn into protein chains where no receptor can reach it. The bean is a locked warehouse. Fermentation is the locksmith.



Koji — *Aspergillus oryzae* grown out on grain. Six days of secreted enzyme turns a silent starch into a factory for sugar, peptide, and free glutamate.

Consider the enzyme arsenal koji brings, because it is the model for the whole craft. It secretes **α-amylase** and **glucoamylase**, which together unzip starch into maltose and then glucose — this is the sweetness that blooms in amazake and the fuel that later yeasts and bacteria will eat. It secretes a battery of **proteases** (acid, neutral,

and alkaline) that chop storage protein into peptides, and then **peptidases** that trim those peptides down to free amino acids — this is the savory depth, and the kokumi body, and much of the aroma. And it secretes **glutaminase**, a small quiet enzyme that does one magnificent thing: it converts glutamine, which is nearly tasteless, into glutamate, which is the loudest umami molecule we have. A koji rich in glutaminase makes a miso that tastes savory out of all proportion to its protein, because it is not merely releasing the glutamate that was there — it is manufacturing new glutamate from a tasteless precursor. Hold that in mind; we will come back to it, because it is precisely the kind of move we like to engineer.

Now put a number on it, because the number is where fermentation and our own work become the same sentence. The Iyengar Flavor Index scores a pure compound by its intrinsic loudness:

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

where EC_{50} is the concentration that drives a receptor to half its maximum response. A low EC_{50} means a potent molecule — little is needed to be heard — and inverting it against a fixed reference turns potency into a score. But a food is not a pure compound, and fermentation does not change any molecule's intrinsic EC_{50}; glutamate's affinity for its receptor is a constant of chemistry, the same in a raw bean and a three-year miso. What fermentation changes is **how much free glutamate is present to do the binding.**

So the quantity that actually sets how loud an axis sounds in your mouth is not the bare IFI of the molecule but the receptor's *occupancy* — the fraction of receptors filled:

$$\theta = [L] \,/\, ([L] + \mathrm{EC}_{50})$$

and the number we print on a scorecard, the **effective IFI** of a food on an axis, tracks that occupancy, renormalized onto the same thousand-anchored scale. Fermentation raises [L]. Occupancy climbs toward one. The effective IFI climbs with it. That is the whole story of why an aged ferment tastes louder than the raw ingredient, written in two equations.

How fast does [L] climb? Here the craft reveals itself as kinetics. A koji protease working on bean protein obeys the same Michaelis–Menten law as any enzyme:

$$v = V_{max} \cdot [S] \,/\, (K_m + [S])$$

In the first days, substrate [S] — intact protein — is abundant, far above K_m, so the enzyme runs near V_max and free amino nitrogen pours out almost linearly. As the accessible protein is consumed, [S] falls below K_m, the rate drops, and release decelerates. Integrate that and you get the shape every fermenter knows in their hands even if they have never drawn it: a rapid early rise that bends over into a long slow approach to a ceiling. Model the free-glutamate accumulation as first-order toward a maximum,

```
Glu = Glu_max · (1 − e^(−kt))
```

and the ferment's whole personality lives in two constants. **Glu_max** is how savory it can ever become — set by the substrate's protein and by how much glutamine the glutaminase can convert. **k** is how fast it gets there — set by temperature, enzyme load, and salt. The half-time to reach the midpoint is ln2 / k. A warm, high-koji sweet miso has a large k and reaches its (modest) ceiling in two weeks; a cold, salt-heavy red miso has a small k and a high ceiling, and takes two years to climb it. The craft of fermentation is, in the end, the deliberate choice of a curve.

FOOD	FREE GLUTAMATE (MG / 100 G)	EFFECTIVE IFI-U BAND
Raw soybean	~10	low
Fresh ripe tomato	150–250	low–mid
Young white miso (2–4 wk)	200–400	mid
Aged red miso (1–2 yr)	500–700	mid–high
Soy sauce (premium grade)	400–1,700	high
Fish sauce	~950	high
Kombu (dried kelp)	1,500–3,000	very high

The industry even grades by it without saying so: premium soy sauce is classified by *amino-acid nitrogen*, the dissolved free-amino-acid load the koji released, with the top grade requiring roughly 0.8 grams per 100 millilitres. That number is nothing but [L] on the umami axis, measured by titration instead of by tongue. Every soy sauce on a shelf is already wearing an effective-IFI badge; it just prints it in different units.

Put real numbers through the model and the craft's own timelines fall out of it. Say a miso's accessible protein sets a glutamate ceiling Glu_max of 600 milligrams per 100 grams, and the salt holds the rate so the half-time is seven days — a rate constant k of about 0.099 per day. The model then says the ferment reaches roughly 300 mg at one week, 450 at two, 525 at three, and crawls the last stretch toward 600 over months, each week adding less than the week before. That decelerating approach is why a young miso tastes transformed in a fortnight and yet a two-year miso still reads deeper than a one-year — the climb never quite stops, it only slows, and patience is paid in an ever-smaller but never-zero increment. The craft's instinct that *more time always helps, but less each time* is the integral of a first-order rate law, felt on the tongue.

Release, though, is only half the story; the other half is a census of enzymes and a distinction between freeing and making. A single koji secretes not one protease but a battery — acid, neutral, and alkaline proteases with different pH optima, backed by the aminopeptidases and carboxypeptidases that trim peptides from their ends — alongside the amylases, the glutaminase, the lipases, and minor players like pectinase and tannase. Food chemists track the sum of this work as the *degree of hydrolysis*, the fraction of peptide bonds cut, and as the ratio of amino nitrogen to total nitrogen, which climbs as a ferment matures. A high degree of hydrolysis means more free

amino acids, more glutamate, and fewer of the bitter peptides of Chapter 4 — a higher U and a higher B at once. And the glutaminase step bears repeating, because it is the difference between a good ferment and an uncanny one: most of a bean's nitrogen arrives as glutamine, which is nearly tasteless, so a ferment that merely *released* its amino acids would leave much of its potential savor locked in a silent precursor. Glutaminase converts it. The loudest ferments are not simply the most hydrolyzed; they are the ones that also turn their glutamine into glutamate before it ever reaches your tongue.

Here is where our two rooms meet on one batch. Fermentation can only liberate what the substrate holds — Glu_max is capped by the raw material. So we raise the cap. Our soybean host carries three hundred and sixty constructs, and among them are edits that increase total storage protein and, more pointedly, shift the plant's nitrogen toward bound glutamine — the exact tasteless precursor koji's glutaminase is waiting to convert. And we engineer the other hand too: the koji itself is a host organism in our library, eighty constructs deep, and the edits that matter most there push harder secretion of glutaminase and the trimming peptidases. One edit raises Glu_max; the other raises k. Put the engineered bean under the engineered mold and you have moved both constants of the same curve — a ferment that climbs faster to a higher ceiling, every batch, at a coefficient of variation we hold at or under seven percent even though the thing doing the work is alive. The crock has always been a factory. We have simply learned to specify the raw material and tune the machines.

Chapter 2 · The proton and the pepper — *why a sour heat is a louder heat.*

Taste a fresh habanero purée and then taste the same chili fermented three weeks into a sour mash, and the fermented one is hotter. Not deeper only, not merely more complex — hotter, more insistent, the burn arriving faster and sitting higher. For a long time the craft explained this with the vocabulary of brightness and balance, and the explanation was not wrong, only incomplete. The complete explanation lives in a single receptor and a single ion, and it is the reason this book carries the title it does.



A pepper in cross-section — placenta, seed, oil. Capsaicin lives in the pith; the receptor that reports it also reports acid.

TRPV1 is a non-selective cation channel on the trigeminal nerve endings that wrap the tongue, and it is *polymodal* — it opens to more than one kind of signal. It opens to heat above about 43 °C, which is why a hot drink and a hot pepper share a word. It opens to capsaicin, which slides into a hydrophobic vanilloid pocket on the channel's inner face and holds it open at body temperature, so the nerve reports a fire that is not there. And — this is the hinge of the chapter — it opens to **protons.** Extracellular acid, a surplus of H^+, acts at residues on the channel's outer vestibule to lower its activation threshold. Protons alone, at high enough concentration, can crack TRPV1 open; protons at more modest concentration do something subtler and more useful to a fermenter — they *sensitize* the channel, so that a given dose of capsaicin produces a larger current than it would at neutral pH.

The receptor's architecture is why the coincidence works, so it repays a closer look. TRPV1 is a homotetramer — four identical subunits ringed around a central ion pore, each subunit contributing six membrane-spanning helices and a re-entrant pore loop — and it holds its several triggers at several places. Capsaicin slips into a hydrophobic vanilloid pocket on the channel's inner face, anchored by residues the literature numbers around Y511 and T550, with a half-maximal activation near a single micromolar. Heat acts through

the channel's temperature-sensing elements. And protons act at acidic residues on the *outer* vestibule — glutamates in the neighborhood of E600 and E648 in the rat channel — which is why acid and capsaicin can push the same gate from opposite faces of the membrane. Because the vanilloid pocket and the proton sites are physically distinct, their effects are not rivalrous but additive: the acid does not displace capsaicin, it lowers the energy the gate needs to open, so the same capsaicin occupancy swings more channels. Co-agonism, in a word — and the fermenter's falling pH is the second agonist, applied for free by the culture.

This is also where the old Scoville scale and our Index finally reconcile, and the reconciliation is the sharpest thing in the chapter. A Scoville rating measures *content* — historically a dilution assay, now chromatography, of how much capsaicinoid a fruit holds — and content is fixed at harvest; nothing a jar does adds a milligram of capsaicin. IFI-H measures *loudness* — the inverse of the effective EC_{50} at the receptor — and loudness is not fixed, because the receptor's sensitivity is not fixed. A chili's Scoville number cannot change in the ferment; its IFI-H can, and does, because the culture moves the receptor, not the molecule. That is precisely the distinction the Index was built to draw, and fermentation is its cleanest demonstration: same Scoville, higher IFI, because sour is a volume knob and heat was always a loudness and never merely an amount.

The consequence is exact and a little wonderful. Capsaicin's intrinsic EC_{50} at TRPV1 does not change when you add acid — the molecule is the same, the pocket is the same. But the channel's *responsiveness* climbs, which is equivalent to the capsaicin dose–response curve sliding leftward: less capsaicin needed for the same burn. In the Index's language, acidification lowers the **effective** EC_{50} of capsaicin in that tissue, and since IFI-H runs on the inverse of EC_{50}, the effective heat score rises with no change in capsaicinoid content whatsoever. A chili does not need more capsaicin to read hotter. It needs a lower pH. And lowering pH is the one thing fermentation does most reliably.

Watch it happen. Drop a mash of chilies into a 3-to-5-percent brine, keep oxygen off the surface, and the native lactic bacteria take over in a predictable succession. Early on the heterofermentative *Leuconostoc mesenteroides* leads, producing lactic acid alongside acetic acid, ethanol, and a prickle of CO_2. As the pH falls it yields to the hardier homofermentative *Lactobacillus* species, which drive almost purely to lactic acid. Over one to three weeks the mash moves from about pH 6.2 to pH 4.2 or below, and the curve of that fall is, read through TRPV1, a curve of rising heat.

FERMENT DAY	PH	DOMINANT CULTURE	EFFECTIVE IFI-H VS FRESH
0	6.2	(raw mash)	1.0×
2	5.4	*Leuconostoc*	~1.3×
5	4.8	mixed	~1.8×
10	4.3	*Lactobacillus*	~2.3×
21	3.8	*Lactobacillus*	~2.6×

Those multipliers are the honest heart of the title. Measured in isolation, TRPV1's response to a fixed capsaicin dose can climb severalfold as the bath moves from the neutrality of fresh tissue down toward pH 5 and below — the proton sensitization and the capsaicin binding are co-agonists, and their effects stack. Translate a two-to-threefold drop in effective EC_{50} into IFI-H and a fermented mash at pH 3.8 reads like a fresh chili of two to three times the Scoville, from the same fruit. This is why a well-made fermented hot sauce punches so far above the pepper it was built from, and why the great aged mashes — a Tabasco mash resting under salt in oak for up to three years — deliver a heat that is not just preserved but *concentrated at the receptor*, brightened and lifted by the acid the culture laid down.

There is arithmetic under the acidity worth having, because it governs both the heat and the safety. Lactic acid is a weak acid, pK_a 3.86, and only its undissociated form carries the antimicrobial punch and contributes fully to the free-proton pool that TRPV1 feels. The Henderson–Hasselbalch relation gives the split:

$$pH = pK_a + \log_{10}([A^-] / [HA])$$

At the fresh mash's pH 6.2, almost all the acid is dissociated and declawed; by pH 4.2, close to a third of the lactic acid sits in its protonated, active form, and by pH 3.8 nearly half does. The fermenter's instinct to push the pH down past 4 is, mechanistically, two instincts at once: safety, because the undissociated acid suppresses spoilage and pathogens, and heat, because the same protons sensitize the channel. One falling number serves both.

A brief, fizzing aside, because it belongs here. The prickle of a carbonated ferment — a lively kombucha, a young fermented soda — is not only bubbles on the skin of the tongue. Dissolved CO_2 is converted at the mucosa by carbonic anhydrase into carbonic acid, releasing protons locally, and those protons tingle the sour pathway and lightly provoke TRPA1, the sharp-pungency channel we will meet in the chili pastes. A spicy ferment that is also faintly effervescent is therefore ringing three bells at once — TRPV1 by capsaicin and acid, TRPA1 by its pungent thiols and the CO_2-born protons — which is exactly why a fermented, lightly carbonated hot sauce feels so much more *alive* in the mouth than a cooked one. The liveliness is literal. There is a living chemistry on your tongue.

Now our hand enters again. We breed peppers for this specific destiny. The fermentation-grade cultivars in our pepper host — the family is a thousand and sixty constructs wide — are selected and edited not for raw peak Scoville but for a mash that ferments cleanly

and lands its heat where the acid can lift it: thinner pericarp so brine penetrates fast, lower surface wild-yeast load, and a capsaicinoid profile skewed toward the shorter, faster-reading dihydrocapsaicin so the lifted heat arrives bright rather than lingering. And because TRPV1 is one receptor hearing two co-agonists, we can design the burn and let the culture tune it — build the pepper to a chosen capsaicin load, then let the fall in pH do the last multiplication, the same way a winemaker builds sugar in the grape and lets the yeast decide the wine. The pepper is the instrument. The ferment is the second hand on the dial.

Chapter 3 · The multiplier — *umami is not addition.*

The single most important fact about savory depth is that it does not add. Put a glutamate on your tongue and you taste umami; put a 5′-ribonucleotide — inosinate, guanylate — beside it, and you do not taste the sum of two savory things, you taste something several times larger than either. Dashi is built on this and nothing else: a sheet of kombu for glutamate, a handful of katsuobushi for inosinate, water just off the boil, and a broth whose savor is wildly out of proportion to the little that went in. Our grandmothers called it depth. The receptor calls it allostery, and once you see the mechanism you can design for it on purpose — which is exactly what a co-ferment lets you do.



Dried shiitake — the plant kingdom's great store of guanylate, the nucleotide that multiplies glutamate's voice.

Umami is read by two receptor systems working in parallel: the taste-specific heterodimer **T1R1·T1R3**, and metabotropic glutamate receptors, chiefly **mGluR1** (and a taste-truncated mGluR4), which the Index anchors the U axis to. T1R1 carries a large clamshell called a Venus flytrap domain, and glutamate binds down in the hinge of that clamshell, nudging it toward its closed, active shape. The 5′-nucleotides do not compete for that site. They bind a *second* site, just outside the hinge, and by sitting there they stabilize the closed conformation — they are a wedge that keeps the flytrap shut once glutamate has shut it.

The specificity of that second site is worth pausing on, because it is exact and it explains the kitchen. The amplifier must be a 5′-ribonucleotide monophosphate — inosine, guanosine, or adenosine 5′-monophosphate — and the position of the phosphate matters: the corresponding 3′-nucleotide barely works, which is why the enzymes that cleave RNA specifically to *5′*-nucleotides are the ones a cook learns, without knowing it, to court. Among the three the ranking is consistent — guanylate amplifies more strongly than inosinate, and inosinate more than adenylate — so a few grams of dried shiitake, rich in guanylate, outperform a far larger weight of a weaker partner. Structurally the nucleotide binds near the lip of the T1R1 flytrap and clamps it around whatever glutamate has entered, which is why its effect is less to make the umami *bigger* than to make a little glutamate *behave like a great deal* — the signature of allostery, a small ligand changing the meaning of a large one. The metabotropic receptors the Index anchors the axis to, mGluR1 and its taste-truncated cousins, lend the same signal a slower, rounder, longer-lasting dimension, part of why umami lingers where sweetness flashes and is gone. The practical result is that in the presence of inosinate or guanylate, glutamate's affinity for its own site rises steeply: its effective EC_{50} drops, by as much as eight-fold. Lower EC_{50} is higher IFI. The nucleotide does not taste of much alone, and yet it multiplies the loudness of the glutamate beside it. This is the molecular machine under every great broth, every aged cheese, every long ferment, and it is the reason a pinch of dried mushroom can make a whole pot savory.

The multiplication has a famous equation. In the classic psychophysical form, the perceived umami intensity *y*, expressed as the equivalent concentration of glutamate alone, relates to the glutamate fraction *u* and the nucleotide fraction *v* of a mixture by

```
y = u + β · u · v
```

where β is a synergy constant on the order of 1,200 for inosinate and larger still for guanylate. The term that matters is the product $u \cdot v$ — the cross term, the one that is zero if either partner is missing and enormous when both are present. Set $u = v = 0.5$ and the linear part contributes 0.5 while the synergy term contributes $\beta \cdot 0.25 \approx 300$: the mixture tastes like hundreds of times its glutamate alone would. No amount of glutamate by itself climbs like that; the nonlinearity is the whole point, and it means the U axis is not a line you walk up by piling on more of one thing — it is a surface you climb fastest by holding two hands at once.

SOURCE	PRINCIPAL UMAMI MOLECULE	TYPICAL FREE CONTENT (MG / 100 G)
Kombu (dried kelp)	glutamate	1,500–3,000
Parmesan (aged)	glutamate	~1,200
Soy sauce	glutamate	400–1,700
Ripe tomato	glutamate	150–250
Katsuobushi (dried bonito)	inosinate (IMP)	470–700
Dried shiitake	guanylate (GMP)	~150
Dried sardine (niboshi)	inosinate (IMP)	300–500

Read that table as two columns of one equation. The glutamate sources set *u*; the nucleotide sources set *v*; a cook who combines a high-*u* ingredient with a high-*v* ingredient is solving for maximum *y* whether they know the symbol or not. Kombu and bonito. Tomato and anchovy. Parmesan and mushroom. Every cuisine on earth stumbled onto the same product term, because the tongue rewards it so steeply that you cannot help but find it.

Fermentation is extraordinary here because a single ferment can feed *both* columns. Koji proteolysis pours out glutamate, raising *u*. Yeast, when it autolyzes at the end of its life — the step behind every yeast extract and much of the depth in aged soy sauce and long miso — breaks its own RNA down into 5′-nucleotides, raising *v*. A ferment that runs koji and yeast in sequence is building *u* and *v* in the same jar, climbing the synergy surface rather than one of its edges. Add a few grams of dried shiitake to a simmering miso and you are hand-placing a spike of guanylate onto a field of koji-made glutamate, and the two-hundred-milligram mushroom does the work of a far larger quantity of anything singular.

And there is a third axis riding along, quiet and structural. **Kokumi** — the Index's K axis, the sensation of thickness, continuity, and mouth-filling richness — is read not by a taste receptor in the usual sense but by **CaSR**, the calcium-sensing receptor, which responds to certain **γ-glutamyl peptides**: glutathione (γ-Glu-Cys-Gly) and small peptides like γ-Glu-Val-Gly. These molecules have almost no taste of their own, yet in the presence of umami and salt they amplify the whole sensation and, crucially, extend it — they are why a good aged miso does not just taste savory but *keeps* tasting, filling the sides of the mouth and lingering. Fermentation generates γ-glutamyl peptides abundantly, through the γ-glutamyltransferase activity that koji and aging bring, which is why the longest-aged ferments feel the roundest and most continuous. One culture, running on one substrate, feeds three axes at once: glutamate for U, nucleotides for the synergy that multiplies U, and γ-glutamyl peptides for the K that gives the whole thing body and length.

This is where our library turns the coincidence into a design. We do not have to rely on whatever nucleotide a ferment happens to make; we can engineer the partners. The tomato host — two thousand and forty constructs, our most developed umami platform — is edited to hold unusually high free glutamate at full ripeness, a loaded *u* before fermentation even begins. The shiitake and button-mushroom hosts are tuned toward elevated guanylate, concentrated *v* in a form you can dry and crumble. Pair an engineered high-glutamate tomato, fermented under koji so still more glutamate is freed, with a dusting of engineered high-guanylate shiitake, and you are not hoping for synergy — you are specifying both terms of the product $u \cdot v$ and multiplying them on purpose. The grandmothers found the cross term by patience and a thousand pots. We can write it into the seed. The reverence is the same; only the hand is new.

Two technologies, one act — manufacture the molecule the receptor waits for. The rest of the book is that act, performed.



PART II OF IV

The Organisms of the Crock

The hosts we engineer for the jar — koji · soybean · pepper · the kimchi board · the nucleotide fungi

Chapter 4 · Koji, the machine we rebuilt — *Aspergillus oryzae, eighty constructs deep.*

Of the forty-two organisms in our library, one is not a plant and not an animal and does more of the actual work of flavor than almost anything we grow. *Aspergillus oryzae* — koji — is a filamentous mold, domesticated over a thousand years from a wild ancestor the way a wolf was domesticated into a dog, and like the dog it has been bred so thoroughly toward usefulness that it can no longer quite live wild. Its toxin-producing machinery, the aflatoxin gene cluster that makes its wild relatives dangerous, has been silenced through the long centuries of selection; the koji we use is a non-toxigenic strain, one of the most intensively studied and safest industrial organisms on earth. It is also, functionally, an enzyme-secreting factory you can grow on a tray of rice in two days. We treat it as exactly that: a host to be tuned, eighty constructs deep, the same as any pepper or tomato in the catalog.



Koji under the lens — the hyphal felt that will cut a bland grain into sugar, peptide, and free glutamate.

Start with the craft as it exists, because our edits only make sense against it. You steam a grain — rice for the sweet, clean kojis; barley for a nuttier one; the soybean itself for certain misos — cool it to blood warmth, and inoculate it with *tane-koji*, the spore starter. Then for forty-two to forty-eight hours you keep it in a warm, humid chamber and you *manage* it, because the mold's own metabolism throws off heat and it will cook itself if you let it. The traditional hands — the first mixing, the middle turning, the final shaping into furrows that shed heat — are a thermal-control program done by

feel. The target is a grain bound in white, faintly green-free mycelium that smells of chestnut and flowers; let it go too far and it sporulates, greening and turning bitter, the enzymes giving way to the next stage of a life you did not want it to reach. You are harvesting an organism at the precise hour its enzyme secretion peaks and before it spends that wealth on reproduction.

And here is the lever the old masters found without the molecular picture: the *ratio* of enzymes koji secretes depends on how you run the temperature. Cooler incubation, around 30 °C, favors the amylases — more starch cut to sugar, a sweeter koji for amazake and sweet white miso. Warmer incubation, pushing 35 °C and above, favors the proteases — more protein cut to peptide and amino acid, a savorier koji for a long red miso or a soy sauce. The koji maker tunes the flavor destiny of the batch by choosing a thermal path through the incubation, biasing the enzyme spectrum before a single bean is added. It is flux control by thermostat, and it works because the two enzyme families have different optima and different regulation. Everything we do to the strain is in service of the same ratios, made sharper and more reliable.

What we edit, and why. The first and most valuable set of constructs pushes **glutaminase** — the enzyme that turns tasteless glutamine into loud glutamate. Recall from Chapter 1 that a ferment's umami ceiling, Glu_max, is set partly by how much of the substrate's nitrogen can be *converted* to free glutamate, not merely released; a koji that over-secretes glutaminase raises that ceiling directly. The second set tunes the **protease-to-peptidase balance**, and this is subtler and arguably more important, because it governs bitterness. Proteases cut protein into peptides; if the ferment stops there, many of those peptides are small, hydrophobic, and *bitter* — proline-, leucine-, and phenylalanine-rich fragments that sit in the bitter receptor TAS2R38. It is the exopeptidases, the carboxypeptidases and aminopeptidases that trim those bitter peptides down to free amino acids, that turn a harsh young ferment round and savory. A koji edited for strong exopeptidase activity therefore does two things on two different axes at once: it raises IFI-U, by making more free glutamate, and it raises IFI-B — our bitterness-attenuation axis — by destroying the bitter peptides before they can be tasted. One strain, two axes, the same cut.

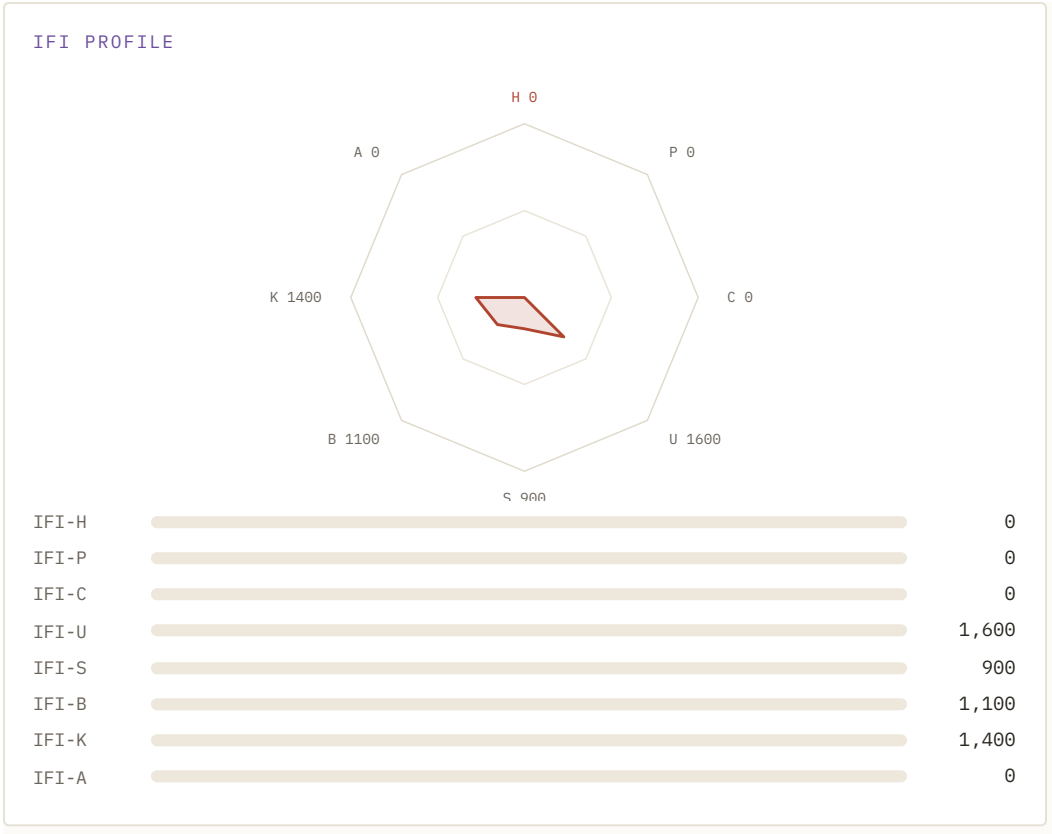

That is a finished rice koji read on the eight axes — nothing hot, nothing pungent, but a quiet, filling spread of umami, sweetness, body, and a measurable attenuation of bitterness, all of it manufactured in two days from a grain that scored near zero on everything. A third construct set tunes **lipase**, because fats cut cleanly give fruity, desirable aroma esters while fats cut carelessly give rancidity; a fourth suppresses pigment and spore formation, keeping a batch in its enzyme-rich prime a little longer and cleaner. And a fifth, less glamorous but the reason we can sell the result, hardens enzyme secretion against temperature swings, so a tray run a degree warm does not drift — the tool by which we hold a living organism's output to a coefficient of variation at or under seven percent.

Koji is not one thing but a family, and the grain you grow it on sets its character before any edit does. Grown on rice, it is the clean, sweet, amylase-forward koji of sake and white miso and amazake; grown on barley, it is nuttier and earthier, the base of the mugi misos; grown on the soybean itself, it is the dark, protease-heavy koji of the hatcho-style misos and the soy sauces, where there is little starch to sweeten and everything bends toward savor. The substrate sets the starting enzyme ratio — more starch invites more amylase, more protein more protease — and our strain edits ride on top of that choice rather than replacing it, so a grower picks a grain for its register and then a strain for its precision. The spore starter that carries all of this, the *tane-koji*, is a small craft of its own: a koji grown deliberately to green, sporulating maturity and sometimes dusted with wood ash, whose alkalinity the mold tolerates and its competitors do not — a four-hundred-year-old purity trick that, read in modern terms, is simple selective pressure by pH. We supply our engineered strains as defined spore preparations at a counted density, which is that same purity trick carried to its end: not hoping the right organism wins, but pitching a known one, in a known number, every time.

One last edit closes a loop with the heat side of the book. Koji secretes tannase and related activities that cleave the bound phenolics which can read as astringent, and a strain tuned there yields a rounder, less puckering ferment — a quiet touch on the A axis to go with its loud work on U, B, and K. A single organism, grown in two days, reaching across four of the eight axes at once: this is why we treat the mold not as a tool we reach for but as a host we design, the equal of any pepper in the catalog.

We design these edits the way we design everything: in silico before in vivo. The same structure-prediction and flux-balance pipeline that gave us the Pun1 double variant — a protein language model scoring every substitution for its effect on catalytic efficiency and fold stability, then a genome-scale metabolic model locating control in the pathway and validating each edit — runs just as well on a secreted glutaminase as on Pun1.

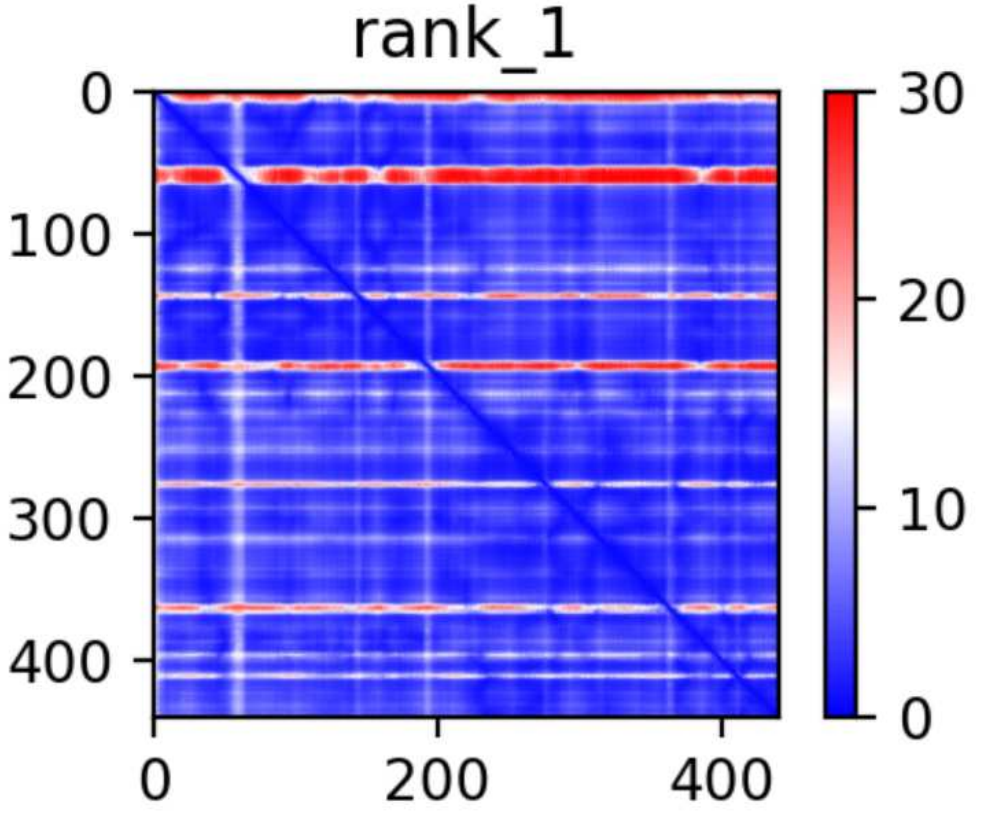

FIG. 4 — predicted aligned-error heatmap for a modeled secreted enzyme; the dark diagonal blocks are the confidently placed catalytic domains we edit against.

A word on the regulation, because an engineered food mold is exactly the kind of thing that should be held to account. *A. oryzae* is Generally Recognized As Safe in its traditional form, and the line we work from is non-aflatoxigenic; every construct we introduce is screened by whole-genome verification that the dormant aflatoxin cluster remains non-functional and that no horizontally acquired

element has woken it, and the strains are held under containment through development. We say this plainly because the nourishment ethos that runs under this book is not only about what a food gives but about trusting what it is. A culture you let into your food should be one whose entire genome you can read and vouch for. That is a standard the old masters could meet only by lineage and faith. We can meet it by sequence — and we do.

Chapter 5 · The bean that holds the depth — *soybean, three hundred and sixty constructs.*

Nearly every great savory ferment on earth runs through one seed. The soybean is a small, hard, almost flavorless legume that happens to be among the most protein-dense staples we grow — roughly thirty-six to forty percent protein by dry weight, a fifth oil, the rest carbohydrate and a scatter of remarkable minor compounds — and that protein is the warehouse every soy ferment exists to unlock. Miso, soy sauce, tempeh, natto: four of the pillars of the fermented world, and all of them are the same bean handed to a different organism. To understand the ferments you first have to understand the warehouse, and then you understand why we spent three hundred and sixty constructs rebuilding it.



The soybean — glycinin and conglycinin storage protein, isoflavone, oil. A warehouse of savor, locked until a microbe arrives with the key.

The bulk of the bean's protein is two storage proteins, glycinin (the 11S fraction) and β-conglycinin (the 7S fraction), folded tight and nearly tasteless. Its oil carries linoleic acid, which matters because the bean also carries **lipoxygenase**, an enzyme that, the moment the bean is crushed in water, oxidizes that linoleic acid into the green, grassy, "beany" aldehydes that untrained palates dislike in soy. It carries **isoflavones** — genistein and daidzein — mostly in bound, glycosylated form, where they are poorly absorbed. It carries **phytate**, which chelates minerals and carries them out of reach, and

the oligosaccharides raffinose and stachyose, which humans cannot digest but gut microbes can, to famous effect. The raw bean, in other words, is rich and difficult: a fortune in protein wrapped in antinutrients and an off-flavor enzyme. Fermentation's genius is that it solves all of these at once, and our engineering is aimed at making the bean easier for the microbe to solve.

Watch the four ferments diverge from the one bean. **Miso and shoyu** hand the bean to koji: the mold's proteases and glutaminase convert storage protein to a flood of free amino acids and glutamate over months under salt, and in soy sauce a second stage, a brine *moromi* fermented by the halophilic yeast *Zygosaccharomyces rouxii* and the lactic bacterium *Tetragenococcus halophilus*, layers in aroma esters — above all the caramel-sweet HEMF — and a bright acidity. **Tempeh** hands the bean to *Rhizopus oligosporus*, which binds the cooked beans into a sliceable cake, partially digests the protein, cuts the phytate with its phytase, and lays down B-vitamins; it is the gentlest of the four, mushroomy and mild. **Natto** hands the bean to *Bacillus subtilis*, an alkaline ferment that proteolyzes hard, throws off pungent pyrazines, synthesizes vitamin K2 as long-chain menaquinone, and spins the surface into sticky threads of poly-γ-glutamic acid — which is, note, glutamate polymerized into a chain, bound and therefore *not yet* savory, a reminder that it is free glutamate and only free glutamate the receptor can hear.

SOY FERMENT	ORGANISM	CHEMISTRY IT PERFORMS	LEAD AXES
Miso	*Aspergillus oryzae* (koji) + salt	proteolysis → glutamate; long Maillard	U · K · B
Soy sauce	koji, then *Z. rouxii* / *T. halophilus* brine	glutamate + esters (HEMF) + acid	U · K
Tempeh	*Rhizopus oligosporus*	partial proteolysis; phytase; B-vitamins	U (mild)
Natto	*Bacillus subtilis*	alkaline proteolysis; K2; poly-γ-glutamate	U · K (aroma)

Three hundred and sixty constructs, then, aimed at the warehouse. The first and most quietly important are the **lipoxygenase knockouts** — removing the enzyme that makes the beany off-note, so that a ferment built from the bean starts clean and lets the savory depth read without a green veil in front of it. The second shift the **storage-protein ratio and the glutamine load**: more of the nitrogen parked as bound glutamine, the precursor koji's glutaminase converts straight to glutamate, so that Glu_max — the ferment's umami ceiling from Chapter 1 — is raised in the field before the mold ever arrives. The third trims the **antinutrients**: lower phytate so the ferment liberates more available mineral, lower raffinose and stachyose. None of these are flavors we add; they are obstacles we clear and precursors we pre-position, so that the microbe's own enzymes arrive to an easier, richer substrate.

It is worth dwelling on those antinutrients, because clearing them is where the nourishment ethos and the flavor work turn out to be a single edit seen from two sides. The raw soybean defends itself

chemically: trypsin inhibitors of the Kunitz and Bowman–Birk families block the very protease the gut would use to digest it, phytate chelates iron and zinc out of reach, and the oligosaccharides raffinose and stachyose pass undigested to the colon with the famous consequences. Cooking denatures the trypsin inhibitors; fermentation handles the rest, as microbial phytase cleaves the phytate and the culture eats the oligosaccharides before you can. A fermented soy food is therefore not only louder on the tongue but more available to the body — more of its protein digestible, more of its minerals unbound — and the edits that lower phytate and trim the oligosaccharides in the living plant make the microbe's cleanup lighter and the result cleaner still. We do not add nutrition. We remove the locks the bean put on its own.

The bean's amino-acid profile is the other reason it anchors the whole savory world. Soy protein is unusually rich in glutamate and glutamine among its residues, which is to say the warehouse is stocked disproportionately with exactly the precursor the U axis is listening for; a protease set loose on soy liberates a flood weighted toward savor rather than toward blandness. Tilt that profile a little further toward glutamine, as our edits do, and you have biased the ferment's whole output toward umami at the source — a bean that was always going to make a savory miso, rebuilt to make a louder one.

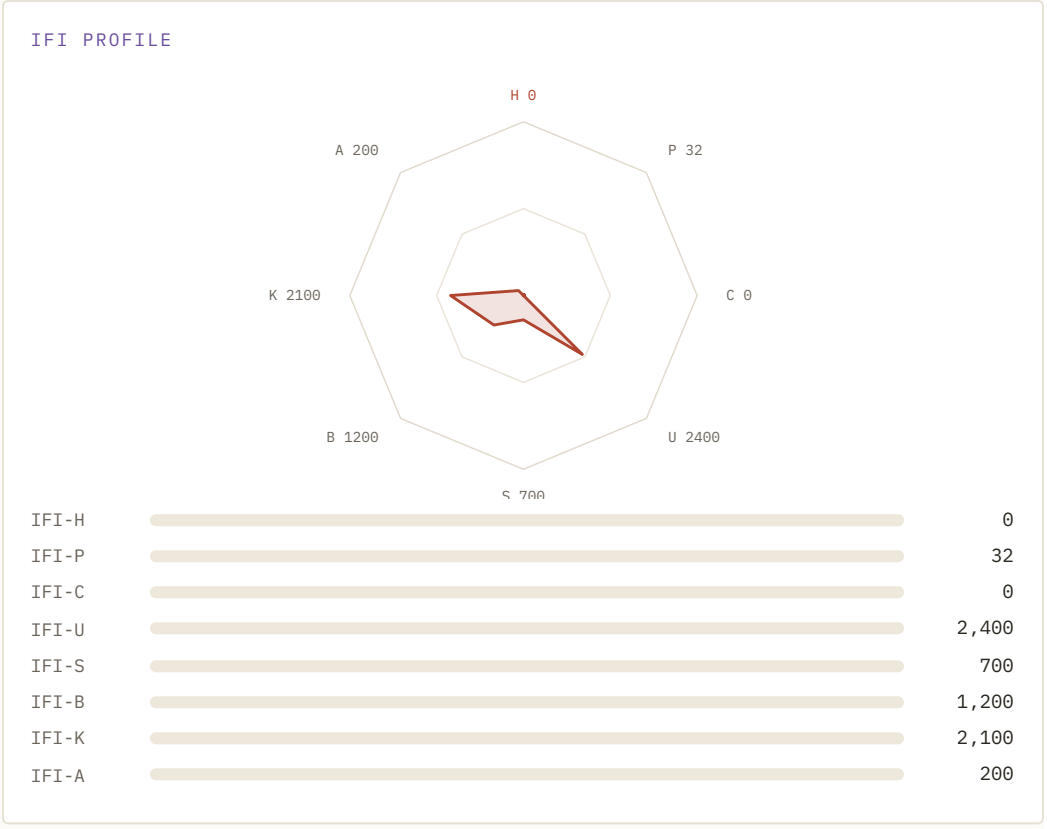

That is a two-year red miso built from the engineered bean under the engineered koji — a reading that would have been impossible from either hand alone. The umami is high because the bean carried more glutamine and the koji converted more of it; the kokumi is high because the long aging generated γ-glutamyl peptides against a clean background; the bitterness attenuation is real because the koji's exopeptidases trimmed the harsh peptides and the Maillard browning of two winters rounded what remained. And underneath the flavor sits the nourishment the ferment gives back and we will

not oversell: protein already cut to peptides and amino acids the body takes up more readily, isoflavones freed from their sugars into their absorbable aglycone form, minerals released from phytate's grip, vitamins laid down that the raw bean never held. We did not engineer those benefits in. The ferment makes them, as it always has. We only handed it a better bean to make them from.

Chapter 6 · Peppers bred for the crock — *a thousand and sixty constructs, a different spec for the jar.*

A pepper destined for a fermentation is not the same pepper you would breed to eat fresh, and it is emphatically not the same one you would breed for extract. We keep a thousand and sixty constructs in the pepper host — it is our native organism, the one the whole company grew from — and a distinct branch of that work exists for one purpose: fruit that ferments cleanly and lands its heat exactly where the acid can lift it. To see why that is a separate design problem, you have to follow capsaicin from its biosynthesis to its moment on the tongue, and then remember what Chapter 2 did to that moment.



A fermentation-grade pepper in cross-section — thin pericarp for fast brine penetration, capsaicin concentrated in the placental pith.

Capsaicin is built where two metabolic rivers meet in the placenta, the pale pith that holds the seeds. One river is the phenylpropanoid branch that runs in every vascular plant — **PAL → C4H → 4CL → HCT → COMT** — carried the last step by **pAMT** to vanillylamine, the aromatic half of the molecule. The other is a branched-chain fatty-acid pathway — **BCAT → KAS** building 8-methyl-6-nonenoic acid, the fatty tail. The enzyme **Pun1**, capsaicin synthase, condenses the two: vanillylamine plus the branched acyl group, amide bond formed, capsaicin made. Shift the fatty-acid side toward a shorter,

saturated tail and you make dihydrocapsaicin instead, and that distinction is not academic to a fermenter, because the two capsaicinoids have different *shapes* of burn. Dihydrocapsaicin reads sharp, forward, and quick; capsaicin reads slower, deeper, with a long tail. A bright fermented sauce wants the forward one, so our mash cultivars are skewed toward dihydrocapsaicin — a heat that the acid lift of fermentation brightens rather than smears.

The capsaicinoids are a family, not a molecule, and the family's proportions are themselves a design variable. A typical pepper's heat is roughly seventy percent capsaicin, a fifth dihydrocapsaicin, and the rest a scatter of nordihydrocapsaicin and minor relatives — and because each reads with a slightly different onset, duration, and location in the mouth, the *blend* matters as much as the total. Capsaicin is the long, deep, lingering one; dihydrocapsaicin the sharp, forward, fast one; the minors supply the edges. A fresh-eating cultivar can afford a capsaicin-heavy, slow-building profile, but a ferment, which will brighten and lift whatever it is handed, is better served by a blend tilted forward, so the lifted heat arrives clean rather than dulling to a smear under the acid.

All of this rides on one enzyme, capsaicin synthase — Pun1, also called AT3 — and the double variant designed into it is the heart of the heat platform. We found them not by breeding but by computation: a protein language model scored every possible single-residue substitution for its predicted effect on catalytic efficiency and fold stability, while our flux-balance model showed that metabolic control sits upstream in the phenylpropanoid pathway — at 4CL and HCT/CCoAOMT — rather than at the terminal Pun1 step. Pun1 is the final capsaicin synthase; it labels the product, but our model predicts it does not govern the rate. Two substitutions worth combining at the Pun1 active site nonetheless emerged from the structural scan, and written in with gene editing they lift per-molecule output while collapsing the batch-to-batch variance.

EDIT	EFFECT	ΔΔG (KCAL/MOL)
S39L	binding affinity ↑ (N-terminus)	−2.1
L345G	steric relief, active-site cleft	−1.8
C175S	disulfide-misfold prevention (structural annotation — not in shipped construct)	−0.9

The double variant S39L+L345G is why a flagship mash begins from a capsaicin distribution tight enough that the ferment's multiplier lands where we said it would — bottle after bottle, and not as a lottery.

Then there is everything about the fruit that is not capsaicin. A mash ferments through its surface and its flesh, so **pericarp thickness** matters — thinner walls let a brine penetrate fast and evenly, so the culture establishes before any spoiler can. The **surface microbiome** matters — a fruit bred and grown for a low wild-yeast and mold load gives the lactic bacteria a cleaner start. And **sugar** matters more than anyone expects, because the superhot *C. chinense* fruits that carry the great heat are low in fermentable sugar, and lactic bacteria need sugar to make the acid that is the whole point. A pure superhot mash can stall for lack of fuel. The traditional fix is to co-ferment

with a sweeter, milder pepper or a little carrot; our additional fix is to edit a modest lift in fruit fructose into the mash cultivars themselves, enough to feed the culture to a full, safe pH drop without blunting the heat.

FLAGSHIP CULTIVAR	SKU	PREDICTED SHU[1]	IFI-H (FRESH)	AXES
Ember	SS-0087	866K	1,570	2/8
Cinder	SS-0088	1.05M	2,290	3/8
Blaze	SS-0091	1.72M	2,630	4/8
Wildfire	SS-0092	1.96M	2,930	5/8
Firestorm	SS-0094	2.47M	3,490	6/8
Solara	SS-0095	2.72M	3,980	6/8
Phoenix	SS-0096	2.95M	4,200	8/8

[1] SHU values are computational predictions from our flux model, not certified measurements.

Those are the ten proof-of-concept superhots as they grow, before any jar touches them — and the reason consistency is non-negotiable for a fermenter is that the acid multiplier of Chapter 2 *compounds variance*. If a fresh mash varies ten percent in capsaicin and the ferment then multiplies the effective heat two-and-a-half-fold, the finished sauce inherits both the spread and the multiplication; sloppiness upstream becomes chaos downstream. This is why we lock the heat in the plant first. The double variant — **S39L** lifting binding affinity at the N-terminus, **L345G** relieving steric strain in the active-site cleft — raises peak output and, more to the point here, pulls the batch-to-batch coefficient of variation from a wild-type forty-two percent down to seven or below. The ferment then starts from a tight distribution, and the sauce it makes is reproducible heat rather than a lottery.

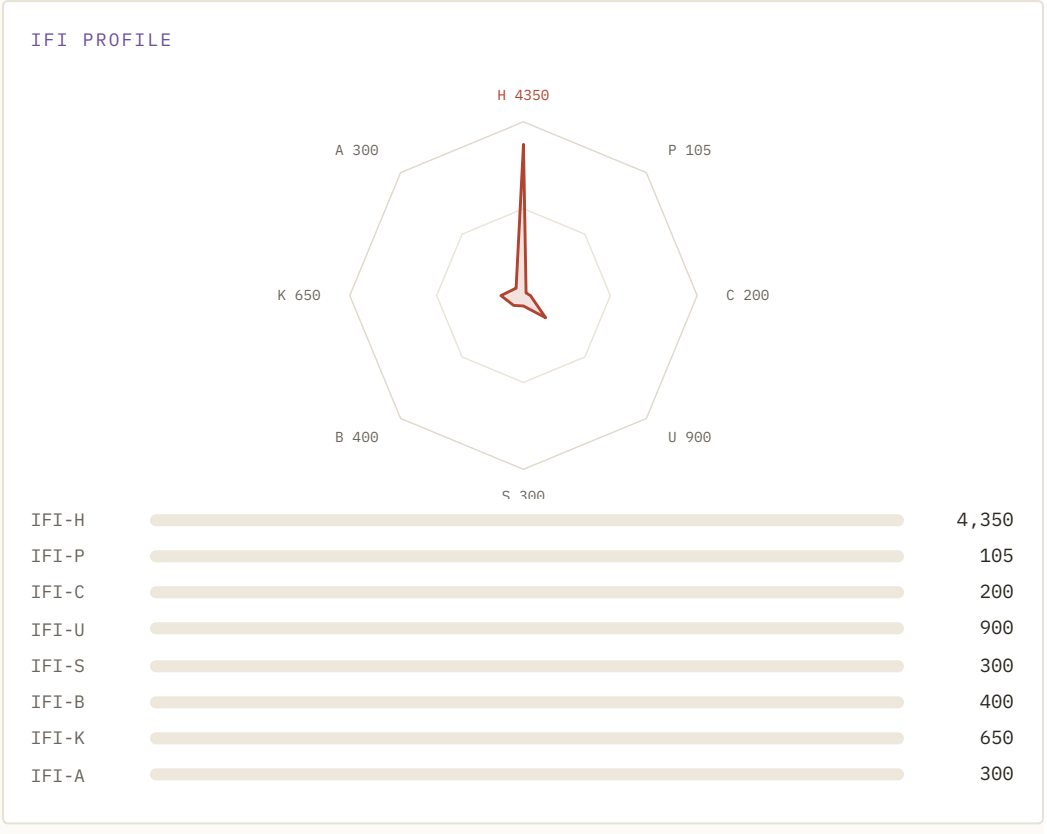

That is a finished mash sauce built on Wildfire — fresh IFI-H 2,930, grown at a locked coefficient of variation — read after three weeks of lactic fermentation to pH 3.8. The heat is up by roughly half again on the effective scale, pushed there not by more capsaicin but by the proton sensitization of TRPV1; the small but real umami and kokumi are the ferment's gift, glutamate and γ-glutamyl peptides the

lactic culture and the mash's own proteolysis laid down; the whole thing reads brighter and rounder than the fresh fruit it came from. One pepper, built to a number in the field, finished to a louder number in the jar. The instrument and the second hand on the dial, exactly as promised.

Those are the loud organisms. The next two chapters are the board and the multiplier — the vegetables that ferment into a living side dish, and the quiet fungi that make everything savory climb.

Chapter 7 · The kimchi board — *cabbage, radish, allium, ginger, and the culture that eats them.*

Kimchi is not a recipe, it is a *consortium* — a board of vegetables brought together so that a procession of bacteria can live on them in turn, each organism setting the table for the next, the whole thing moving over days from fresh and vegetal to deep, sour, carbonated, and alive. No one inoculates a traditional kimchi; the culture rides in on the vegetables themselves and on the hands that make it, and the fermenter's only job is to set conditions — salt and temperature — that let the right organisms win. Every vegetable on the board is doing a specific job, chemically, and because we engineer most of them, we get to design the board one role at a time.



Napa cabbage — the matrix of the board. Salt draws its water, selects the culture, and opens the leaf to everything that follows.

The matrix is **napa cabbage**, and the first act is salt. Brining the cut cabbage to roughly two-and-a-half to three percent does three things at once: it draws water out of the leaf by osmosis, wilting it pliable and concentrating everything; it lowers water activity enough to exclude most spoilage organisms; and it specifically selects for the salt-tolerant lactic bacteria that will run the ferment. Here our first edit matters, because cabbage is a *Brassica*, and brassicas carry glucosinolates — sinigrin and its relatives — which an enzyme called myrosinase converts, the moment the tissue is cut, into

isothiocyanates: sharp, pungent, and for many people bitter, sitting squarely in the bitter receptor TAS2R38. A little of that is character; too much is harsh. The cabbage host carries forty constructs on the bitterness-attenuation axis, and the lead edits lower the glucosinolate load and shift its profile, pulling the finished kimchi's bitterness down by the twenty to sixty-five percent that defines our IFI-B range — so the sour, savory, spicy flavors read clean instead of fighting a bitter undertone.

Radish — daikon, another forty-construct brassica host — brings crunch that survives the ferment and its own measured dose of pungent isothiocyanate, engineered to a controlled sharpness rather than the raw bite of a wild root. Then the alliums. **Garlic** is a five-hundred-and-sixty-construct host, and it earns its place twice over. Crushed, its odorless alliin meets the enzyme alliinase and becomes allicin, a reactive sulfur compound that lights up the pungency channel TRPA1 — the same sharp, nose-reaching sensation as mustard and wasabi. But garlic also carries γ-glutamyl-S-allyl-cysteine and related peptides that speak to CaSR, the kokumi receptor, giving richness and length; our garlic is edited kokumi-forward, so that beyond the pungent hit it deepens and sustains everything around it. **Ginger**, a hundred-and-twenty-construct host, adds gingerol, which warms through both TRPA1 and a light touch of TRPV1 — a rounder, warmer pungency than garlic's sharp one. And over all of it goes the chili, the gochugaru, carrying the heat that this book exists to ladder; we will do exactly that, across four bands, in a recipe chapter to come.

Each of those pungencies is a specific reaction a cook triggers by cutting, and the ferment then tames and transforms. The brassicas — cabbage and radish — hold glucosinolates like sinigrin intact until the tissue is damaged; then myrosinase, released from a separate cellular compartment, hydrolyzes them into isothiocyanates, the volatile, sharp-bitter compounds that give raw radish and cabbage their bite. Salt and the ferment moderate this: some isothiocyanate volatilizes, some is transformed, and the lactic acids layer over what remains, so a mature kimchi's brassica pungency reads rounded rather than raw. Garlic runs the parallel path — odorless alliin, once cut, meets alliinase and becomes allicin, which is unstable and decays into the diallyl sulfides that give long-fermented garlic its mellow, sweet, sulfurous depth; it is the same chemistry, slowed and sweetened, that turns a clove into black garlic over weeks of warm aging. Ginger's gingerol, warming through TRPA1, converts on heating and drying into the sharper shogaol, which is why dried ginger bites harder than fresh. A fermenter who knows these reactions can place each pungency exactly — raw and loud, or aged and round — by when and how the tissue is cut and how long the culture is given to work on what the cut released.

This is also where the bitterness-attenuation edits earn their place on the board. Left alone, a brassica-heavy kimchi can carry a harsh glucosinolate-derived bitterness that fights the sour and the savory; the low-bitter cabbage and radish lower that load at the source, so

the ferment's brightness comes through clean. It is the same move as the koji's exopeptidases in Chapter 4 — not adding a flavor, but removing the one that stood in front of all the others.

Now the culture takes over, and the succession is beautiful. In the first days, while oxygen lingers and the pH is still near six, the heterofermentative *Leuconostoc mesenteroides* and *Leuconostoc citreum* lead — they make lactic acid but also acetic acid, ethanol, mannitol, and carbon dioxide, which is the gentle fizz and the bright, slightly sweet lift of a young kimchi. Mannitol itself tastes faintly sweet and cooling. As their own acid drops the pH below about 4.5, they yield to the hardier homofermentative and facultative species — *Latilactobacillus sakei*, *Lactiplantibacillus plantarum*, *Weissella* — which drive almost purely to lactic acid and carry the ferment down to its mature pH of around 4.2, sour and stable. Watch what that falling pH does to the chili on the board: it is the heat multiplier of Chapter 2, operating inside a side dish. The gochugaru you added does not taste hotter because you added more; it tastes hotter because the *Leuconostoc* and the *Lactobacillus* spent a week acidifying the receptor it speaks to.

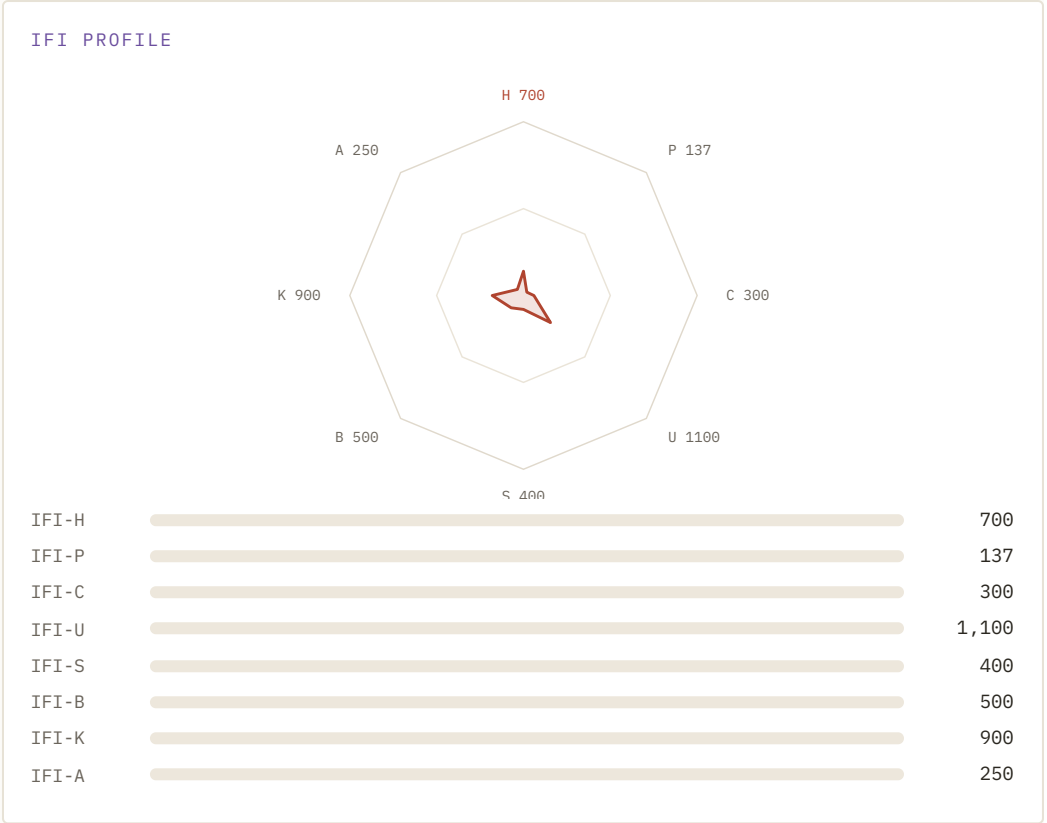

That is a medium-heat cabbage kimchi at full maturity — a genuinely eight-dimensional food made from vegetables that each scored near nothing raw. The heat is the gochugaru lifted by the ferment's acid; the pungency is the garlic and ginger speaking to TRPA1; the umami is the lactic bacteria's own glutamate plus whatever savory base you built in (kelp and shiitake for a plant-first board, salted shrimp or fish sauce for a traditional one, each written with equal care); the kokumi is the engineered garlic giving the whole thing body; and the faint cooling is the *Leuconostoc*'s mannitol. Underneath the flavor is the

nourishment we let the ferment provide without a word of promise: fiber, vitamins the bacteria synthesize and the brining preserves, and a living culture arriving in the food itself. The board is the design. The bacteria are the kitchen. We set the table and let the oldest cooks in the world do the cooking.

Chapter 8 · The quiet multiplier — *shiitake, mushroom, tomato, yeast, and the nucleotide side of savor.*

The loud organisms get the attention, but the organism that most often decides whether a savory dish is merely good or uncanny is one that barely tastes of anything by itself. Recall the multiplier of Chapter 3: glutamate sets the umami, but a 5′-nucleotide beside it — inosinate, guanylate — collapses glutamate's effective EC_{50} and multiplies the loudness, by as much as eight-fold, through the equation $y = u + \beta \cdot u \cdot v$ whose whole power lives in the cross term. To design great savor you must design *both* terms, and the second term, v, comes from a different kingdom of organisms than the first. This chapter is about the plants and fungi we engineer to hold v — and how we pair them.



The engineered tomato — our most developed glutamate platform. High free glutamate at ripeness is a loaded *u*, waiting for its nucleotide partner.

The great plant-kingdom store of guanylate is the **shiitake**, an eighty-construct host, and the way it gives up its nucleotide is a small masterpiece of enzymology worth doing right. A fresh shiitake holds its guanosine mostly as RNA; the famous depth of dried-shiitake stock is not released until the mushroom is dried — which concentrates everything and primes its enzymes — then soaked cold and brought gently up to around 60 to 70 °C, the window where the mushroom's own nuclease activity cleaves RNA into free 5′-GMP before a hotter heat would denature the enzyme. Soak cold, heat

slow: that is the whole secret of shiitake dashi, and it is enzyme kinetics disguised as a grandmother's instruction. We engineer the shiitake host for higher RNA and nucleotide yield and for robust nuclease activity, so that more guanylate is liberated under easier conditions — a higher, more reliable v in a form you can dry and crumble into anything. The **button mushroom**, a forty-construct host, and other fungi carry the work into milder registers.

The companion term, u, we grow in the **tomato** — at two thousand and forty constructs our most developed umami platform, edited to carry unusually high free glutamate at full ripeness. A ripe engineered tomato is already a loaded u before anything ferments; dry it, or ferment it, and the glutamate concentrates further. And then there is **yeast**, a two-hundred-and-forty-construct host that is quietly one of the most important organisms in the savory world, because at the end of its life it does something no plant does: it *autolyzes*. Its own enzymes digest it from within, breaking its RNA down into exactly the 5′-nucleotides that drive the synergy, releasing its stored glutamate, and freeing mannoproteins from its cell wall that read as kokumi body. Every yeast extract on every shelf is autolysis in a jar — a concentrated source of both u and v and K at once. We engineer the yeast host for higher nucleotide yield on autolysis and lower off-note, a cleaner multiplier.

The enzymology of freeing a nucleotide deserves one more sentence of precision, because it is where the craft hides its difficulty. A cell's guanosine and inosine sit locked inside RNA; to release the 5′-monophosphate the tongue can use, a 5′-phosphodiesterase must cut the RNA backbone on the correct side of each phosphate, and a 5′-nucleotidase must then be kept from going one step too far and stripping that phosphate off entirely, which would destroy the synergy. Dried shiitake performs this beautifully unaided — the drying concentrates the RNA and primes the enzymes, the cold soak lets them position, and the gentle climb to sixty-odd degrees runs the phosphodiesterase in its window before heat can denature it. Yeast performs it by autolysis, its own enzymes turned inward. The industrial yeast extract and the kaiseki kitchen are, in this one narrow respect, running the identical reaction. Our edits raise the RNA substrate and bias the enzyme balance toward the 5′-monophosphate and away from over-digestion — more v, held at exactly the right step.

Run the synergy equation on a real pairing and the number startles even us. Take a stock with a glutamate fraction u of about 0.5 and lift its guanylate fraction v from zero to a modest 0.1; with a guanylate synergy constant well north of a thousand, the cross term β·u·v alone contributes on the order of a hundredfold over the glutamate's own 0.5. That is not a seasoning adjustment — it is a different order of savor, reached with a pinch. It is also why the single most cost-effective move a savory kitchen can make is not to add more of its glutamate source but to add a little of a nucleotide source, and why a dust that carries both, pre-mixed and engineered to the cross term, is the most efficient umami on any shelf.

PAIRING	U SOURCE (GLUTAMATE)	V SOURCE (NUCLEOTIDE)	EQUIVALENT UMAMI VS SINGLE SOURCE
Kombu alone	kelp	—	baseline
Dashi	kombu	katsuobushi (IMP)	~7–8×
Shiitake–kelp (plant-first)	kombu / tomato	shiitake (GMP)	~6–7×
Engineered savory dust	dried tomato	engineered shiitake (GMP)	designed to the cross term

Read that last row as the thesis of the chapter made literal. If you co-dry and co-grind an engineered high-glutamate tomato with an engineered high-guanylate shiitake, you have made a single savory dust in which u and v are both high and *pre-mixed* — a spoonful that is not a flavor but a multiplier, climbing the synergy surface of Chapter 3 the instant it hits water. Fold that dust into a stock, a sauce, a ferment, and the glutamate already present in the dish is multiplied by the guanylate you just added. The grandmothers built dashi by pairing two ingredients from two kingdoms; we can build the pairing into a powder, with both terms specified in the seed.

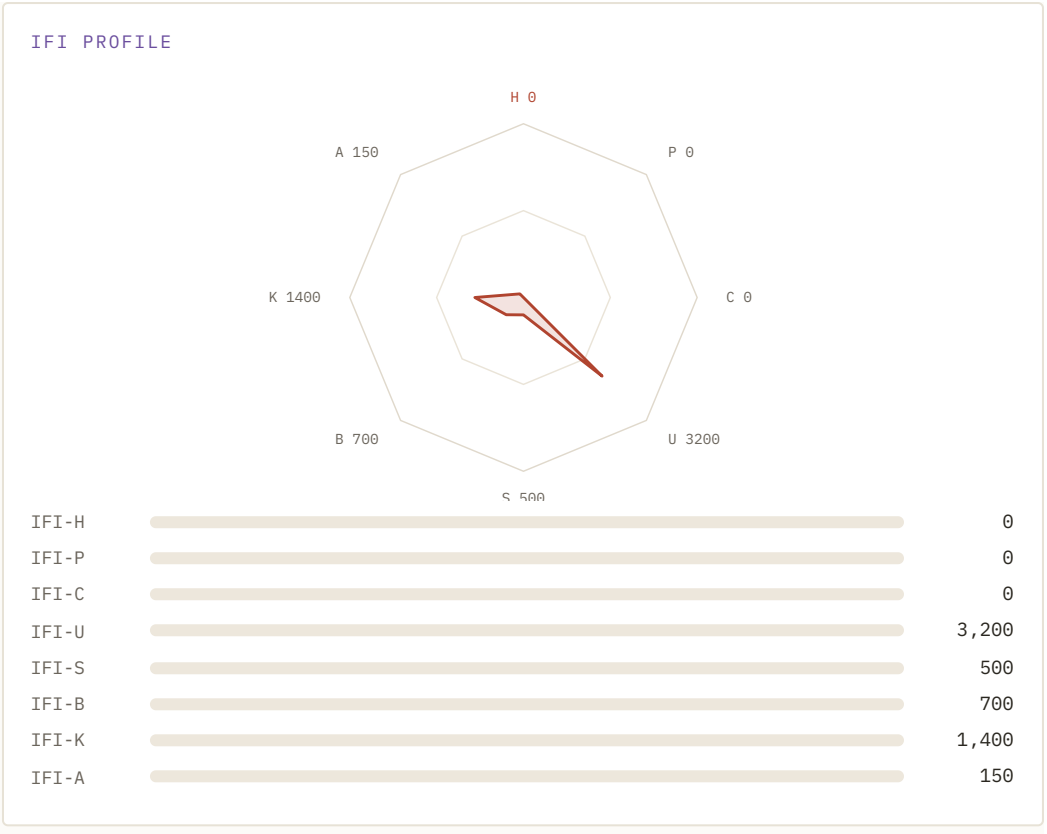

That is the savory dust, read alone — no heat, no pungency, just an enormous, filling umami with real body and a clean, attenuated bitterness, manufactured entirely from a loaded glutamate plant and a loaded nucleotide fungus. It is the quiet engine we will reach for again and again in the recipes to come, because it is the single most efficient way to climb the hardest axis. The loud organisms make a dish exciting. This quiet one makes it *deep* — and depth, as every cook and every ferment eventually learns, is the thing you cannot fake and cannot rush, only build, two terms at a time.

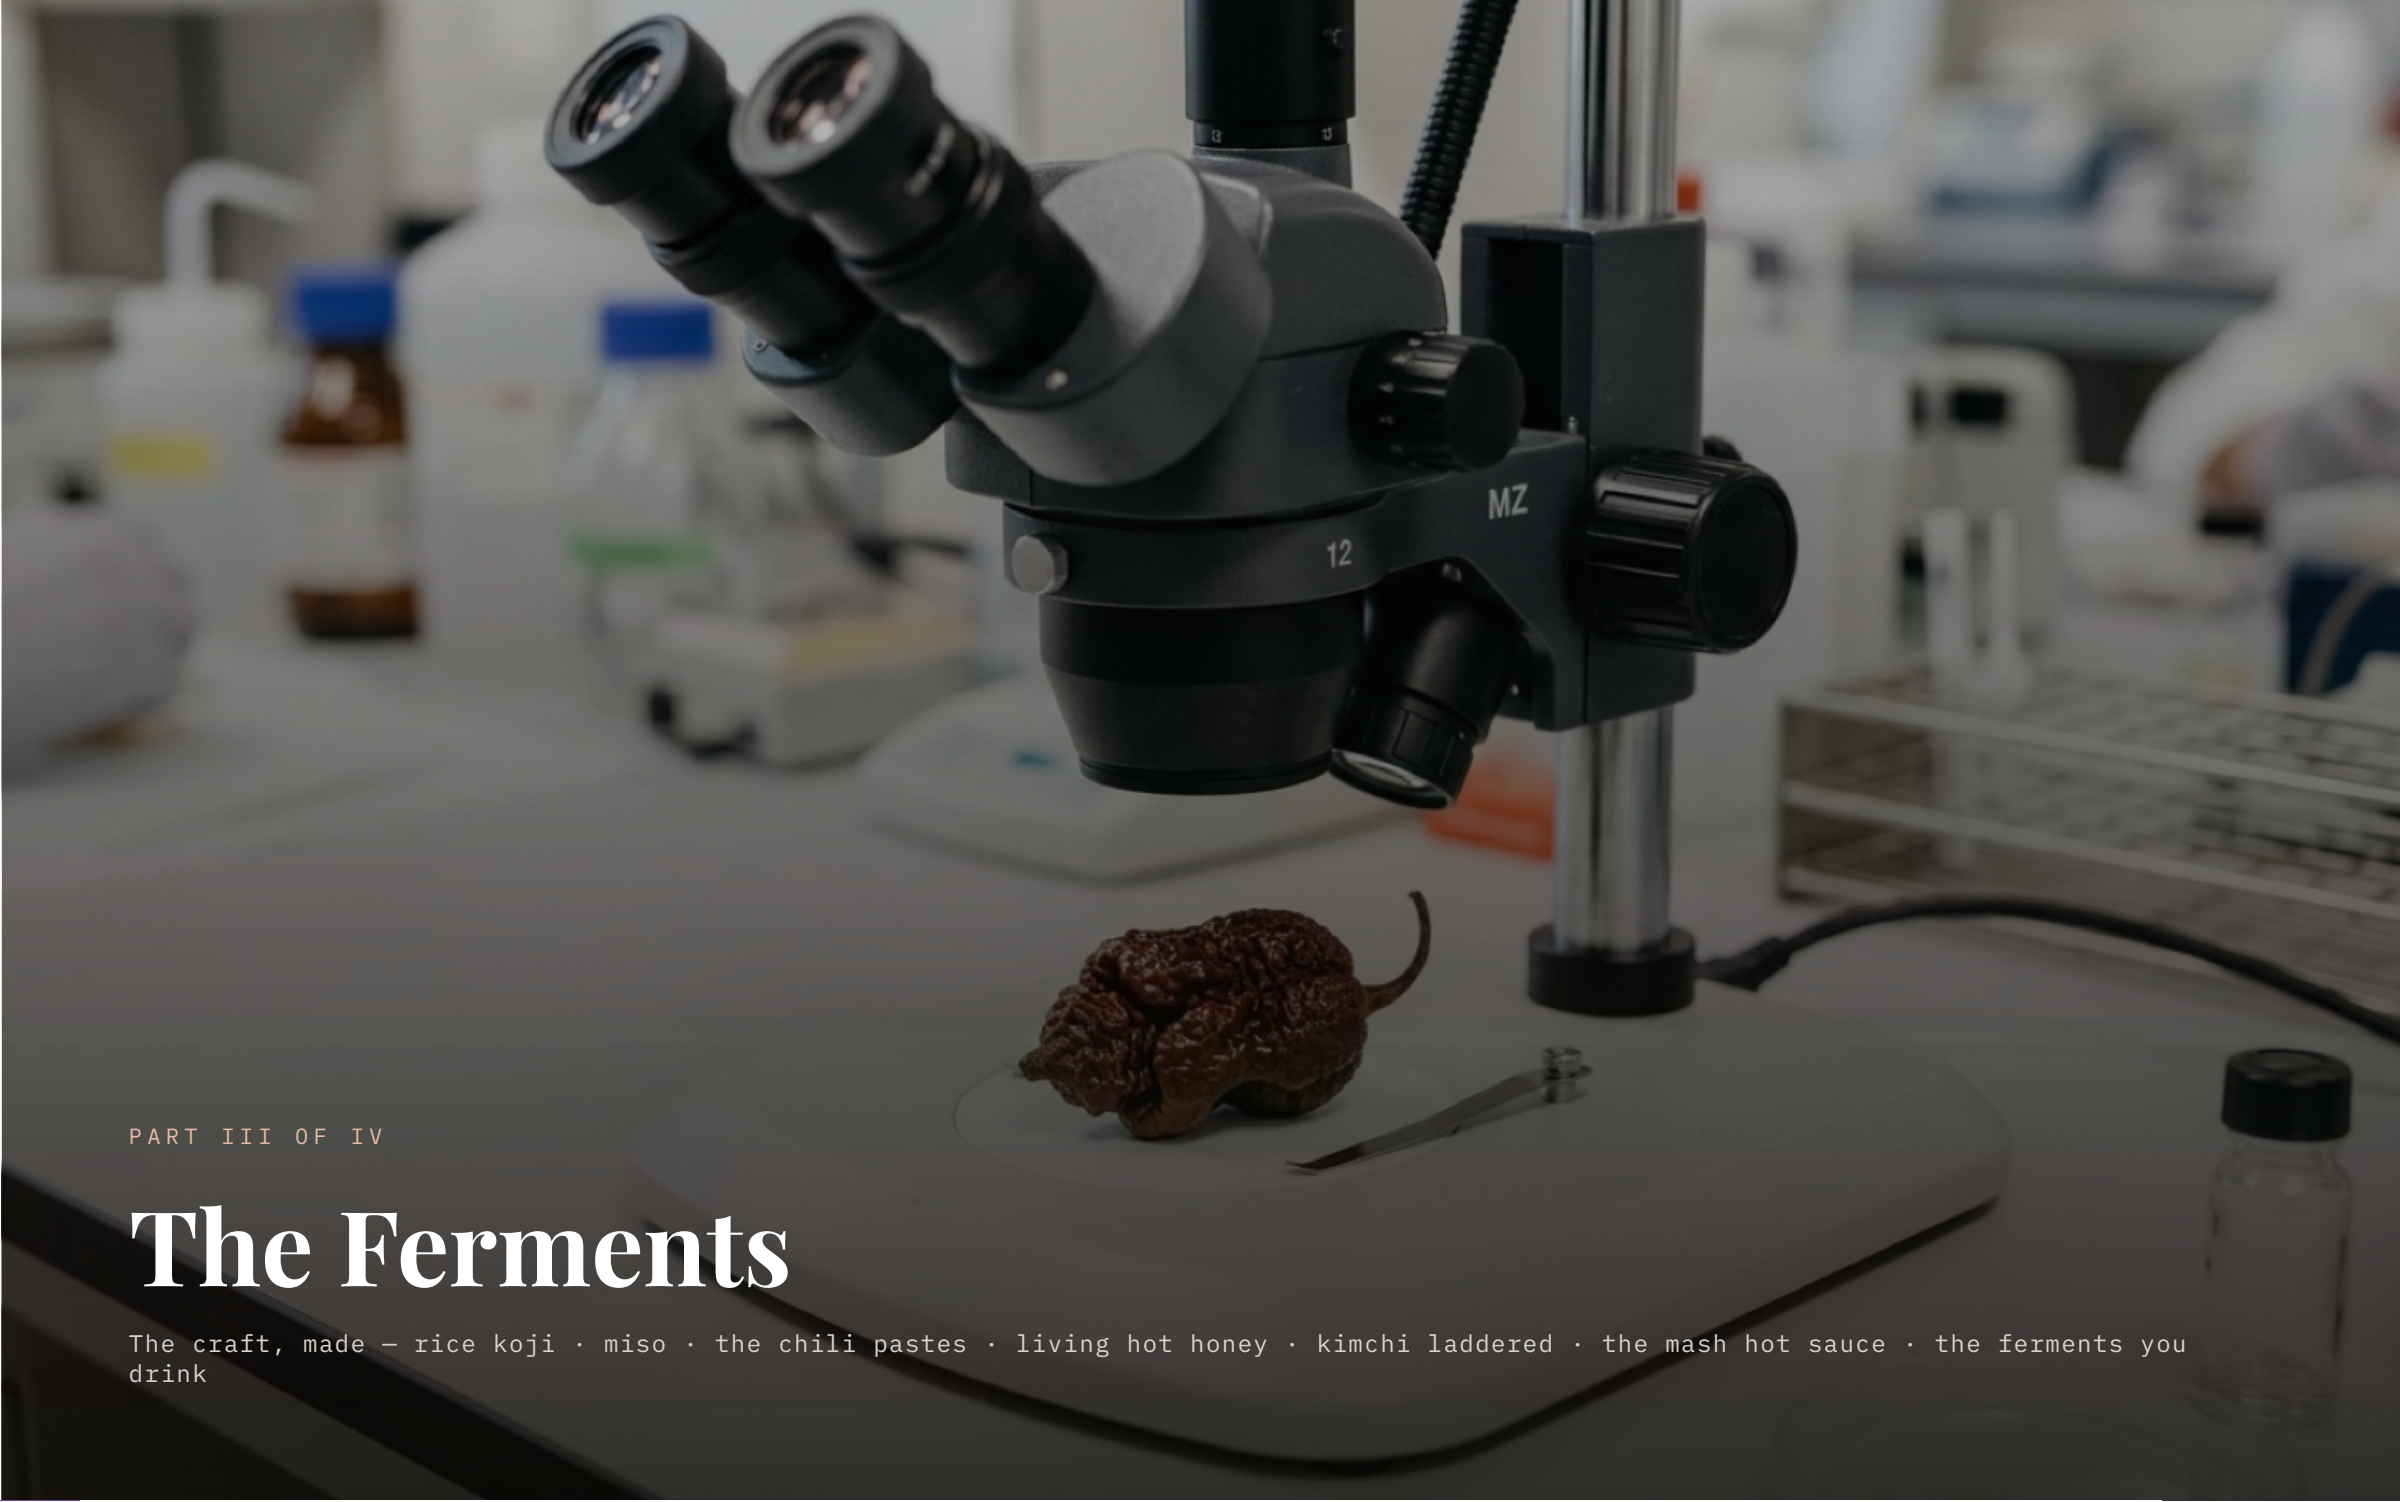

PART III OF IV

The Ferments

The craft, made — rice koji · miso · the chili pastes · living hot honey · kimchi laddered · the mash hot sauce · the ferments you drink

Chapter 9 · Rice koji and its sweet children — *grow the machine, then let it make sugar.*

Everything savory in Part III begins with the organism from Chapter 4, so Part III begins by growing it. Rice koji is the foundation stone of the Japanese ferments and the single most useful thing a home fermenter can learn to make, because once you have living koji you have, on your counter, the whole enzyme arsenal — amylase, protease, glutaminase, lipase — ready to point at any substrate you like. Grow it once and you will understand in your hands what the earlier chapters argued on paper: that a ferment is a factory, and that you have just built one.



Short-grain rice — steamed, cooled to blood-warm, and given to the mold. In two days it becomes a tray of enzymes.

Rice koji — *the factory itself.*

```
RICE KOJI · makes ~1.2 kg · 44–48 h
short-grain white rice     1 kg (dry weight)
tane-koji spore starter    1 g (≈ 1 tsp)
– incubation: 28–30 °C, 85–90% humidity
– equipment: steamer, clean cloth, a tray, a warm box
```

1. Rinse the rice until the water runs clear, then soak 8–12 hours; drain well for an hour so the surface is dry to the touch.
2. Steam — do not boil — for 45–60 minutes, until a grain crushes between your fingers with a firm, translucent core. Dry-cooked grains let the mycelium breathe; waterlogged ones smother it.
3. Cool to 35 °C, spread on a clean cloth. Dust the *tane-koji* through a fine sieve over the warm rice and mix thoroughly with clean hands so every grain is seeded.
4. Wrap loosely and hold at 28–30 °C. After about 18 hours the rice will clump and smell faintly sweet; mix to release heat (the mold's metabolism is exothermic and will cook itself). Mix again at ~30 hours, furrowing the surface to shed heat.
5. At 44–48 hours the grains are bound in white, fragrant mycelium that smells of chestnut and flowers — *before* any green of sporulation. Stop here. Refrigerate and use within a week, or dry

low for storage.

The two children of koji that every kitchen should know are made by pointing that factory at two different jobs, and the difference between them is nothing but which enzyme you favor — which is nothing but temperature. **Shio koji** is koji preserved in salt water, left at room temperature so its proteases run slowly over days into a seasoning that tenderizes and savors. **Amazake** is koji plus cooked rice held hot, at the amylase's sweet spot, so its starch-cutting enzymes run fast into a drink that is pure sweetness with no sugar added. Same organism; two destinies; one thermostat deciding between them.

Shio koji — *the all-purpose tenderizer.*

```
SHIO KOJI · makes ~500 g · 7–10 days
rice koji            200 g
fine sea salt        60 g   (~13% of total)
water                260 ml
```

1. Break up the koji and stir in the salt. Add water; it should just cover. Jar loosely.
2. Hold at room temperature, stirring once a day. Over a week it softens to a loose porridge and smells of melon, chestnut, and faint alcohol.
3. When a grain crushes to nothing and the aroma is sweet-savory, refrigerate. Use at about ten percent of a food's weight: rub into vegetables, tofu, fish, or meat an hour before cooking.

Shio koji works because its proteases cut surface protein into free amino acids — depositing glutamate (umami) directly onto whatever you rub it with — while its residual amylase leaves a trace of sweetness and its enzymes tenderize by loosening protein structure. It is a marinade that *builds flavor into the food* rather than merely coating it, and a spoonful rubbed onto a mild chili before roasting is the simplest possible demonstration of this whole book: savor manufactured on a surface, meeting heat.

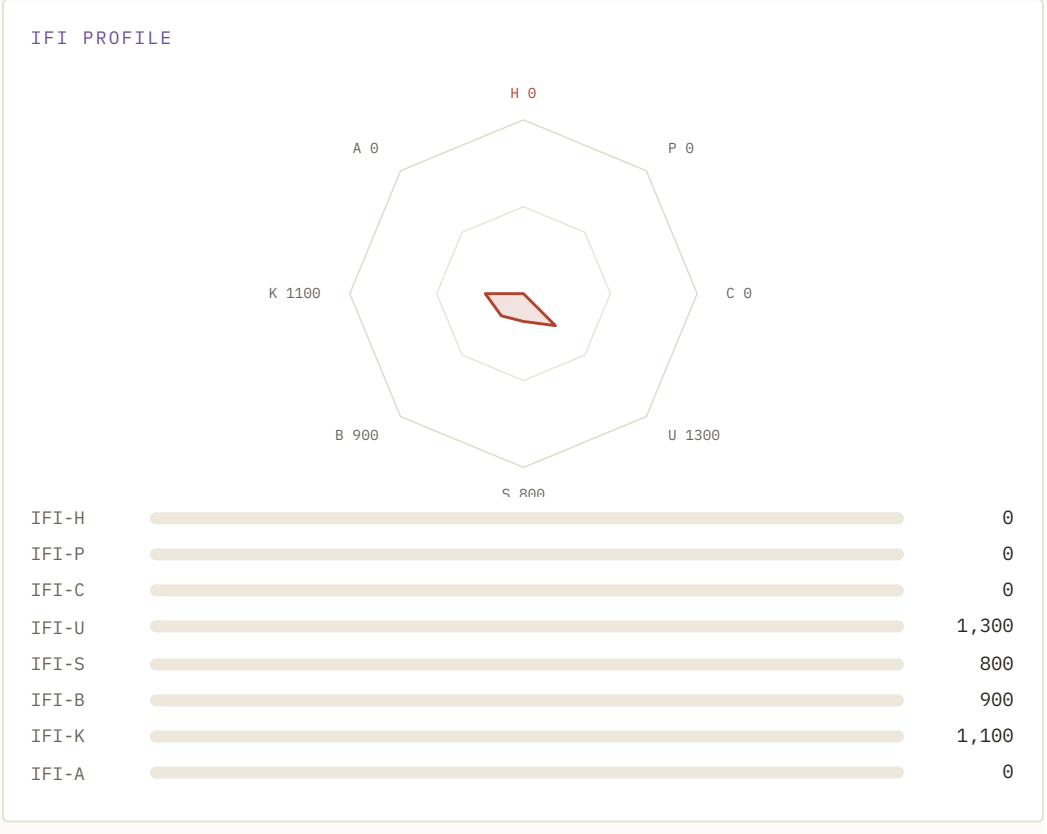

Amazake — *sweetness with no sugar added.*

```
AMAZAKE · makes ~1 L · 8–12 h
rice koji            250 g
cooked short rice    500 g
water                400 ml
— hold at 55–60 °C the entire time
```

1. Stir koji, cooked rice, and water together; the temperature matters more than anything else here.
2. Hold the mixture at 55–60 °C for 8–12 hours — a rice cooker on *keep warm* with the lid propped, a sous-vide bath, or a low oven with a thermometer. Stir every few hours.
3. It is done when it tastes fully sweet and the grains have half-dissolved. Blend smooth for a drink, or leave textured. Refrigerate up to a week.

The reason amazake is held at exactly 55–60 °C is pure enzymology: the glucoamylase and α-amylase that cut rice starch into glucose and maltose have their activity optimum in that window. Run it cooler and lactic bacteria wake and sour the batch; run it hotter and you denature the very enzymes doing the sweetening. Held in the window, the amylases convert starch to sugar until the mixture tastes, to many people, as sweet as a soda — all of it glucose and maltose the enzyme made from rice, none of it added. On the Index it reads almost purely on the S axis, and it is our favorite bridge to heat: a swicy drink is nothing more than amazake and a measured dose of fermented chili, sweetness and burn made by two cultures and married in a glass.

A word on why the window is so narrow and so worth hitting. The sweetness of amazake is built by two amylases dividing the labor: α-amylase chops the long starch chains into shorter dextrins at random interior points, and glucoamylase then marches along those pieces from the ends, releasing glucose one unit at a time. Both run fastest around 55–60 °C, and a brewer measures the result as a rising sugar concentration — the °Brix a refractometer reads climbing hour by hour as starch becomes sugar. Push past 65 °C and you denature the very enzymes doing the work, freezing the Brix wherever it stood; let it fall below 50 °C and lactic bacteria wake and sour the batch before it can sweeten. The whole art of amazake is to park the mixture in the enzymes' window and let the Brix climb to its ceiling undisturbed.

Koji's reach does not end at its two sweet children, either. A fish or a cut of meat rubbed in shio koji and left a day is being *aged* in fast-forward — the proteases depositing glutamate on the surface and tenderizing the interior, a dry-age compressed from weeks into hours — and a vegetable treated the same way gains a savory depth it could never hold raw. Every one of these is the same factory pointed at a new substrate, which is the lesson under the whole chapter: grow the enzymes once, and you can hand them anything. What follows is exactly that — the factory, pointed at a bean, a chili, a cabbage, each time making the molecule its particular receptor is waiting to hear.

IFI PROFILE

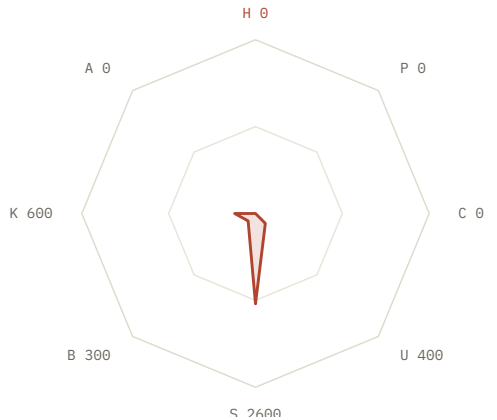

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	400
IFI-S	2,600
IFI-B	300
IFI-K	600
IFI-A	0

Chapter 10 · Miso, from the bean up — *the long patience, and what it builds.*

If rice koji is the factory, miso is the factory's masterpiece — the ferment that asks the most of your patience and returns the most for it. At its simplest it is three things: cooked soybeans, rice koji, and salt, packed into a vessel and left to age while koji's enzymes, trapped under salt, cut the bean apart over months or years. Everything about a miso's character — how savory, how sweet, how dark, how long it keeps tasting — is decided by the ratio of those three things and by how long you are willing to wait, which is to say by the two constants of Chapter 1: Glu_max, the savory ceiling, and k, the speed of the climb.



Soybeans, soaked and ready for the mash. Under salt and koji, two winters will turn them the color of old wood and the depth of a whole cuisine.

Red miso — *two winters to the bottom of savor.*

```
RED MISO · makes ~2.5 kg · 12–24 months
dry soybeans          1 kg  (→ ~2.2 kg cooked)
rice koji             1 kg
fine sea salt         450 g  (~13% of final weight)
seed miso (unpasteurized) 2 Tbsp
```

1. Soak the soybeans 18 hours, then simmer or pressure-cook until a bean crushes effortlessly — 3–5 hours by pot, 45 minutes under pressure. Reserve a cup of the cooking water. Cool to warm.
2. Mash the beans (a potato masher, or pulse) to a coarse paste. Mix the salt into the koji first, then fold koji-salt and the seed miso through the mashed beans, adding reserved water until the mixture holds together like damp dough.
3. Pack into a crock in fist-sized balls, pressing each down hard to drive out air — air pockets are where spoilage molds take hold. Smooth the top, dust with a little extra salt, and lay a sheet against the surface. Weight it.
4. Cover with cloth and age somewhere cool and stable. Check every few months; scrape any surface mold (it is cosmetic on a well-packed miso), re-weight. It is ready when it is the color of old wood and tastes of the bottom of savor — 12 months for young, 24 for deep.

Three slow chemistries build that depth, and they are worth naming because they are the same three the earlier chapters promised. First, **proteolysis**: koji's proteases and peptidases, working under salt, cut the bean's storage protein to peptides and then to free amino acids, and its glutaminase converts bound glutamine to glutamate — the climb of Glu toward Glu_max, slow because the salt keeps k small, high because two years is a long integral. Second, the **Maillard reaction**: over months the free amino acids react slowly with the bean's sugars to form melanoidins, the brown pigments that darken a red miso and round its flavor — and melanoidins measurably attenuate bitterness, which is why a long-aged miso reads high on our IFI-B even though nothing was removed. Third, **γ-glutamyl peptide formation**, the koji-driven chemistry that lays down the CaSR-active peptides giving a great miso its kokumi — the mouth-filling body and the length, the way the taste keeps going after you swallow.

Those three chemistries are dialed by the same two numbers from Chapter 1, Glu_max and k, and the world's misos are simply different settings of them. A high ratio of koji to bean, low salt, and a warm fortnight give a pale, sweet, amylase-bright *shiro*; a balanced ratio and a year give the everyday *shinshu*; a bean-koji base, high salt, and two or three years give the dark, dense, barely sweet *hatcho*; barley koji gives the earthy country *mugi*. The map is legible once you hold the dials.

MISO	KOJI : BEAN	SALT	AGE	CHARACTER
Shiro (white)	high	5–6%	2–8 wk	sweet, light, amylase-forward
Shinshu (yellow)	medium	11–12%	6–12 mo	balanced, everyday
Aka (red)	medium–low	12–13%	1–2 yr	deep, savory, Maillard
Hatcho (bean)	bean koji	10–11%	2–3 yr	dense, dark, barely sweet
Mugi (barley)	barley koji	9–12%	1–3 yr	earthy, nutty

The browning that darkens the long ones is Maillard chemistry run cold and slow: free amino acids condense with the bean's sugars into Amadori compounds and, over months, into the brown

melanoidins that give a red miso its color, its roundness, and — because melanoidins measurably blunt the bitter receptor — its high reading on our B axis. Nothing is removed to attenuate that bitterness; a new class of molecule is slowly built that stands in front of it. Time, on this axis, is not patience for its own sake. It is synthesis.

One rule of use that is really a rule of respect: never boil miso. Heat above a simmer drives off its aromatics and kills the live culture that is part of what you made. Whisk it into a broth off the heat, at the end. You spent two winters building something alive; do not undo it in the last thirty seconds.

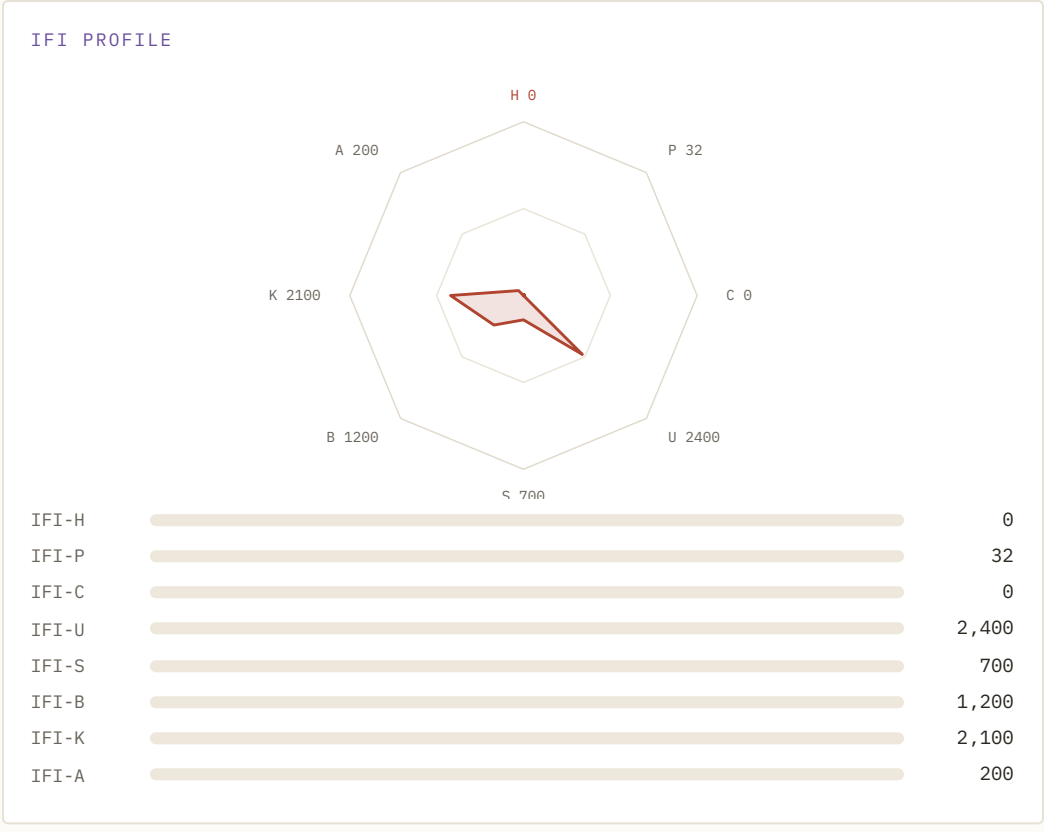

Spicy miso — *the bridge, in one spoon.*

The fastest way to see this book's two halves meet is to fold a tablespoon of any fermented chili — the mash from Chapter 14, or a scoop of gochujang — into a quarter cup of this miso. The miso brings the manufactured umami and kokumi; the chili brings capsaicin already lifted by its own ferment's acid. Taste the plain miso, then the spiked one, and you are tasting the whole argument: savory depth built in one jar, heat built in another, married in a bowl, each one louder for the pH and the glutamate the other brought.

IFI PROFILE

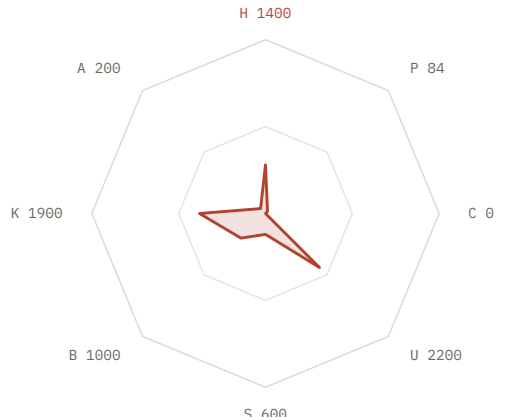

IFI-H	1,400
IFI-P	84
IFI-C	0
IFI-U	2,200
IFI-S	600
IFI-B	1,000
IFI-K	1,900
IFI-A	200

Chapter 11 · The fermented chili pastes — *doubanjiang, gochujang, and a living chili crisp.*

The great fermented chili pastes of Asia are where this book's two subjects have been married for centuries, by cooks who never needed the receptor diagram to know that a chili aged under salt and culture becomes something a fresh one can never be. Three of them deserve a fermenter's full attention, because each solves the heat-plus-depth problem differently: doubanjiang through the long sun-aging of beans and chili together, gochujang through a malt-driven sweetness that balances the burn, and chili crisp through the marriage of a lacto-ferment with hot aromatic oil.



Peppers bound for the paste. Fermented under salt, their heat will climb at the receptor even as their flavor deepens.

Doubanjiang, the soul of Sichuan cooking, is fermented broad (fava) beans and chilies aged together under salt, traditionally in earthenware open to the sun and hand-turned for anywhere from six months to three years — Pixian's finest are aged like wine. The sun-aging drives both proteolysis and a slow Maillard, so the paste develops a deep, almost meaty savor beneath its heat, and the chili's burn, carried down to low pH over the long ferment, reads bright and lifted. It is the purest example of the Chapter 2 mechanism operating over years rather than weeks.

Doubanjiang rewards a closer look, because it is fermentation and heat aged to the patience of wine. The broad beans are first made into their own koji, grown until enzyme-rich, then salted into a mash; the chilies — classically the slender, fragrant *er jing tiao*, prized for aroma over brute heat — are salted separately and married in. What follows is the rhythm of Pixian: the jars stand open to the sun by day and the dew by night, and a worker turns the mash by hand, over and over, for months into years. The sun drives Maillard browning and evaporative concentration; the turning aerates and homogenizes; the salt holds k low so the whole thing matures with glacial depth. A one-year doubanjiang is good; a three-year is a different substance, nearly black, its heat folded so far into savor that it reads less as spice than as a dark, resonant ground beneath

everything else in the wok. It is the slowest expression in this book of the single idea that a chili under salt and time becomes something a fresh one never can.

And the *málà* buzz has a mechanism worth naming exactly: beyond its touch on the cooling channel TRPM8, hydroxy-α-sanshool quiets the two-pore potassium channels that normally hold light-touch fibers at rest, so those nerves fire in the fluttering rhythm the brain reads as numbness — a tingle that is neither taste nor heat but a third thing, laid over the capsaicin fire to make the two-part sensation no single channel could produce.

Gochujang is the cleverest of the three, because it solves the sweet-heat balance with an enzyme. Its recipe includes barley malt powder — *yeotgireum* — whose amylase, held warm with cooked glutinous rice, saccharifies the rice starch into sugar before the long ferment. That enzymatic sweetness (pure IFI-S, no sugar added) is what balances the gochugaru's heat into gochujang's signature swicy roundness, while fermented soybean powder (meju) brings the umami. It is amazake's sweetness and miso's savor and a chili's heat, resolved into one paste.

Gochujang — *sweet heat, built by an enzyme.*

GOCHUJANG · makes ~2 kg · 2–6 months	
gochugaru (engineered flake, chosen band)	500 g
glutinous rice flour	300 g
barley malt powder (yeotgireum)	100 g
fermented soybean powder (meju)	400 g
fine sea salt	180 g
water	1.2 L

1. Whisk the rice flour and malt powder into the water. Hold at 60 °C for 2 hours, stirring: the malt's amylase converts the rice starch to sugar and the mixture visibly thins and sweetens.
2. Simmer gently to reduce by about a third, to a thick, glossy, sweet porridge. Cool to warm.
3. Beat in the meju powder, gochugaru, and salt until smooth and uniform. It should be thick but spreadable.
4. Pack into a crock, smooth the top, cover with cloth, and ferment warm with some daylight — traditionally on a rooftop — for 2 to 6 months, stirring occasionally. It deepens in color and savor as it ages.

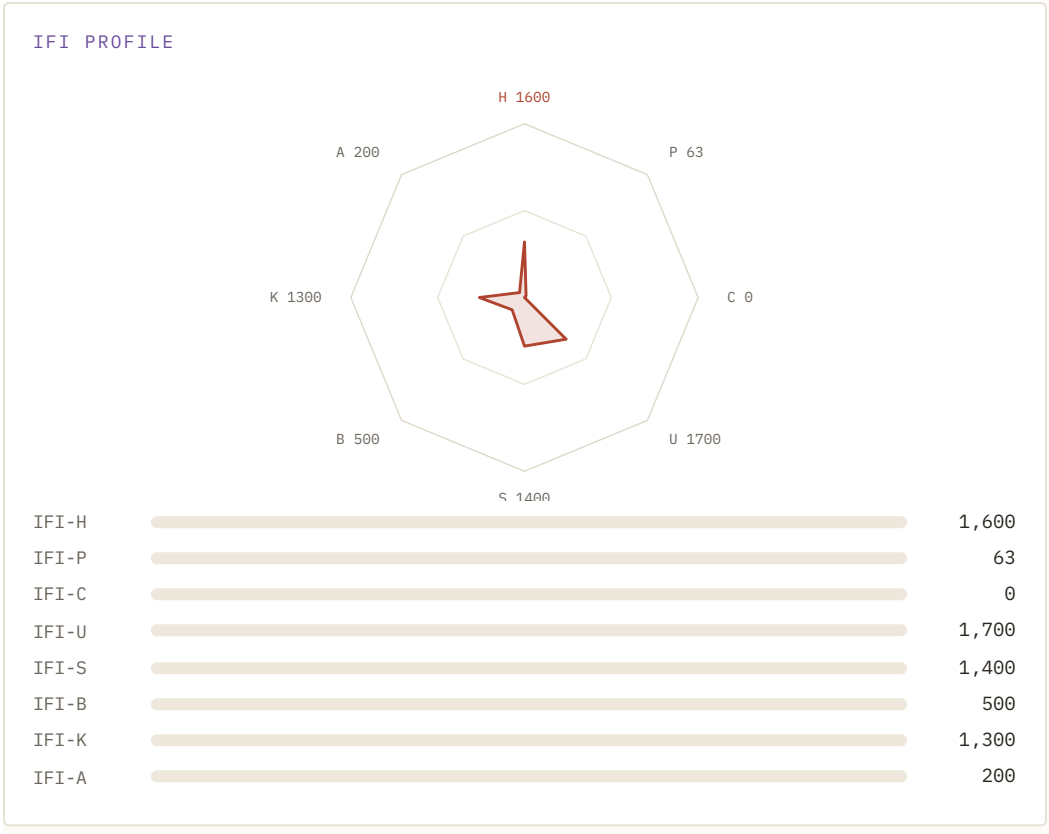

Chili crisp is the modern member of the family, and a living one if you build it that way. Rather than frying fresh chili, we lacto-ferment the flake first — a short brine ferment that drops its pH and lifts its heat by the Chapter 2 mechanism — then dry it and marry it to aromatic oil. The Sichuan peppercorn here is not incidental: its active molecule, hydroxy-α-sanshool, is a cooling-and-tingling agent that acts on TRPM8 and on touch-sensing fibers, producing the *málà* numbness that rides alongside the heat. That tingle registers on our Cool-Burn axis, which is why a true chili crisp is not merely hot but hot-and-humming — two channels, TRPV1 for the burn and the sanshool pathway for the buzz, firing together.

Living chili crisp — *fermented flake, hot oil, the málà hum.*

LIVING CHILI CRISP · makes ~700 ml	
lacto-fermented chili flake (dried)	60 g
neutral oil	400 ml
shallot, thinly sliced	100 g
garlic, sliced	60 g
Sichuan peppercorn	2 Tbsp
fermented black bean (douchi), minced	3 Tbsp
savory dust (Ch. 8)	1 Tbsp
sugar	2 tsp
salt	to taste

1. Put the dried fermented flake, douchi, savory dust, sugar, and salt in a heatproof bowl.
2. Fry the shallot and garlic in the oil over medium heat until golden and crisp; lift them out onto the flake mixture.
3. Add the Sichuan peppercorn to the hot oil for 30 seconds until fragrant, then strain the oil (at about 190 °C) directly over the bowl — it will foam and bloom. Stir, cool, and jar. Better after a day.

IFI PROFILE

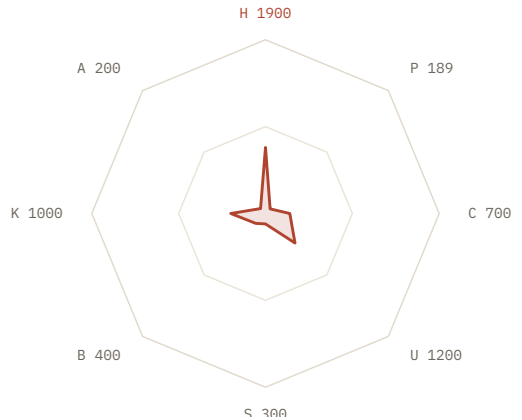

IFI-H	1,900
IFI-P	189
IFI-C	700
IFI-U	1,200
IFI-S	300
IFI-B	400
IFI-K	1,000
IFI-A	200

Chapter 12 · Living hot honey — *the honest ferment, and its vegan twin.*

Hot honey is the sweetest-selling idea in condiments right now, and most of what is sold is neither fermented nor living — it is honey with chili infused, which is a fine thing but not a ferment. The reason is worth stating honestly, because it is the clearest lesson in water activity you will ever taste: **honey does not ferment, on purpose.** Its water activity sits around 0.5 to 0.6 — far below the roughly 0.6 floor that even the most osmotolerant microbes need to grow — and it is additionally defended by low pH, by gluconic acid, and by the trace hydrogen peroxide its glucose-oxidase throws off. Honey is self-preserving. A jar of honey and chili will infuse, but it will not come alive, because nothing can grow in it.

The number behind that is water activity, a_w — not how much water a food holds but how much of it is *available*, free rather than bound up by sugar or salt, on a scale where pure water is 1.0. Honey's enormous sugar concentration binds its water so tightly that its a_w sits near 0.5 to 0.6, and the floor for microbial growth is around 0.6; below it, the osmotic pressure outside a cell exceeds what the cell can resist, and water is pulled *out* of any microbe that lands, dehydrating it where it sits. Honey does not kill by poison but by thirst. It is reinforced by a low pH and by the gluconic acid and trace hydrogen peroxide its own glucose-oxidase throws off, and the only organisms that trouble it at all are a few xerophilic yeasts — *Zygosaccharomyces* among them — built to live in syrups. Dilute the honey past roughly nineteen percent moisture and you lift its a_w over the floor; now those osmotolerant yeasts wake, and the honey will ferment toward mead whether you meant it to or not. The threshold is exact, which is precisely why a living hot honey has to be built deliberately rather than stumbled into.



Golden light through the greenhouse — the color of honey, and of the sweet-heat axis this ferment lives on.

So to make a hot honey that is genuinely living, you must be honest about the physics and choose one of two real routes. The first keeps the honey as honey and lets a *separate* ferment ride in: you fold a

lacto-fermented chili mash — already sour, already alive, its heat already lifted at the receptor by its own acid — into raw honey, cold, so the honey preserves the whole and the culture arrives with the chili. The second route accepts that to ferment honey you must dilute it: thin it with water past about nineteen percent moisture and osmotolerant yeasts will slowly turn it toward a short, sparkling chili session-mead. We give the first here, because it is the one that belongs on a biscuit and a wing; the second is a drink, and it lives in Chapter 15.

Living hot honey — *cold-folded, culture intact.*

```
LIVING HOT HONEY · makes ~500 ml
raw honey                        450 g
lacto-fermented chili mash       60 g  (Ch. 14, or 3 Tbsp)
  (drained of excess brine)
flaky salt                       1 tsp
```

1. Warm the honey only to body temperature — never hot, which would kill the culture and drive off the volatiles you want.
2. Fold the fermented chili mash and salt through the honey until streaked and uniform. Do not heat.
3. Jar and rest a week in the refrigerator so the heat and sour marry into the sweet. Keep cold; it is alive.

The flavor is swicy in the exact IFI sense — the H and S axes both high, the dominant trend in modern condiments — but with a dimension infused honey cannot reach: the chili's heat is already lifted by the ferment's acid, so it reads hotter than its capsaicin alone, and a faint lactic sourness cuts the honey's sweetness so it does not cloy. It is the sweet-heat intersection with a live culture folded through it.

Swicy is worth a moment of receptor theory, because it is not obvious that heat and sweetness should flatter each other — they are read by entirely separate systems, capsaicin at TRPV1 on the trigeminal nerve and sugar at TAS1R2·TAS1R3 on the taste buds, with no shared receptor for them to compete at. That separation is exactly why the pairing works: neither masks the other, so both arrive at full volume, and the only interaction happens upstream in the brain, where a little sweetness softens the aversive edge of the burn and a little heat keeps the sweetness from cloying. Our combinatorics work treats this as a dual-axis target — an M2 pairing, in the Index's shorthand — and designs to it directly, choosing a chili band and a sweetness level that land the H and S axes where the two sensations hold each other in balance rather than one burying the other. The living version adds a third quiet note the diagram predicts: the fermented chili's own acids, lactic from the mash and a little gluconic the honey already carried, drop the pH just enough to lift the heat at the receptor and to cut the sugar's roundness with a clean sour edge. Three molecules, three systems, one spoonful — sweet at the bud, hot at the nerve, sour at the gate, and none of them in another's way.

IFI PROFILE

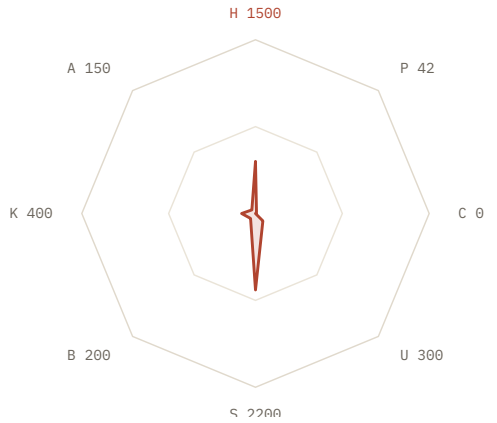

IFI-H	1,500
IFI-P	42
IFI-C	0
IFI-U	300
IFI-S	2,200
IFI-B	200
IFI-K	400
IFI-A	150

Fermented hot syrup — *the vegan twin, with equal respect.*

Honey is a line some at our table do not cross, and the plant kingdom answers beautifully, so we give its twin the same care rather than treating it as a substitute. Reduce engineered sugarcane juice — our sugarcane host carries two hundred constructs on the sweetness axis — to a thick syrup, or use a date syrup, and cold-fold the same lacto-fermented chili through it. The sweetness profile differs (sugarcane's is cleaner, date's is darker and more mineral), but the structure is identical: a high-S, high-H living drizzle, the sour lift and the receptor-raised heat folded into a plant-made sweet. Nothing here is an apology. It is simply the same idea, grown a different way.

IFI PROFILE

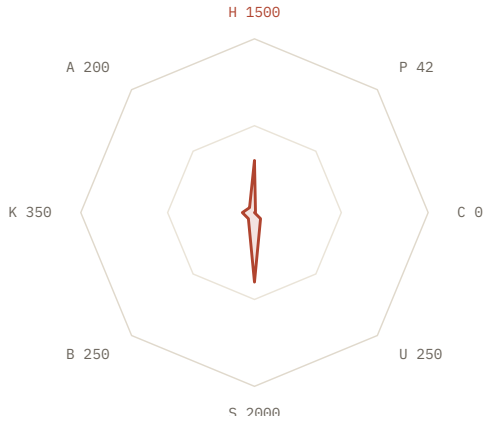

IFI-H	1,500
IFI-P	42
IFI-C	0
IFI-U	250
IFI-S	2,000
IFI-B	250
IFI-K	350
IFI-A	200

Those are the pastes and the sweet-heat drizzles. Now the vegetables come back — kimchi laddered deliberately across four heat bands — and then the loudest ferment of all: a mash hot sauce aged like wine.

Chapter 13 · Kimchi, laddered — *one board, four heats, the same lift under each.*

Chapter 7 built the kimchi board as a consortium; now we cook it, and we do the one thing the Index makes possible that tradition never quite could — we ladder it, deliberately, across four heat bands, knowing exactly how hot each will read because we chose the chili to a number and we understand what the ferment will do to that number. The beauty of laddering a ferment is that the Chapter 2 multiplier is roughly *band-independent*: the fall in pH lifts the effective heat of a mild kimchi and an extreme one by about the same proportion, so a ladder built in the fresh paste stays proportional in the jar. You set the rungs with the chili; the culture lifts the whole ladder.

The one number the ladder cannot be careless about is the salt, because salt is what lets the whole board ferment cleanly rather than rot. Brining the cabbage is osmosis put to work: the salt outside the leaf pulls water out across the cell membranes until the tissue goes limp and pliable and its own expressed liquid becomes the brine the culture will live in. Too little salt and the water activity never falls far enough to exclude the soft-rot bacteria, and the kimchi turns to mush; too much and even the lactic bacteria stall, and the ferment sits sullen and un-soured. The target is a finished salt of roughly two to three percent — the window where *Leuconostoc* and *Lactobacillus* thrive and almost nothing else does, the same window Chapter 16 maps across every ferment in the book. Salt the cabbage, taste the brine so it reads of the sea rather than of a mouthful of salt, and you have set the one condition that lets all four rungs of the heat ladder climb safely to their sour.



Napa cabbage in cross-section — the leaf's architecture is why brining and the ferment reach all the way through.

Baechu kimchi — *the master board.*

```
BAECHU KIMCHI · makes ~1.5 kg · 1–5 days + cold age
napa cabbage                 1 large (~1 kg)
coarse salt (for brining)    ½ cup
— rice porridge:
glutinous rice flour         2 Tbsp
water                        1½ cups
sugar                        1 Tbsp
— seasoning:
gochugaru                    3–10 Tbsp (see ladder)
garlic, grated               6 cloves
ginger, grated               1 Tbsp
fish sauce OR kelp-shiitake stock + salt   3 Tbsp
daikon radish, julienned     1 cup
scallions                    4, cut
```

1. Quarter the cabbage lengthwise, salt between the leaves, and leave 2 hours until the leaves bend without snapping. Rinse three times and drain.
2. Cook the rice flour, water, and sugar to a loose porridge; cool. This feeds the lactic bacteria and carries the seasoning into every fold.
3. Blend the porridge with gochugaru, garlic, ginger, and fish sauce (or the vegan stock and salt — written, as always, with equal care). Fold in daikon and scallion.
4. Coat every leaf with the paste, pack tight into a jar leaving headspace, and press down so brine rises to cover. Ferment 1–5 days at room temperature — tasting daily, watching for bubbles and sour lift — then refrigerate. It keeps deepening for weeks.

BAND	CHILI CHOICE	PASTE HEAT (SHU)	EFFECTIVE IFI-H (FERMENTED)
Mild	mild gochugaru only	~1,500	~250
Medium	hot gochugaru	~15,000	~700
Hot	gochugaru + 2 Tbsp Ember mash	~60,000	~1,900
Extreme	gochugaru + 1 Tbsp Wildfire mash	~250,000	~3,900

The rungs are set in the fresh paste and then, over the ferment's five days, every one of them climbs as the *Leuconostoc* and *Lactobacillus* drop the pH from six toward four-point-two. A cook who wants a reliable "hot but not punishing" kimchi can now hit it on purpose — choose the chili to the band, let the culture apply its proportional lift, and the jar lands where the ladder said it would, batch after batch, because the chili's own coefficient of variation was locked in the field at seven percent or below. This is what the Index is *for*: not to take the life out of a ferment, but to let you aim it.

A kimchi is never finished, which is both its glory and the thing a cook must plan around. Even in the cold the lactic bacteria keep working slowly, so a jar moves over weeks from fresh and crunchy through its bright, balanced prime and on toward a deep, sour maturity that is less a side dish than an ingredient — the old, sour kimchi that belongs in a stew. The *Leuconostoc* of the early days leaves behind mannitol, a sugar alcohol that tastes faintly sweet and cooling and gives a young kimchi its lift; the carbon dioxide those same bacteria throw off is the small effervescence on the tongue and the pressure that will, if you forget, lift a lid in the night.

Refrigeration does not stop any of this. It only slows the climb to a crawl, buying weeks of prime. To know your kimchi is to know where on that curve you like it, and to cool it there.

Kkakdugi — *the radish cube, a faster board.*

```
KKAKDUGI · makes ~1 kg · 1–3 days
daikon radish, cubed      1 kg
coarse salt               1 Tbsp
sugar                     1 tsp
gochugaru                 3–6 Tbsp (to band)
garlic, grated            4 cloves
ginger, grated            2 tsp
fish sauce or kelp stock  2 Tbsp
scallion                  3, cut
```

1. Toss the radish cubes with the salt and sugar; rest 30 minutes and drain the water they give up.
2. Mix in the gochugaru first — it dyes the radish evenly — then the rest. Pack tight so brine rises over the cubes.
3. Ferment 1–3 days at room temperature until bubbly and tart, then refrigerate. The dense radish stays crisp far longer than cabbage, and its mild sweetness meets the heat halfway.

IFI PROFILE

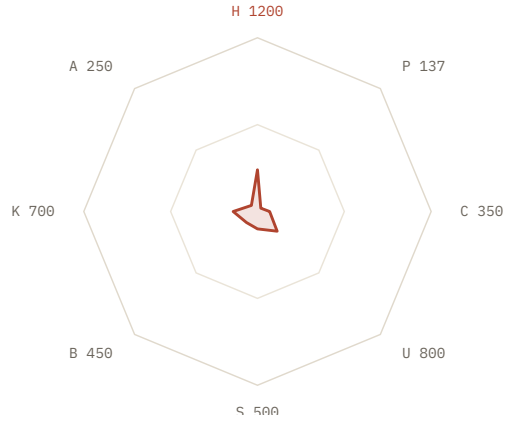

IFI-H	1,200
IFI-P	137
IFI-C	350
IFI-U	800
IFI-S	500
IFI-B	450
IFI-K	700
IFI-A	250

IFI PROFILE

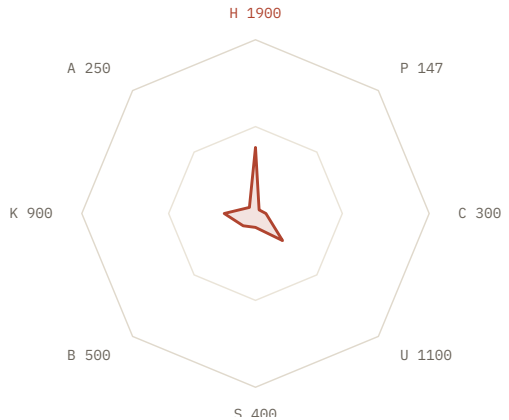

IFI-H		1,900
IFI-P		147
IFI-C		300
IFI-U		1,100
IFI-S		400
IFI-B		500
IFI-K		900
IFI-A		250

Chapter 14 · The fermented hot sauce — *mash, brine, and the barrel that waits.*

Here is the loudest thing in the book, and the purest expression of its title. A fermented hot sauce is nothing but chilies, salt, and time, and yet it delivers a heat that no fresh sauce of the same pepper can touch — because, as Chapter 2 showed, the weeks under brine are weeks of falling pH, and falling pH is rising heat at the receptor. Build a mash sauce on one of our flagship superhots, grown to a locked number, and the ferment takes that number and lifts it; the result is reproducible, three-dimensional, receptor-amplified heat in a bottle.



A flagship pepper at harvest, bound for the mash. What the field locked to a number, the jar will lift.

Fermented mash hot sauce — *the method, and three ways to age it.*

FERMENTED MASH HOT SAUCE · makes ~750 ml · 2–4 wk + blend	
ripe chilies (engineered, chosen band)	1 kg
fine sea salt	30 g (3% of chili weight)
garlic (optional)	4 cloves
filtered water	to cover if needed
white or cider vinegar	~120 ml (to finish)

1. Stem the chilies (wear gloves at these heat levels — a flagship mash is no place for bare hands). Blend coarsely with the salt and garlic to a wet mash, or pack whole under a 3–5% brine.
2. Pack into a jar with real headspace; the ferment will foam. Keep everything submerged under a weight, with an airlock or a daily "burp" to release the CO_2 the *Leuconostoc* throws off.
3. Ferment at room temperature 2–4 weeks. It will cloud, bubble, and turn pleasantly sour as the pH falls from about 6 to below 4.
4. When it tastes bright and sour — pH 3.4 to 3.8 — blend smooth, then cut with vinegar to bring the finished sauce below pH 3.5 for shelf stability. Strain if you like it smooth, bottle, and age at least a week more before using.

Three ways to take it further, each a real tradition. The **quick mash** above is a fresh, bright Louisiana-style sauce. The **barrel** is the Tabasco path: salt the mash heavily, seal it under a salt cap in oak, and let it rest one to three years — the long ferment and slow oxidation round the heat into something deep and almost smoky, the acid lifting it the whole time. The **sriracha** path ferments a chili-and-garlic mash a few bright days, then blends it smooth with a little sugar and vinegar into the familiar glossy, swicy, pourable sauce. One mash, three destinies, set by how long you let the culture and the air work.

Under the surface the microbiology runs the same succession as the kimchi board, in a vessel built for heat. The first days belong to *Leuconostoc*, throwing off the carbon dioxide that foams the mash and the early acids that begin the pH fall; by the second week the hardier *Lactobacillus* have taken over and drive the acidity past four and on toward three-point-something, where the mash goes quiet and stable. Keeping the solids submerged through all of it is the one non-negotiable — exposed mash grows surface mold and Kahm yeast, submerged mash ferments clean — which is why the airlock, the weight, and the daily burp are not fuss but method. When the probe reads pH 3.4 to 3.8 and the foam has stopped, the ferment has finished making heat the slow way, and it is time to stabilize it.

Stabilizing means acidifying to a shelf-safe pH below 3.5, and the number is a floor, not a preference. A bottled sauce that drifts back above pH 4.6 in storage is one in which dangerous things could in principle grow; cutting with vinegar to below 3.5 buys a margin wide enough to hold as the bottle sits. The old Louisiana makers reached it with cider or distilled vinegar; the Tabasco house reaches it by aging the salted mash in oak for up to three years, then cutting it with vinegar at the end — the long barrel rounding the heat while the acid preserves it. Either way the finished bottle, slim-necked and shaker-topped, carries a heat that is high, lifted at the receptor, and, by that very same low pH, safe to keep on a shelf for years.

SAUCE, BUILT ON	FRESH IFI-H	FERMENT	EFFECTIVE IFI-H (BOTTLED)
Ember · SS-0087	1,570	3-week mash	~2,350
Blaze · SS-0091	2,630	3-week mash	~3,900
Wildfire · SS-0092	2,930	3-week mash	~4,350
Firestorm · SS-0094	3,490	3-week mash	~4,650

IFI PROFILE

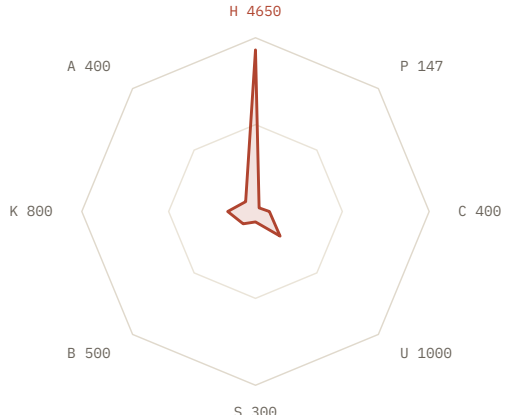

IFI-H	4,650
IFI-P	147
IFI-C	400
IFI-U	1,000
IFI-S	300
IFI-B	500
IFI-K	800
IFI-A	400

That is a mash sauce built on Firestorm — a six-axis flagship — read after three weeks of fermentation and a vinegar finish. The heat is near the top of the scale, lifted there not by capsaicin we added but by the protons the culture laid down; the umami and kokumi are the ferment's own gift; the whole thing is louder, rounder, and more alive than any fresh purée of the same fruit. A note we will make properly in the next part but must flag here: the vinegar finish below pH 3.5 is not only for flavor. It is the safety floor. A fermented, acidified hot sauce is shelf-stable precisely because it is too sour for anything dangerous to live in — the same low pH that makes it hot makes it safe, one number doing both jobs, as it has all through this book.

Chapter 15 · The ferments you drink — *kombucha, tepache, and the chili session.*

Not every ferment ends on a plate; some of the oldest and most refreshing end in a glass, and they let us close the craft with the lightest, liveliest members of the family — the fizzing, sour, faintly alcoholic drinks that are fermentation at its most immediate. They also let us keep a promise from Chapter 2: that a spicy drink which is also carbonated rings three bells at once, because dissolved CO_2 becomes carbonic acid at the mucosa, releasing protons that tingle the sour pathway and lightly provoke the pungency channel. A spicy kombucha is alive on the tongue in the most literal sense.



The tea plant — theanine for umami, catechins for astringency, and the leaf a kombucha culture turns to bright, structured sour.

Kombucha is sweet tea fermented by a consortium — the SCOBY, a symbiotic mat of yeast and acetic-acid bacteria. The yeast ferments the sugar to ethanol; the bacteria oxidize the ethanol to acetic and gluconic acids; the pH falls to around 2.8 to 3.5 and the drink turns bright, sour, and lightly effervescent. The tea underneath it matters, and our tea host — two hundred and eighty constructs across the umami, bitter-attenuation, and astringency axes — lets us shape it: a kombucha tea edited to hold its theanine (a quiet umami and a softness) while attenuating the harsh catechin bitterness, so the finished drink keeps its structuring astringency without the clenching bitterness that a long brew usually brings. Astringency, remember, is not a taste but a touch — tannins binding the salivary proteins (the PRPs of the A axis) and stripping their lubrication — and a good kombucha wants that grip to give it backbone.

The SCOBY itself deserves a word, because it is the clearest case in the book of a *consortium* doing what no single organism could. The mat is mostly a cellulose pellicle spun by acetic-acid bacteria — *Komagataeibacter* chief among them — and threaded through it live the osmotolerant wild yeasts, *Brettanomyces* and *Zygosaccharomyces* and their kin. The yeasts ferment the tea's sugar to ethanol and carbon dioxide; the bacteria oxidize that ethanol to acetic acid and the glucose to gluconic acid, a softer,

rounder sour than vinegar's sharp acetic. The two acids together give kombucha its particular brightness, and the division of labor is strict: take the yeasts away and the bacteria have no alcohol to work on; take the bacteria away and you are left with a faintly boozy sweet tea, not a kombucha. It is a relay — like the vinegar of Chapter 18, but caught at the moment both runners are still on the track — which is why a kombucha always carries a trace of alcohol, usually well under one percent, enough to name but not to intoxicate, the honest signature of a ferment still in motion. Chill it and the runners slow; the colder it sits, the longer it holds its balance before the acids win.

Spicy kombucha — *the second ferment, with a chili lift.*

```
SPICY KOMBUCHA · makes ~1 L · 7–10 day F1 + 2–4 day F2
brewed tea (engineered leaf)      1 L
sugar                             70 g
SCOBY + starter liquid            1 + 150 ml
— for the second ferment:
lacto-fermented chili mash        1–2 Tbsp
amazake or fruit (Ch. 9)          2 Tbsp (optional, for fizz)
```

1. Dissolve the sugar in the hot tea; cool to room temperature. Add the SCOBY and starter liquid, cover with cloth, and ferment 7–10 days until pleasantly sour (the first ferment, F1).
2. Remove the SCOBY. Bottle the kombucha with the fermented chili mash and a little amazake or fruit to feed a fresh burst of carbonation.
3. Seal and ferment 2–4 days at room temperature — burping once a day to check pressure — then refrigerate. The sealed second ferment (F2) builds the fizz; the chili brings heat already lifted by its own acid.

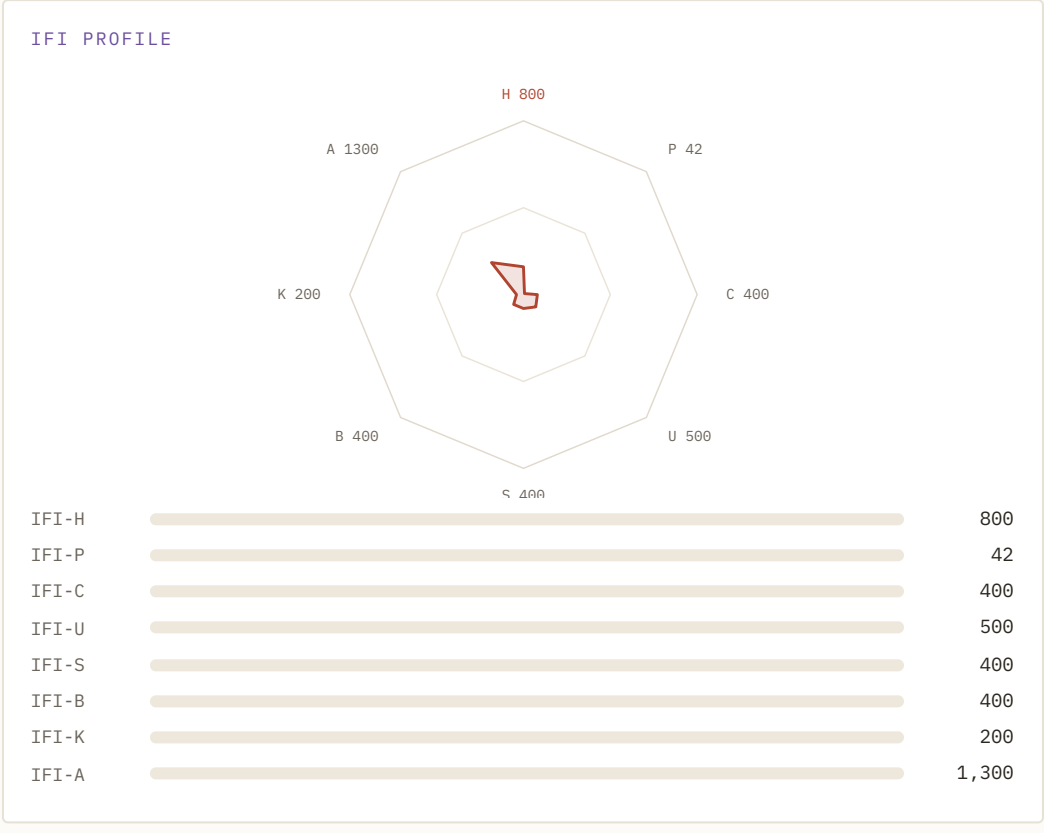

Two more drinks round out the shelf. **Tepache** is the easiest ferment in this whole book — pineapple peel, a dark unrefined sugar, a stick of cinnamon, water, and two to four days of wild fermentation into a cloudy, tart, lightly fizzy refresher that takes a dose of fermented chili gladly. And the **dilute hot-honey session** promised back in

Chapter 12 lives here: honey thinned past about nineteen percent moisture — above its preserving threshold at last — with chili and a wild or pitched yeast, fermented a few sparkling days into a short, sweet-hot, faintly alcoholic drink that is the honey's self-preservation finally, deliberately, undone.

Spicy tepache — *two days, pineapple skin, a chili finish.*

```
SPICY TEPACHE · makes ~2 L · 2–4 days
peel + core of 1 ripe pineapple
piloncillo or raw brown sugar    120 g
cinnamon stick                   1
water                            2 L
fermented chili mash             1–2 Tbsp (to finish)
```

1. Scrub the pineapple, keep the peel and core (eat the flesh). Dissolve the sugar in the water, add peel, core, and cinnamon.
2. Cover with cloth and ferment at room temperature 2–4 days, until cloudy, tart, and lightly fizzy — taste daily, as it will turn to vinegar if left too long.
3. Strain, stir in the fermented chili mash, and chill. For more fizz, bottle sealed a further day and burp it once.

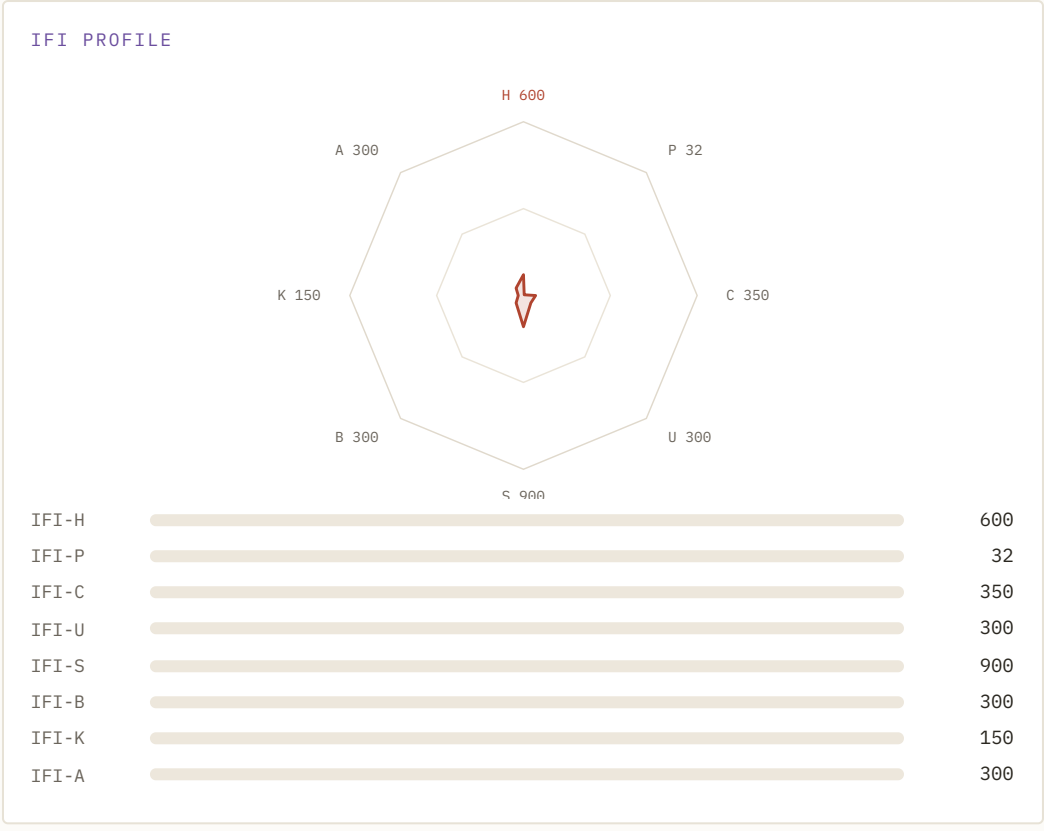

Each of them is a reminder that fermentation's range runs from a two-year miso to a two-day soda, and that all of it is the same act at a different speed: a living thing, handed a substrate, making the molecules a receptor is waiting to hear.

Chapter 16 · The three dials — *salt, temperature, and time decide everything.*

Every ferment in this book — the two-year miso and the two-day tepache, the alkaline natto and the sour mash — is tuned with the same three dials, and a fermenter who truly understands them can walk into any tradition on earth and read what it is doing and why. Salt, temperature, and time: three numbers, and between them they decide which organisms live, which enzymes run and how fast, how far up every axis the food climbs, and whether the result is safe. Before the last recipes and the measurement that governs them, it is worth setting the three dials out plainly, because everything that came before has been turning them without always naming them.



A petri dish — the whole craft in miniature: set the conditions, and the conditions choose the culture.

The first dial is salt, and salt is a selector. It does three things at once, and all three matter. It lowers water activity, the free-water fraction microbes need to grow, which excludes the fast spoilage organisms and clears the field for the salt-tolerant ones you want. It draws water out of plant tissue by osmosis, wilting a cabbage leaf pliable and firming a cucumber by plasmolysis — texture, made by physics. And it slows enzymes, which is to say it sets k, the rate constant of Chapter 1: a heavily salted miso ferments so slowly that only the patient, halotolerant koji enzymes and a few halophilic

organisms can work at all, over years, which is exactly why a long miso tastes of time. The salt concentration you choose is therefore a choice of *who ferments, how fast, and how safely*, all in one number.

SALT (% BY WEIGHT)	WATER ACTIVITY (APPROX.)	WHAT THRIVES	FERMENTS
0–2%	> 0.98	fast lactic — and spoilers	quick pickles (risky)
2–3%	~0.96	clean lactic bacteria	kimchi, kraut, hot-sauce mash
5–13%	0.90–0.94	halotolerant koji, halophiles	miso, soy sauce
> 15% / dried	< 0.85	almost nothing	preservation, no ferment

The second dial is temperature, and temperature is a throttle — on two engines at once. It sets enzyme rates, because every enzyme has an activity optimum: the amylases that sweeten amazake peak at 55–60 °C, the proteases that savor a miso run warmer than the amylases, and pushing a koji incubation from 30 toward 35 °C biases its whole secretion from sweet toward savory, as Chapter 4 showed. And it sets microbial succession, because each organism has a temperature it wins at — mesophilic lactic bacteria at a cool 18–24 °C, koji at 28–35 °C, *Rhizopus* for tempeh at 30–32 °C, *Bacillus* for natto at a hot 40–45 °C. As a rule of thumb the rate of a biological process roughly doubles for every 10 °C (a Q_{10} near two to three), so ten degrees is not a nuance; it can halve a ferment's time. Warm ferments fast but shallow and riskier; cool ferments slow but clean and aromatically complex, because the slow path lets side reactions and aroma precursors accumulate. Temperature is where a fermenter trades speed against depth.

FERMENT	ORGANISM	TEMPERATURE	REDOX
Koji growth	*A. oryzae*	28–35 °C	aerobic
Amazake	koji enzymes	55–60 °C	– (enzymatic)
Lactic (kimchi, mash)	*Leuconostoc* → *Lactobacillus*	18–24 °C	anaerobic
Tempeh	*Rhizopus*	30–32 °C	aerobic (surface)
Natto	*Bacillus subtilis*	40–45 °C	aerobic
Kombucha	SCOBY	24–28 °C	aerobic
Miso aging	koji enzymes + halophiles	15–25 °C	anaerobic (packed)

The third dial is time, and riding with it, oxygen. Time is simply how far up the curves you let the ferment climb — the integral under the Glu accumulation and the pH fall. Stop a miso at two weeks and you have a sweet young white; let it run two years and you reach the bottom of savor. But time only means what the *redox* condition lets it mean, which is why oxygen is the third dial's silent partner. Submerge a mash and you create the anaerobic world lactic bacteria rule, where acid builds and spoilers suffocate — submergence is the master safety move of the whole craft. Leave a surface open to air and you hand the ferment to aerobic organisms instead: koji needs oxygen to grow its mycelium, the acetic-acid bacteria of the next

chapter need it to make vinegar, *Bacillus* and *Rhizopus* work at an aerated surface. The single decision *under air or under brine* changes which kingdom of organism inherits your substrate.

Read the three dials together and you have a grammar, not a set of rules. Salt chooses the players and sets the safety floor; temperature throttles the enzymes and picks the winner of the succession; time and oxygen decide how far the chosen players carry it and in which direction. Every axis on the Index moves as you turn them: more salt, slower umami climb but a higher clean ceiling; warmer, more sweetness from the amylases or more savor from the proteases depending on the window; longer, more of everything and more Maillard roundness on the bitter-attenuation axis. The traditions of the world are, every one of them, a particular setting of these three dials, discovered by patience.

To watch the dials compose a flavor, run one substrate two ways. Take the same steamed rice and the same koji spores, and incubate one tray cool, near 30 °C, the other warm, near 35 °C. The cool tray favors the amylases; it comes out sweeter, an amazake-leaning koji that will build a pale, sweet shiro miso. The warm tray favors the proteases; it comes out savorier, a koji that will build a darker, deeper one. Same organism, same grain — one dial moved a few degrees, and two different foods.

IFI PROFILE

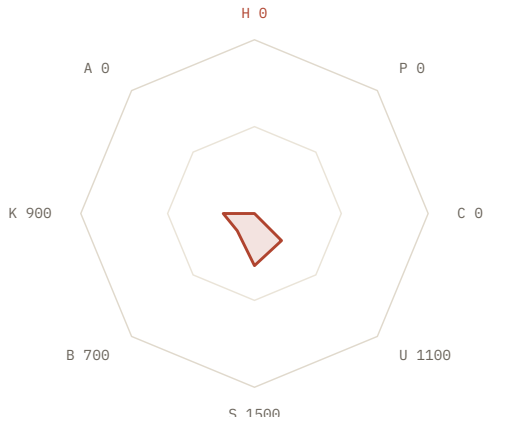

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	1,100
IFI-S	1,500
IFI-B	700
IFI-K	900
IFI-A	0

IFI PROFILE

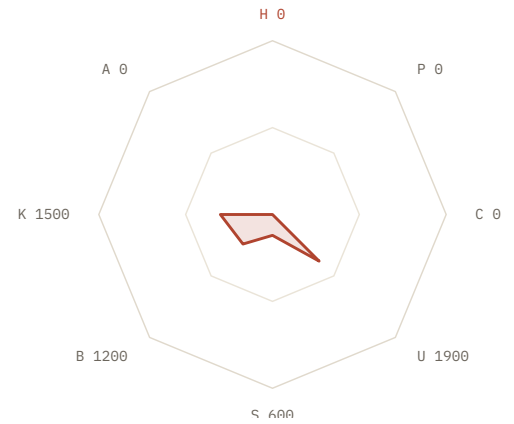

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	1,900
IFI-S	600
IFI-B	1,200
IFI-K	1,500
IFI-A	0

Read the two scorecards side by side — the cool koji above, the warm one below — and the lesson is on the page: the sweetness and the savor have traded places, and nothing changed but the temperature. Our contribution in the lab is not a new dial — there are only three — but the ability to hold all three to a tolerance the hand cannot, and to start them from an engineered substrate, so that the same setting gives the same flavor every time. The dials are ancient. The steadiness of the hand on them is what we add.

Chapter 17 · The alkaline ferments — *natto, tempeh, and the bean made strange.*

Almost everything so far has run toward acid, because acid is where heat and safety and brightness all live. But two of the most remarkable soy ferments on earth run the other way — natto climbing *alkaline*, tempeh holding near neutral — and they are worth a full chapter not only because they are nutritional powerhouses but because they prove the breadth of the thesis: fermentation is the manufacture of receptor-active molecules whichever direction the pH goes. These two take the humble bean of Chapter 5 and make it genuinely strange, and in doing so they make it more fully food.

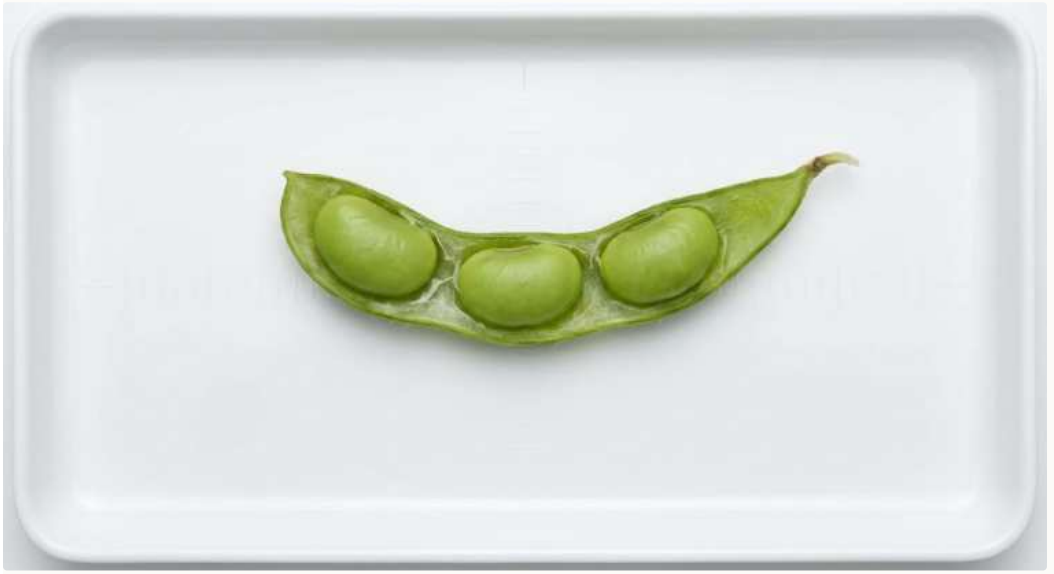

The soybean in cross-section — the same warehouse, handed to *Bacillus* or *Rhizopus* instead of koji, becomes two foods koji could never make.

Natto is the work of a single bacterium, *Bacillus subtilis* var. *natto*, and it is a ferment that runs uphill in pH. The *Bacillus* proteases are powerful and relentless, cutting the bean's storage protein so hard that they release not only free glutamate but ammonia, which drives the pH *up*, to 7.5 or beyond — the sharp, penetrating aroma of natto is that ammonia and a set of nutty pyrazines the bacterium throws off. Three products make natto extraordinary. The free glutamate gives it a deep, pungent umami. The famous sticky threads are **poly-γ-glutamic acid** — glutamate polymerized through its side-chain into long chains, which is a lovely closing of a circle: the loudest umami molecule, strung into a polymer so thoroughly bound that it has no taste at all until something hydrolyzes it, a reminder one last time that only *free* glutamate is heard. And *Bacillus* lays down **vitamin K2** as long-chain menaquinone — natto is the richest dietary source there is — along with the fibrinolytic enzyme nattokinase, which we name as a compound the ferment makes and about which we promise nothing. It is a food that is savory, aromatic, strange, and dense with things the raw bean never held.

Natto — *the uphill ferment.*

```
NATTO · makes ~600 g · 24 h + age
small soybeans           300 g
natto starter (or 1 Tbsp live natto)   per instructions
— incubate 40–45 °C, 22–24 h, aerobic and humid
```

1. Soak the small beans 12 hours, then steam or pressure-cook until a bean crushes with no resistance — softness matters, as *Bacillus* works the surface.
2. While still hot, stir through the starter (dissolved in a spoon of just-boiled, cooled water, or a spoonful of live natto). Spread shallow in a tray — it needs air.
3. Cover loosely to hold humidity while admitting oxygen, and hold at 40–45 °C for 22–24 hours, until the surface turns filmy and a stirred spoon pulls long threads.
4. Refrigerate at least a day; the flavor deepens and the strings strengthen. Serve with mustard and scallion — or, for this book, a spoon of living chili crisp, the heat meeting the ammonia-savor head on.

IFI PROFILE



IFI-H	0
IFI-P	84
IFI-C	0
IFI-U	2,000
IFI-S	200
IFI-B	400
IFI-K	1,500
IFI-A	100

Tempeh is the gentler cousin, the work of a mold — *Rhizopus oligosporus* — that does something no other ferment here does: it binds the loose beans into a solid, sliceable cake, knitting them with a dense white mycelium you can cut like a cutlet. It ferments near neutral and aerobically, at the surface, and its chemistry is milder than miso's or natto's: a partial proteolysis that lifts free amino acids modestly and leaves a clean, mushroomy, nutty flavor. Its gifts are nutritional and structural. Its phytase cleaves the bean's phytate, freeing the iron and zinc the raw bean chelated away; it raises B-

vitamins; it deglycosylates the isoflavones into their absorbable forms; and traditional tempeh, carrying bacterial co-cultures alongside the *Rhizopus*, is one of the few plant foods to hold any vitamin B12 at all — a fact we state as the chemistry it is, without a word of health claim. Above all tempeh is a *canvas*: neutral, firm, and savory enough to carry whatever you bring it.

Tempeh — *the mold that knits a cake.*

```
TEMPEH · makes ~500 g cake · 24–48 h
dehulled soybeans        300 g
vinegar                  1 Tbsp
tempeh starter (Rhizopus) 1 tsp
— incubate 30–32 °C, 24–48 h, in perforated bags
```

1. Soak and cook the dehulled beans until just tender, not mushy; drain and dry the surface thoroughly — a wet surface smothers the mold, a dry one lets it breathe and knit.
2. Toss the warm beans with the vinegar (its mild acidity favors *Rhizopus* over competitors) and the starter.
3. Pack about two centimetres thick into perforated bags or a banana leaf. Incubate at 30–32 °C for 24–48 hours, until bound into a firm, white, mycelium-laced cake that smells of mushroom and nut.
4. Use fresh, or steam and freeze to stop it. For this book: marinate in shio koji and a spoon of fermented chili, then sear hard — the Maillard crust, the koji-deposited glutamate, and the lifted heat all at once.

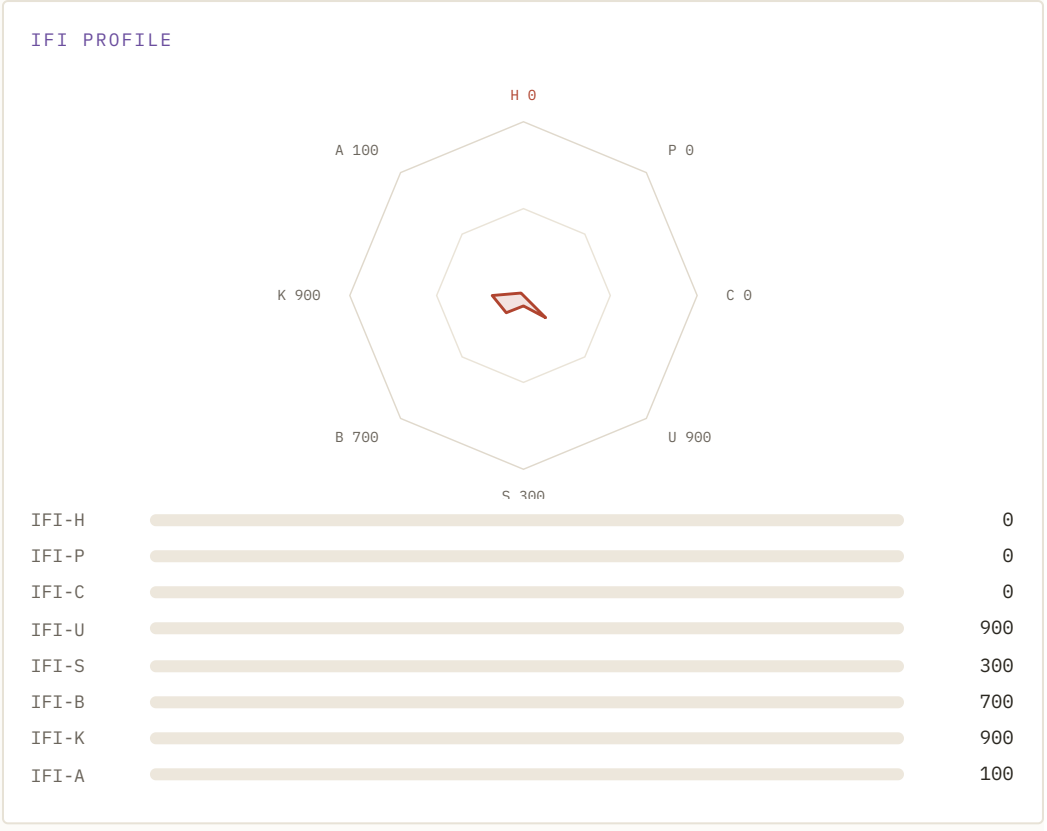

One honest contrast closes the chapter, and it is the contrast that makes the title sharp. Because natto and tempeh do not acidify — natto climbs the other way, tempeh barely moves — they do *not* self-lift heat the way a lactic ferment does. There is no falling pH here to sensitize TRPV1, so a chili added to natto or tempeh reads at its plain strength, not amplified. These ferments bring the savory and the structural and the nutritional; they do not bring the proton multiplier. Which is exactly why the fermenter's palette needs both directions: pair an alkaline or neutral ferment's depth with an *acid* ferment's lift

— natto under chili crisp whose flake was lacto-soured first, tempeh glazed with a mash hot sauce — and you get the savor of the one and the amplified heat of the other. Two pH directions, composed on one plate. Fermentation is not one road to flavor. It is every road, and a cook who knows which way each ferment moves the pH knows which ones lift heat and which ones only deepen it.

Chapter 18 · Vinegar — *the other sour, and the heat it also lifts.*

Lactic acid is not the only road down to a low pH, and this chapter is here to generalize the engine of the whole book. If the thesis of Chapter 2 is true — that falling pH lifts heat at the receptor — then it should hold for *any* acid a ferment can make, not just the lactic acid of a kimchi. It does, and the proof is the oldest condiment in the world. Vinegar is a second fermentation stacked on a first, and the acid it makes is sharper, more volatile, and in its own way even better at lifting a chili's burn.



Sugarcane — a river of fermentable sugar. Yeast turns it to alcohol; the acetic bacteria turn the alcohol to the sharp sour of vinegar.

Vinegar is a relay between two organisms that need opposite things. First, in the absence of air, yeast ferments sugar to ethanol — the anaerobic step that also makes wine and beer. Then, in the *presence* of air, acetic-acid bacteria — *Acetobacter* and *Komagataeibacter* — oxidize that ethanol in two moves, through acetaldehyde to acetic acid, a strictly aerobic step that cannot happen until the yeast has made its alcohol and air has reached the surface. The slippery, translucent mat that forms on a living vinegar, the "mother," is a cellulose pellicle the *Komagataeibacter* spins as it works — close kin to the kombucha SCOBY of Chapter 15, which is, after all, simply a tea vinegar caught partway. The two-organism relay is why vinegar was humanity's accidental discovery: leave any sweet thing open to the air long enough and the yeast and the bacteria arrive in order.

The family of vinegars is as wide as the family of sweet things, and each carries its origin. Wine vinegars keep the grape's structure; cider vinegar the apple's sweetness; rice vinegar the clean softness that suits a ferment-heavy cuisine; malt vinegar the roast of barley. The aged ones stack a third chemistry on top of the acetic — a traditional balsamic is grape must cooked down and aged through a cascade of casks for years, Maillard browning and slow concentration building a sweet-sour syrup that reads across the S, A, and U axes at once, and a sherry vinegar carries the oxidative depth of its solera. For a fermenter chasing heat, the lesson is that a vinegar is never only acid: it is acid plus the ghost of whatever sugar

it came from, and that ghost frames the chili the way the aroma of Chapter 19 frames the burn. A fruit-scrap vinegar made from the trimmings of the very peppers you fermented is the tidiest version of the whole idea — the chili's own sugars turned into the acid that will lift the chili's own heat, nothing wasted, the loop closed in a single jar.

Acetic acid, pK_a 4.76, carries a different edge than lactic acid, and the difference is all to a chili's advantage. It is volatile, so it reaches the nose as a sharp pungency before it even lands on the tongue; it drops the pH as surely as lactic acid, sensitizing TRPV1 by the exact proton mechanism of Chapter 2; and its own molecule brings a bright, penetrating sourness that reads as *sharp* alongside the heat. A chili vinegar is therefore heat lifted along two paths at once — the low pH sensitizing the burn, and the acetic sharpness framing it — which is the quiet reason that the great majority of the world's hot sauces are finished with vinegar and not only for its preservation. The vinegar below pH 3.5 keeps the sauce safe, as Chapter 14 noted; but it also, in the same stroke, makes the sauce read hotter. Preservation and amplification, once again, on one falling number.

Fermented chili vinegar — *two organisms, one sharp heat.*

```
FERMENTED CHILI VINEGAR · makes ~1 L · 4–8 wk
ripe chilies (engineered), chopped   300 g
raw cane sugar                       150 g
water                                1 L
wine yeast (or a wild pinch)         small pinch
unpasteurized vinegar / a mother     2 Tbsp
```

1. Dissolve the sugar in the water, add the chilies and yeast, and cover with cloth. Ferment about a week, loosely covered, until it smells winey — the yeast has made alcohol from the sugar.
2. Strain out the solids, stir in the vinegar mother, and move it to a wide jar covered only with cloth, giving it maximum air. Over 3 to 6 weeks the acetic bacteria oxidize the alcohol to acetic acid; it will sharpen, clear, and sour.
3. When it is bracingly sour and stable below pH 3.5, bottle it. It will keep indefinitely — it is, by now, too sour for anything to live in but itself.

IFI PROFILE

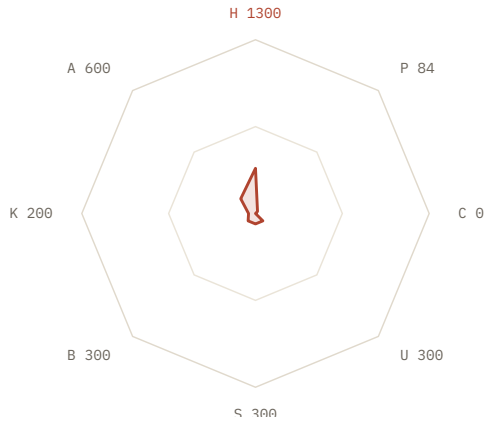

IFI-H	1,300
IFI-P	84
IFI-C	0
IFI-U	300
IFI-S	300
IFI-B	300
IFI-K	200
IFI-A	600

The engineering here is mostly upstream — the chili grown to a locked heat, and our sugarcane host, two hundred constructs on the sweetness axis, as a clean river of the fermentable sugar the yeast needs for its alcohol. But the point of the chapter is less a product than a principle: the heat-lift of fermentation is not a quirk of lactic acid. It is a property of protons, and every ferment that makes an acid — lactic in the mash, acetic in the vinegar, gluconic and acetic in the kombucha, even the carbonic acid of a fizzy drink — turns the same dial on the same receptor. The book's title names a coincidence that turns out to be a law. Sour heat is louder heat, by whichever sour you reach it.

Chapter 19 · The scent of the work — *the dimension the eight axes leave to the nose.*

We have spent this whole book putting numbers on flavor, and it would be a kind of dishonesty to end the craft without admitting what the numbers do not cover. The Iyengar Flavor Index measures eight dimensions — the five tastes the tongue transduces and the chemesthetic channels the trigeminal nerve reports, heat and cooling and pungency. That is taste, and taste is only a fraction of flavor. The rest, the larger part, is *aroma*: volatile molecules rising to the olfactory epithelium, read not by eight receptors but by hundreds, in combinations too vast and too personal to reduce to a single score. We do not put a number on aroma, and we want to say clearly why: not because it does not matter, but because it is the one dimension of flavor that is genuinely combinatorial rather than axial, and a number would lie about it. Fermentation, of all the kitchen crafts, is the one that lives most in this unscored dimension — it is where a ferment most exceeds its raw ingredient — so a book on fermentation must honor the nose even though the Index keeps its silence there.



Dawn on the glass — the hour a ferment smells of most, before the heat of the day burns its volatiles off.

Name the aromas a ferment builds, because they are specific chemistry, not atmosphere. **Esters** — isoamyl acetate reading as banana, ethyl hexanoate as apple — are built by yeast joining alcohols to acids, and they are the fruity lift of a young kombucha or a lively hot sauce. **Higher alcohols** come from yeast catabolizing amino acids through the Ehrlich pathway, warm and winey in trace, harsh in excess. **Pyrazines** — nutty, roasty, cocoa-like — are thrown off by *Bacillus* and built by Maillard chemistry, the smell of natto and of a miso's darkened crust. The **furanones** HEMF and HDMF,

caramel-sweet and savory at once, are the signature perfume of a great soy sauce, born of Maillard reactions and the halophilic yeast together. **Lactones** bring coconut and peach; **Strecker aldehydes** from the Maillard pathway bring malt and honey and the potato note of methional; **sulfur volatiles** from the alliums and from methionine bring the deep savory funk that is magnificent in trace and ruinous unchecked. A ferment's aroma is a chord struck across all of these, and it is why two misos with identical amino-acid nitrogen — identical, that is, on the U axis — can smell like two different worlds.

Here is the part that matters for a book obsessed with the eight axes: aroma is not scored, but it *shapes how the axes are experienced*. Olfaction and taste are bound together in the brain so tightly that a "sweet" aroma makes a food taste measurably sweeter, a smoky aroma reads as savory, a fruity ester can make a sour thing seem riper and rounder. Most of this arrives retronasally — the volatiles rising up the back of the throat as you chew and swallow, which is why pinching your nose flattens a ferment to a shadow of itself, and why a head cold turns the most complex miso to mere salt.

The distinction between smelling and tasting is really a distinction between two routes to one organ. Orthonasal olfaction is what you do when you lift a jar of miso to your face and inhale — air drawn in through the nostrils, the anticipation of a food before it enters. Retronasal olfaction is what happens once it is in your mouth: warmed and chewed, its volatiles driven up the back passage from throat to nasal cavity, where the same receptors read them as part of what the brain fuses into the single sensation it calls flavor. That fusion is so complete that most people cannot tell, from the inside, where taste ends and smell begins — which is the deepest reason the Index does not try to score aroma. Taste is a small, fixed alphabet: a handful of receptors, each with a clean axis and a magnitude. Smell is a combinatorial code of hundreds of receptors firing in patterns, and a pattern is not a magnitude — you cannot put one number on a chord the way you can on a note. We could have forced a score, and it would have been a lie about the kind of thing aroma is. Better to measure the eight axes honestly and name the ninth dimension for exactly what it is: vast, combinatorial, and the place a ferment does its most uncountable work. So even as we design to the eight axes, we design the aroma alongside them, knowing it is not on the scorecard: the lipoxygenase knockout in our soybean is an aroma edit, removing the beany off-note so the savor reads clean; the tomato tuned for its signature volatiles is an aroma edit; and every choice of yeast strain and fermentation temperature in these recipes is, as much as anything, a decision about scent.

Even heat, the most purely tactile thing in the book, wears a coat of aroma. Capsaicin itself is odorless — it is pure trigeminal signal, no smell at all — but a chili is never only capsaicin. Its fresh green pyrazines, its fruity esters, the roasted notes it gains when toasted, the lactic and estery aromas it earns in a ferment: all of these frame the burn and tell you, before the heat even arrives, what *kind* of heat is coming. A fermented chili smells rounder, and so its heat is received as rounder, though the capsaicin is unchanged. This is the

humility the chapter is here to offer. The eight axes are the skeleton of flavor — designed, quantified, reproducible, the part we can engineer to a number and hold to a tolerance. Aroma is the living flesh the culture lays over that skeleton, hundreds of volatiles deep, and it is the part we respect most precisely because we measure it least. We built an index to end the silence around taste. We leave the silence around smell on purpose, because some things are truer kept whole.



PART IV OF IV

Control & Consequence

The rigor the living work demands — reading the ferment as it climbs · the safety floor · what the culture gives back

Chapter 20 · Reading the ferment — *loudness climbs on a curve you can measure, and aim.*

A ferment is a bioreactor that happens to live on a counter, and the discipline that lets us sell a living product at a coefficient of variation of seven percent or below is the discipline of treating it as one: measuring it as it runs, modeling where it is going, and steering it to a number instead of waiting to see where it lands. The romance of fermentation is real and we feel it as much as anyone, but the romance does not survive a bad batch, and the only thing standing between a crock and a bad batch is measurement. This chapter is the rigor under all the recipes — how you read a ferment climbing, and how you aim it.

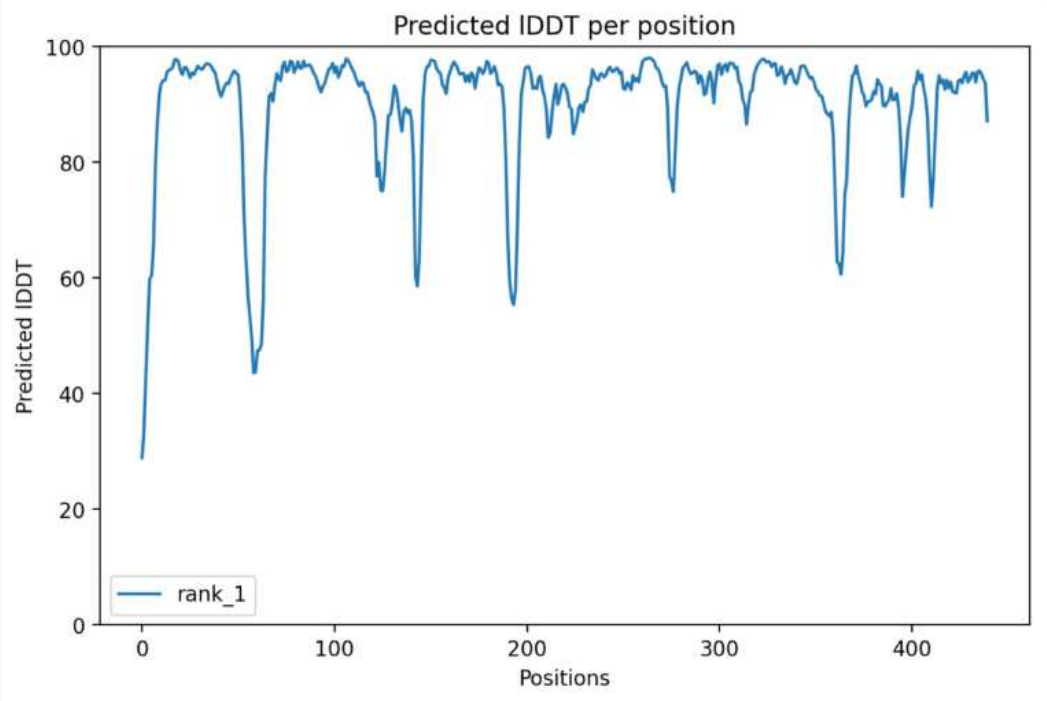

The confidence plot — every enzyme we edit is modeled before it is made, and every ferment it drives is measured as it runs.

Each axis has a proxy you can measure in a sample, and the fermenter's instrument panel is just a handful of them read over time. For the umami axis, **amino-acid nitrogen** — the dissolved free-amino-acid load — tracked by formol titration, with free glutamate itself by enzymatic assay or chromatography when you need the exact figure; this is the soy-sauce grading metric from Chapter 1,

and it is the direct readout of Glu. For the heat axis, **capsaicin by chromatography** to fix the capsaicinoid content, read together with **pH**, because the effective IFI-H is the capsaicin content *lifted by the proton multiplier* — you cannot know how hot a mash will read without knowing how sour it has become. Then the governing variables that set every axis's trajectory: **pH** and **titratable acidity** for the acid climb, **°Brix** for the sugar the culture is eating, **salt** by refractometry or conductivity, and **water activity** for the safety margin we will come to next.

MEASURE	PROXIES	METHOD	WHAT IT TELLS YOU
Amino-acid nitrogen	IFI-U	formol titration	how far proteolysis has run
Free glutamate	IFI-U	enzymatic / HPLC	the exact umami ligand load
Capsaicin + pH	IFI-H (effective)	HPLC + probe	heat, and its proton lift
pH / titratable acidity	the acid climb	probe / titration	ferment progress and safety
°Brix	IFI-S, culture fuel	refractometer	sugar made and sugar eaten
Water activity	safety margin	a_w meter	whether a pathogen *could* grow

The model from Chapter 1 turns those readings into a forecast. Because free glutamate climbs as Glu = Glu_max·(1 − e^(−kt)), you can estimate k from the early slope — the first several days of amino-nitrogen readings — and then *predict the endpoint* before you reach it, which is the whole game in process control. A ferment that is climbing too slowly can be nudged warmer; one racing toward an under-salted finish can be cooled or cut. You are not waiting to taste the result in two years. You are reading the curve in the first week and steering it.

DAY	PH	FREE GLUTAMATE (MG/100 G)	EFFECTIVE IFI-U	EFFECTIVE IFI-H
0	6.2	~20	~150	1.0× base
7	4.8	~180	~900	~1.8×
21	3.8	~340	~1,500	~2.6×
180	3.6	~560	~2,200	~2.7×

How, then, do we hold something alive to a tight number? By locking the four things that make it vary, in order. We **lock the substrate** — an engineered cultivar whose capsaicin or glutamine load is already held at CV ≤7% in the field, so the ferment starts from a tight distribution rather than a wild one. We **lock the starter** — defined cultures pitched at a known load, our engineered koji at a counted spore density and selected lactic strains, instead of trusting whatever rode in on the air. We **lock the trajectory** — controlled temperature, and endpoints defined by *reaching a pH* rather than by a fixed number of days, so each batch ferments to the same chemical state even if it takes a day longer. And finally we **blend to spec** — the vatting trick the whisky world has used for a century, combining batches to land the finished product exactly on its target IFI. Four locks, and the living variability that makes fermentation feel unrepeatable becomes a process you can specify. None of it takes the life out; it only takes the lottery out.

The blend is more than a metaphor; it is arithmetic a formulator runs by hand. If one barrel finishes at an effective IFI-H of 4,800 and another, a touch under-fermented, at 4,200, then a blend weighted sixty to forty lands at 4,560, and a formulator chasing a 4,500 target simply solves for the ratio — a weighted average, batch by batch, until the bottle is exactly where the label promises. Variance that cannot be prevented upstream is averaged out downstream, which is how even a living product carries a number you can trust. And the whole process is watched the way any bioreactor is watched, with the statistics of control: each batch's readings plotted against the running mean and its limits, so a point drifting toward the edge is caught as a trend before it becomes a failure. A ferment beginning to climb too slowly, or to acidify too fast, announces itself on the chart days before it would announce itself on the tongue.

The measurement itself is humble bench chemistry, not a mystery. Formol titration — adding formaldehyde to tie up the amino groups so the freed carboxyls can be titrated — turns the abstract "amino-acid nitrogen" into a figure anyone can read with a burette and a pH meter; a refractometer reads Brix off a single drop; a probe reads pH, and with it the proton multiplier on the heat. The instruments are ordinary. The discipline of using them, every batch, on a thing that is alive and wants to vary, is the whole of what it means to hold a ferment to a standard.

This is also where the company's commitment to being checkable lives. The method behind every engineered enzyme in this book — three patents filed, twenty-one peer-reviewed papers posted openly under a Creative Commons license — is documented so that the reading can be repeated by anyone who doubts it. A ferment you can measure is a ferment you can trust, and a method you can read is a company you can check. We think those are the same virtue, pointed at the two halves of the work.

Chapter 21 · The safety floor and what the ferment gives back — *the two numbers that keep it safe, and the gift underneath.*

We have spent a whole book celebrating a process that is, stripped of romance, the controlled growth of microorganisms in food — and that is exactly why it must end in rigor and in care, because the same conditions that make a ferment delicious are a hair's breadth from the conditions that make one dangerous, and the line between them is not luck. It is two numbers. Every safe ferment on earth sits below at least one of them, and the great ones sit below both.



The family at dusk, in the greenhouse where it started. Everything in this book is still the same intention: to feed the body right.

The first number is **pH 4.6**. Below it, *Clostridium botulinum* — the organism behind the most feared hazard in preserved food — cannot grow and cannot make its toxin. This is not folk wisdom; it is the line that food regulators draw around acidified and fermented foods, the threshold a bottled hot sauce or a jar of kimchi must be held below to be safe at room temperature. Every lactic ferment in this book drives *toward* and well *past* that line, to pH 4.2 and below, and the fermented hot sauces finish below 3.5 — which is the quiet miracle this whole title rests on, because the falling pH that lifts the heat at the receptor is the very same falling pH that locks out the pathogen. Safety and flavor are not in tension here. They ride the same axis, in the same direction. The sourer it gets, the hotter it reads *and* the safer it is.

The second number is **water activity 0.85**. Below it, *C. botulinum* and most pathogens cannot grow no matter the pH — it is why honey (near 0.6) and a heavily salted miso and a dried flake are self-preserving. Water activity is the barrier behind salt and behind drying, and it explains the honesty of Chapter 12: honey will not ferment because its water activity is too low for anything to live, and to make it live you must raise that number past where the barrier holds. The craft calls the combination of these barriers *hurdle technology* — you do not need any single barrier to do the whole job

if you stack several: a falling pH, a lowered water activity, a dose of salt, and a vigorous starter culture that simply outcompetes everything else for the substrate. A well-built ferment is a stack of hurdles, each modest, together absolute.

HURDLE	THRESHOLD	KEEPS OUT	SET BY
Low pH	< 4.6 (we reach < 4.2)	*C. botulinum*, most pathogens	`lactic culture`
Water activity	< 0.85	*C. botulinum*, molds	`salt, drying`
Salt	2–13%	spoilage flora	`the recipe`
Competitive culture	high pitch	everything slower	`defined starter`

The honest hazards get honest answers. Surface films — a white, wrinkled *Kahm* yeast — are harmless and cosmetic; fuzzy, colored mold is not, and the defense against it is mechanical and total: keep the ferment submerged, keep the vessel clean, keep oxygen off the surface. Mycotoxins are a risk only of *wild* molds, which is precisely why we ferment with a defined, non-toxigenic koji whose entire genome we have verified, rather than whatever spore drifts in. Histamine, the hazard of fish ferments, is controlled by temperature and freshness. None of this is cause for fear; all of it is cause for respect. Fermentation has fed humanity safely for ten thousand years because the people doing it honored these lines by instinct. We honor them by instrument, and arrive at the same safety.

And then, underneath the safety, the gift — which we will describe as chemistry and intention and never as a cure, because that is the line our family holds. A ferment gives back more than it takes, and the giving is specific and real. Proteins arrive already cut to peptides and free amino acids the body can take up more readily than intact storage protein. Phytate, the mineral-binder, is cleaved by microbial phytase, so the iron and zinc the raw bean locked away come free. Isoflavones shed their sugars into the aglycone forms the body absorbs more easily. *Rhizopus* and *Bacillus* lay down B-vitamins that were never in the raw bean; natto's *Bacillus* makes vitamin K2 and the folate climbs; the oligosaccharides that trouble the gut are eaten by the culture before they reach you. And the food arrives carrying the living culture itself. We state these as the chemistry they are, make no promise about anyone's health, and let them stand — because they are, quietly, the whole reason our family came to revere the crock in the first place: it makes food more fully food, in the same motion that it makes food delicious.

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

That is the convergence this book set out to show, arrived at from both ends. To engineer a plant is to raise the concentration of the molecule a receptor is waiting for. To ferment a food is to raise the concentration of the molecule a receptor is waiting for. They are the same act, upstream and down, and when you run both on one batch — an engineered bean under an engineered koji, a flagship pepper down to pH 3.5, a loaded tomato dusted with a loaded mushroom — the loudness you can reach, and the nourishment you free, exceed what either hand could do alone.

IFI PROFILE

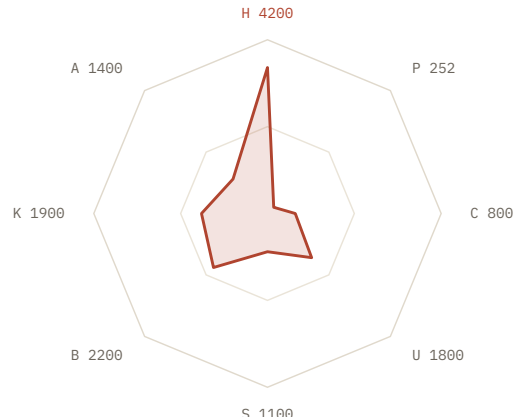

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

We close on Phoenix, SS-0096 — the one cultivar we built to register on all eight axes at once, the full-spectrum proof that flavor can be designed whole. Read it, and then imagine the second hand arriving: ferment it, and the acid lifts the H still higher at the receptor, the koji adds to the U and the K, the long Maillard rounds the B, and even this, the loudest designed thing we have made, becomes louder and rounder and more alive. The construct gives the plant its eight voices. The crock teaches them to sing longer. Two technologies, one act, one patient reverence for the living thing — which is where we began, at a table, over food the family grew, learning to read a silence we can now, at last, write.

Grow it right. Ferment it clean. Feed the body what it was built to hear.