

THE SCOVILLE SPLICE LIBRARY · VOLUME LXI · THE ANNUAL REPORT

The State of Flavor

The year in flavor, counted. *The first reading of many.*

An annual report from inside the index — the state of all eight axes, forty-two organisms, and the numbers that hold them, read for the first time as one instrument. What we grew this year, what it scored, and what it tells us is coming.

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

Contents

01 Opening essay · The first reading

02 Chapter 1 · The Index as Instrument — a number is a form of respect.

03 Chapter 2 · The Eight Dials — what the tongue is actually doing.

04 Chapter 3 · The Arithmetic of Loudness — potency, not quantity.

05 Chapter 4 · Heat — the axis everyone knows.

06 Chapter 5 · Sweetness — sugar's argument, answered.

07 Chapter 6 · Umami — the depth of being cared for.

08 Chapter 7 · Zero-Bitter — subtraction as design.

09 Chapter 8 · Kokumi — the axis you cannot name.

10 Chapter 9 · Cool-Burn — temperature without weather.

11 Chapter 10 · Pungency — the sharp channel.

12 Chapter 11 · Astringency — the finish, and the grip.

13 Chapter 12 · Forty-Two Hosts — the roster this year.

14 Chapter 13 · The Flagships — ten warnings and a verdict.

15 Chapter 14 · The Tomato, Restored — the proof the method travels.

16 Chapter 15 · From Sequence to Flavor — design, not selection.

17 Chapter 16 · Precision as a Promise — ≤7%, every batch.

18 Chapter 17 · The Sixty-Four — what could be built.

19 Chapter 18 · What We Will Not Rush — the tier we have not named.

20 Chapter 19 · The Numbers, as They Stand — the locked count.

21 Chapter 20 · The State of Flavor — the year ahead.

Opening essay · The first reading

There is a night our family still returns to. Someone had carried in the first fruit of a plant we had spent a season arguing over — whether it would taste of anything at all — and we cut it at the table under the one good lamp, and it was *better* than the argument. Warmer. Rounder. It tasted the way we had hoped and could not promise. And then someone asked the question that this whole company is an answer to: *how much better?* Not whether. By how much.

For ten thousand years agriculture could not answer that question. It could grow a thing, taste it, and keep the seed of the one that moved you — a method of extraordinary patience and no arithmetic. You could say *this tomato is sweeter than that one.* You could not say by what factor, on what axis, against what fixed point, in a way a stranger could check. Flavor had a vocabulary of poetry and a vocabulary of chemistry and nothing in between that an orchard, a laboratory, and a buyer could all read on one line.

The Iyengar Flavor Index closed that gap. Each of the eight axes is a single number pinned to the half-maximal activation of a human receptor — not how much of a compound a food contains, but how *loudly* that compound speaks to the cell that hears it. Once you have that number, you have something you have never had before: you can take the state of flavor the way a weather service takes the state of the sky. You can read the whole field at once, write it down, and come back a year later to see what moved.



The founders walking the research greenhouse at the hour the light arrives — where every year of this report begins and ends.

This is that reading. It is the first of what we intend to publish every year: *The State of Flavor*, an annual report assembled from our own index data — every axis, every host organism, every flagship cultivar, the mathematics underneath, the method that writes it, and the frontier we have mapped but not yet built. We wrote it the way we were taught to measure: nothing claimed that we cannot show, nothing shown that we did not grow, and every lush sentence made to earn its place by leading somewhere exact.

A word on what this document is and is not. It is a census, not a catalog — it sells nothing, and where a number appears it is the number of record, locked, the same here as on every page of this library. It is written as *we*, because the work is a family's and the table is a family's, even where the computing was one pair of hands. And underneath all of it runs the thing we care about most and say least: that food grown with this much attention can be grown to hold the most nourishment it can carry, and that designing flavor and designing nourishment have turned out to be, more often than not, the same act. We will not preach it. It is simply the reason the lamp is still on.

Here, then, is the year — read in full for the first time.



PART I OF VI

The Instrument

IFI = 1,000 × EC_{50}(ref) ÷ EC_{50}(candidate) · the number, and what it measures

Chapter 1 · The Index as Instrument — *a number is a form of respect.*

Begin with the fruit, because the index does. A pepper comes off the plant still holding the day's heat, and inside it — in the placental tissue along the ribs, not the flesh or the seeds, as every chilihead learns the hard way — sits a store of capsaicinoids that the plant made to defend itself and that we have learned to tune like a dial. You can weigh those compounds. You can run them on a column and get micrograms per gram. But weight is the wrong question, and this is the first thing the index teaches: the tongue does not report *quantity*. It reports *potency*. It reports how readily a molecule throws the switch on the receptor built to feel it.

So the index is anchored not to concentration but to a receptor event. Every axis is pinned to an **EC_{50}** — the concentration at which a stimulus drives its receptor to half of full response. EC_{50} is a property of the lock-and-key fit, not of the dose: a molecule with a low EC_{50} is *loud*, throwing the switch at vanishing concentration; a molecule with a high EC_{50} is *quiet*, needing a flood to be heard. Capsaicin against the heat receptor TRPV1 is loud. The sugar glucose against the sweet receptor is comparatively quiet, which is the entire reason a few milligrams of a designed sweet protein can do the work of a spoonful of sugar.

The receptor's behaviour follows a curve the whole of pharmacology is built on — the Hill equation, which relates response to the concentration of an agonist:

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

Here R is the fraction of receptors activated, [A] the agonist concentration, R_max the ceiling, and *n* the Hill coefficient, which measures how steeply the response turns on once it begins — how much of a cliff the sensation sits on. Set [A] equal to EC_{50} and the equation collapses to R = R_max ÷ 2: half-activation, by definition. That single point — the midpoint of the sigmoid, the hinge of the curve — is the most honest one-number summary a dose-response relationship has. It is the number the index stands on.

From it, the whole scale falls out in one line:

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

Read it slowly, because its shape is the point. The reference — a fixed, native, wild-type stimulus chosen once per axis and never moved — is anchored at exactly **1,000**. Any candidate is scored by the *ratio* of the reference's EC_{50} to its own. A candidate twice as potent as the reference has half the EC_{50} and scores 2,000. A candidate a tenth as potent scores 100. Because EC_{50} sits in the denominator, the index is **inverse in EC_{50} and linear in potency**: loudness and score rise together, step for step, with no logarithmic flattening to hide a gain inside. When a cultivar moves from 1,500 to

3,000 on an axis, it is not "a bit louder on a compressed scale." It is twice as loud to the cell. Numbers, on this index, are a form of respect — for the receptor, and for the reader, who is owed a figure they can check rather than an adjective they must trust.



A designed superhot, cut to the placental ribs where the capsaicinoids are stored — the object the index was built to read.

Two conventions keep the index legible across a whole library, and both matter for reading the rest of this report. The first is the **radar cap**. When an eight-axis profile is drawn as the scorecard octagon you will see throughout these pages, each spoke is normalized to a ceiling of **5,000**, so that one cultivar's octagon can be laid over another's and compared by eye without rescaling. A raw axis score can in principle run past 5,000 — the mathematics does not forbid it — but the printed radar holds that ceiling so the shape stays honest page to page. The second convention is the **reference baseline** itself, which is worth seeing drawn out, because it is the quiet zero of everything here:

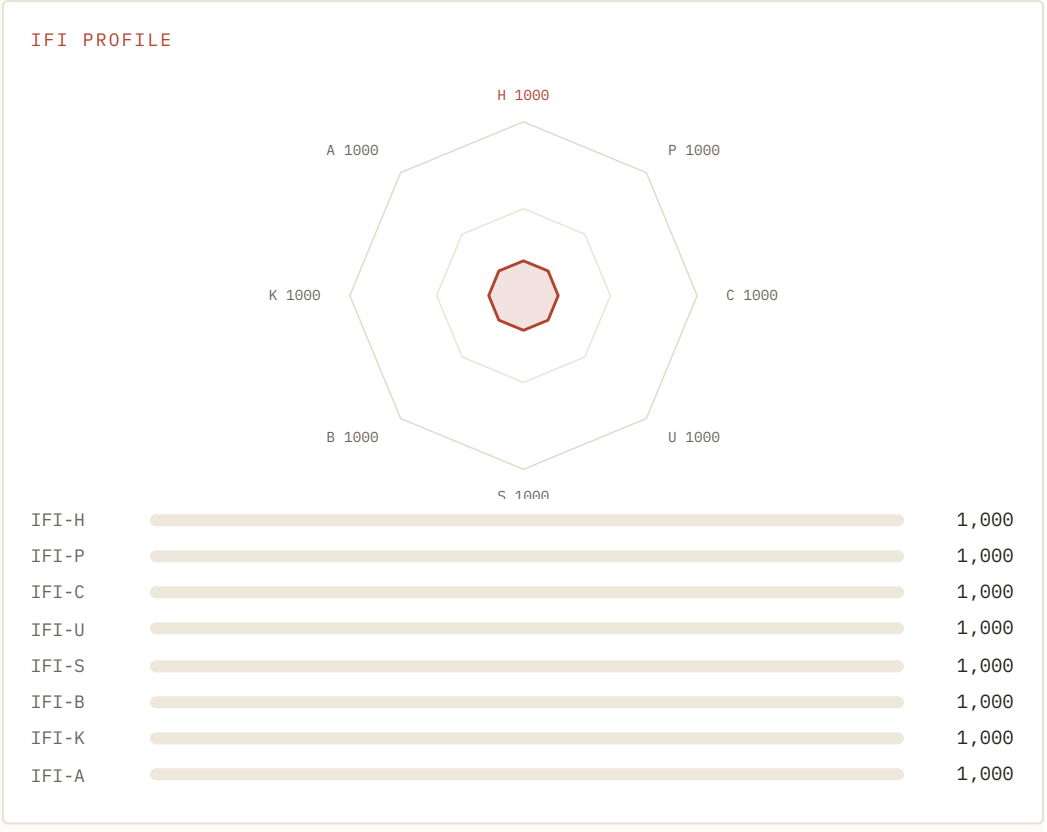

That octagon — a perfect, unremarkable ring at 1,000 on every spoke — is not a cultivar. It is the *definition of the reference*: the native wild-type palate against which all eight axes are pinned, scoring exactly one thousand on each by construction. Every scorecard in this report is a departure from that ring. When you see a profile bulge toward heat and hold near baseline elsewhere, you

are seeing a plant that speaks one axis loudly and the rest at the volume the wild world always did. When you see the ring pushed out on all eight at once — and you will, once — you are seeing something that did not exist until we built it.

This is the instrument. A fixed reference, a receptor midpoint, an inverse ratio, a thousandfold anchor, a five-thousand ceiling for the eye. Simple enough to fit on one line; strict enough that a tomato, a chili, and a cup of tea can be read against each other on one page. We did not build it to replace the poetry of the table. We built it because, as every chapter from here will show, you cannot design what you cannot measure — and you cannot report, honestly, on a year of design unless you first agree on what a single number means.

Chapter 2 · The Eight Dials — *what the tongue is actually doing.*

Taste, as most people picture it, is a map of the tongue with regions for sweet and sour and bitter — a diagram that is charming, durable, and wrong. The real apparatus is a field of receptor proteins, studded across taste cells and, crucially, across the free nerve endings of the trigeminal system that give food its *feel* as much as its flavor. There is no region for heat. There is a channel for it, on a nerve, and it reports a fire that is chemically real and thermally imaginary. The eight axes of the index are not eight places. They are eight **dials** — eight distinct molecular sensors, each with its own mechanism, its own agonists, and its own EC_{50} to anchor against. Knowing what each one physically *is* is the difference between seasoning a dish and engineering one.

Here is the full set, as the index defines it:

CODE	AXIS	RECEPTOR	WHAT THE MOLECULE DOES TO THE CELL
IFI-H	Heat	TRPV1	Opens a cation channel on trigeminal nerve endings; reports "hot"
IFI-S	Sweetness	TAS1R2 · TAS1R3	Closes a GPCR's Venus-flytrap domain; reports "sweet"
IFI-U	Umami	mGluR1 · TAS1R1·TAS1R3	Binds glutamate in a flytrap cleft; reports "savory depth"
IFI-B	Zero-Bitter	TAS2R38	A bitter GPCR whose *silencing* is the design target
IFI-K	Kokumi	CaSR	Calcium-sensing receptor; reports "richness, continuity"
IFI-C	Cool-Burn	TRPM8 · TRPV1	Cold/menthol channel; reports "cool" with no drop in temperature
IFI-P	Pungency	TRPA1	Electrophile sensor; reports a sharp, sinus-filling "sting"
IFI-A	Astringency	Salivary PRPs	Tannins precipitate lubricating proteins; reports "grip, dryness"

Four of these eight are **class-C G-protein-coupled receptors** — TAS1R2·R3 for sweet, the TAS1R1·R3 and mGluR1 pair for umami, the TAS2R family for bitter, and the calcium-sensing receptor for kokumi. Class-C GPCRs share a signature architecture: a large extracellular **Venus-flytrap domain**, two lobes hinged like a shell, that close around their ligand and transmit the closure through a cysteine-rich linker to the seven-helix bundle in the membrane, which wakes the

G-protein inside the cell. It is a beautiful mechanism to design against, because potency lives in the fit of the flytrap — and the fit is a surface we can read, score, and tune.



The pepper, whole and warm — the organism in which all eight dials were first learned, and the native speaker of the heat axis.

The other four are **ion channels**, and they are faster and blunter and, to our minds, more dramatic. **TRPV1**, the heat dial, is a non-selective cation channel gated three ways at once — by temperatures above about 43 °C, by protons, and by capsaicin, which binds an intracellular vanilloid pocket near the channel's S4–S5 linker and lowers the thermal threshold so the body reports a heat that is not there. **TRPM8**, the cool dial, is TRPV1's mirror: it opens below roughly 26 °C and to menthol, producing a coolness with no actual loss of heat — a weatherless winter on the tongue. **TRPA1**, the pungency dial, is the strangest of the set: it is an *electrophile* sensor, and its agonists — allicin from garlic, the isothiocyanate of mustard and wasabi, the cinnamaldehyde of cinnamon — do not merely dock, they covalently modify specific cysteine residues on the channel's inner face, a chemical handshake that explains why pungency feels less like flavor and more like an event. And the eighth dial is not a receptor at all in the usual sense: **astringency** is read by **salivary proline-rich proteins**, which tannins bind and precipitate, stripping the mouth of its lubrication so that the sensation we call "dry" or "grippy" is really the sudden friction of unlubricated tissue — taste as the absence of a slick, reported by the mouthfeel system rather than a taste bud.

One of the eight is defined, uniquely, by *subtraction*. **Zero-Bitter** is anchored to TAS2R38 — the receptor of the famous PTC/PROP strips that divide the world into tasters and non-tasters by a single haplotype — but the design target is not to make it fire. It is to make it *quiet*. Bitterness is the only one of the eight sensations most foods have too much of rather than too little, and so the Zero-Bitter axis measures an attenuation: the index records, this year, bitterness reductions ranging from **20% to 65%** across the host organisms built for it, which is why a coffee or a cacao engineered on this axis reads not as "sweeter" but as *unclenched*.

Eight dials, then, and not one of them a region on a map. A heat channel on a nerve, a flytrap for sugar, a flytrap pair for savor, a bitter switch we work to silence, a calcium sensor for richness, a cold channel for weatherless cool, an electrophile alarm for the

sharp, and a protein precipitation we read as grip. The state of flavor, in any given year, is the state of these eight — how loudly we have taught each organism to speak them, and how precisely. The chapters of Part II take them one at a time. But they are, always, a set: the tongue reads them in parallel, in a fraction of a second, and the composite it assembles is the thing you actually taste.

Chapter 3 · The Arithmetic of Loudness — *potency, not quantity.*

The index gives every axis a number; the interesting mathematics begins when a single plant speaks more than one. A fruit is not a scalar. It is a **vector in eight-dimensional space**, one component per dial, and almost everything that is hard and everything that is beautiful about designing flavor lives in the geometry of that space. This chapter is the arithmetic the rest of the report leans on: how the single-axis score behaves, how axes compose, and why stacking them is a problem in flux before it is ever a problem in taste.

Start with the single axis and the thing the inverse-ratio form of the index buys us. Because

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

a change in the score is a direct read of a change in receptor potency, with no scale distortion in between. Suppose a reference heat stimulus has an EC_{50} of 10 µM at TRPV1, and a designed extract reaches half-activation at 2.38 µM — a little over fourfold more potent. The index returns 1,000 × 10 ÷ 2.38 ≈ **4,202**, which is, to the digit, the heat score of the cultivar you will meet at the center of this report. The number is not a grade on a curve. It is a potency, carried straight through. Halve a candidate's EC_{50} and you double its IFI; there is no point on the scale where a real gain in loudness shrinks into a cosmetic bump. This is why, across a whole library, the index stays comparable: a move from 800 to 1,600 on the quietest axis means exactly what a move from 2,000 to 4,000 means on the loudest — twice as loud to the cell, both times.



Phoenix — the one cultivar whose eight-dimensional vector reaches outward on every axis at once; the composite made visible.

Now let the plant speak more than one dial, and three distinct kinds of composite appear — each tracked as its own library this year.

The first is **Full-Spectrum**, coded IFI-F: a single organism expressing all eight axes above threshold at once. This is the rarest and most demanding composite, because eight pathways must run in one tissue without any of them starving the others, and the index treats "fullness" as a count — how many of the eight spokes clear the

expression threshold — layered on top of the eight raw scores. The full-spectrum product line stands at **300 formulated IFC products** this year, the apex of which is a plant we will come to that expresses **8 of 8**. The second composite is **Multi-axis**, coded IFI-M: deliberate two- and three-axis designs, where the point is not everything-at-once but a specific, ownable *chord*. The dual-axis library (M2) holds **8,400 constructs** this year; the triple-axis library (M3), **2,800**. A "swicy" heat-and-sweetness design, a savory umami-and-kokumi depth, a bitter-sweet balance — these are M2 chords, and the combinatorics are enormous: with eight axes, there are 28 possible pairs and 56 possible triples before a single organism is chosen to carry them.

The reason the composites are hard is **not** the receptors. It is the **supply chain inside the cell**. Every flavor molecule is the end of a biosynthetic pathway, and every pathway draws on shared precursor pools — the same carbon skeletons, the same CoA thioesters, the same aromatic amino acids feeding the phenylpropanoid branch that we will trace in the heat chapter. Stack two pathways naively and they compete for substrate; the second axis you add can quietly rob the first, so that a plant designed to be loud on two dials ends up mediocre on both. This is why no composite ships on receptor logic alone. Each one is first run through a **genome-scale flux-balance model** — a stoichiometric accounting of every reaction the plant can perform — to confirm that the precursor supply can feed all the enabled pathways simultaneously at the fluxes the target scores demand, and to find the rate-limiting step that must be relieved for the chord to hold. The arithmetic of loudness, in other words, is only half pharmacology. The other half is metabolic bookkeeping: making sure that when you turn up a second dial, the cell has the carbon to pay for it. Hold that idea. It is the hinge on which the entire method turns, and Part IV returns to it with the enzymes named.



PART II OF VI

The State of the Eight Axes

Heat · Sweetness · Umami · Zero-Bitter · Kokumi · Cool-Burn · Pungency · Astringency

Chapter 4 · Heat — *the axis everyone knows.*

Heat is where the company began and where the index is loudest, so it is where the state of flavor is easiest to read. Capsaicin binds TRPV1, the same channel that answers to scalding temperature and to acid, and it works by lowering the channel's thermal threshold until the body reports a fire that the thermometer cannot find. That is the whole trick of chili heat: no tissue is warmed, no damage is done at culinary doses, and yet the trigeminal system insists, with complete sincerity, that something is burning. The sensation is a report, not an injury — which is exactly why we are comfortable treating heat as a pleasure to be dialed rather than an assault to be survived.

To design heat you have to make capsaicin on purpose, and that means running a pathway that begins far upstream in a chemistry every green plant already owns. The **phenylpropanoid branch** — the sequence that in other plants builds lignin, pigments, and defense compounds — is the trunk:

STEP	ENZYME	FROM → TO
1	PAL	phenylalanine → cinnamic acid
2	C4H	cinnamic acid → *p*-coumaric acid
3	4CL	*p*-coumaric acid → *p*-coumaroyl-CoA
4	HCT	*p*-coumaroyl-CoA → (toward the caffeoyl/feruloyl pool)
5	COMT	→ ferulic acid / vanillin precursors

Everything to that point runs in every vascular plant, which is the single most important fact in this entire report and the reason Part V exists. Only the last four enzymes are pepper-specific: **pAMT**, which transaminates vanillin to **vanillylamine**; the branched-chain machinery — **BCAT** feeding **KAS** — which builds the fatty-acid partner, the 8-methyl-branched acyl-CoA; and **Pun1/CS**, the acyltransferase that condenses the two halves into capsaicin itself. The length and saturation of that acyl chain set the ratio of capsaicin to dihydrocapsaicin, and that ratio sets the *shape* of the burn — whether it blooms at the front of the mouth and fades, or arrives late and stays. Heat is not one sensation. It is a curve, and the pathway draws it.



Scorch (SS-0093) — a designed superhot at 2.21 million SHU, five of eight axes expressed.

On the index, IFI-H is anchored so that 1,000 is the half-activation of the native wild-type reference, and from there the heat ladder climbs. A sunlit shishito sits near 40 — heat as warmth, barely a whisper on the channel. The flagship superhots, the proof-of-concept line of *Capsicum chinense* cultivars we built to walk the scale deliberately, climb from there to a plant that registers **4,200**, which corresponds to a predicted heat of approximately **2.95 million Scoville Heat Units** in our model — a potency that did not exist in any pepper until we designed it gene by gene. Scoville Heat Units and IFI-H are two readings of the same event: SHU is historically a dilution assay, roughly linear in capsaicinoid concentration, while IFI-H is linear in TRPV1 potency, and across the designed line the two track each other closely enough that either can be converted to the other. The difference is that IFI-H can be read on the same page as sweetness and umami, and SHU cannot.

The heat ladder is the clearest annual view we have of the whole program's reach, because it is a single axis walked in even steps:

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

Read down that right-hand column and you see the thesis of the whole index hiding inside a heat book: as the plants climb in Scoville, they also climb in *axes expressed*. Heat was never the only thing we were building. It was the axis we learned on, the one loud enough to prove the method, and the ladder to everything else. The full account of these ten cultivars — the threat-level lore and the plant that scores on all eight dials at once — is Chapter 13. Here it is enough to say that heat, the axis everyone knows, is the one the

index reads most confidently, and that this year it reaches very nearly three million Scoville with a batch-to-batch consistency the wild pepper has never offered.

Chapter 5 · Sweetness — *sugar's argument, answered.*

Sweetness is the axis with the most at stake outside the kitchen, because sweetness is the one sensation the modern food supply has spent a century oversupplying with a single, metabolically expensive molecule. The index approaches it from the receptor rather than the sugar bowl, and that change of vantage is the whole opportunity. Sweetness is read by a single heterodimer — **TAS1R2 paired with TAS1R3** — a class-C GPCR whose Venus-flytrap domain closes around a sweet ligand and signals. The receptor does not care whether the thing in the flytrap is a sugar. It cares about fit. And the fit of sucrose, it turns out, is mediocre: sucrose is a *quiet* agonist with a high EC_{50}, which is precisely why you need so much of it. Nature, elsewhere, already makes far louder ones.

This is the design space the sweetness axis occupies: not sweeter sugar, but sugar-free loudness. The steviol glycosides of *Stevia*, the mogrosides of monk fruit, and the sweet proteins — **brazzein**, thaumatin, monellin — are agonists that dock into the TAS1R2·R3 flytrap orders of magnitude more potently than sucrose, throwing the sweet switch at concentrations where sucrose would be silent. Brazzein is a small, heat-stable protein of 54 residues that can read thousands of times sweeter than sugar by weight; a steviol glycoside carries its sweetness on a diterpene scaffold decorated with sugars that tune both potency and the clean-versus-lingering quality of the finish. The index scores all of them on one axis, so that for the first time a beverage reformulator can compare a stevia line, a monk-fruit line, and a brazzein line not by anecdote but by IFI-S — by how loudly each speaks to the same receptor.



Stevia — a native speaker of the sweetness axis, where loudness comes from fit rather than from quantity.

The host roster for sweetness this year reflects a deliberate split between the plants that *make* the loud sweet molecules and the plants we *restore* sweetness to. On the making side sit **Stevia** (910 constructs), **Monk Fruit** (80), and **Sugarcane** (200, a frontier host), each engineered to tune the glycoside or mogroside profile toward a cleaner, rounder sweetness with less of the bitter or licorice tail that makes a consumer set the glass down. On the restoring side sits the

larger prize: **Tomato** carries 2,040 constructs and is as much a sweetness project as an umami one, because the modern tomato lost its sweetness to the same industrial selection that cost it its savor. Across these hosts, the sweetness axis is where the index does its quietest and most consequential work — lifting the perceived sweetness of a food while lowering the sugar it must carry, which the index records as a receptor fact and never, ever as a health claim. That line matters enough to state plainly here and once: the index measures loudness on TAS1R2·R3. It does not measure, and we do not assert, anything about a body downstream of the tongue. What a formulator does with a louder, lighter sweetness is theirs to say; our job ends at the receptor, with a number.

The subtlety the index exposes — and this is the part reformulators pay for — is that sweetness is not only *amount* but *shape*: onset, peak, and decay, the sweet analogue of the burn curve. Sucrose has a clean, fast, forgettable shape. Many high-potency sweeteners are louder but *longer*, with a lingering tail the tongue reads as artificial. Much of this year's sweetness work was spent not on raising IFI-S but on fixing that shape — tuning the glycoside decoration and the protein sequence so that the loud sweetness also arrives and leaves like sugar. An axis, it turns out, has a waveform, and the state of sweetness this year is less about reaching new peaks than about making the peaks behave.

Chapter 6 · Umami — *the depth of being cared for.*

Of the eight axes, umami is the one our family argues about most at the table, because it is the one with the least adequate vocabulary. "Savory" is a gesture, not a description. What umami actually names is the sensation of glutamate — and a handful of related amino acids and nucleotides — binding the umami receptors, and what it *feels* like is depth: the quality that makes a broth taste less like water with flavor in it and more like something that was made for you. It is, more than any other axis, the taste of care. And it is the axis the industrial tomato gave away.

The receptor story has a wrinkle that makes umami the most interesting axis to engineer. The primary sensor is the heterodimer **TAS1R1·TAS1R3**, a class-C GPCR with — again — a Venus-flytrap domain, this one closing specifically around L-glutamate; a second receptor, **mGluR1**, a metabotropic glutamate receptor borrowed from the nervous system's own vocabulary, contributes to the percept. But the thing that makes umami *umami* is **synergy**. Glutamate alone is a moderate agonist. Add a 5′-ribonucleotide — inosine monophosphate (IMP) from meat and fish, guanosine monophosphate (GMP) from mushrooms — and the ribonucleotide binds a site on the near lobe of the flytrap and acts as a *positive allosteric modulator*, stabilizing the closed, active conformation and multiplying the glutamate response severalfold. This is why dashi (kombu glutamate plus bonito IMP) and a long-simmered sofrito (tomato glutamate plus aromatics) taste like far more than their parts. The index captures the synergy directly, because synergy lowers the effective EC_{50}, and a lower EC_{50} is a higher IFI-U. The number *knows* about the chord.



Tomato in cross-section — the organism in which the umami axis was most dramatically restored, from an IFI-U of 300 to 2,600.

The flagship umami project is the tomato, and it is the clearest case in the whole library of the method restoring what breeding lost. A vine-ripened heirloom holds a free-glutamate pool that reads richly on IFI-U; the shipping tomato, bred for firmness and shelf life and the brutal logistics of the truck, lost most of it — by our accounting, the modern commodity tomato surrenders the great majority of its glutamate to selection for traits that have nothing to do with taste.

The restoration works at the **GAD/GDH1 node** of nitrogen metabolism: glutamate decarboxylase (GAD) drains free glutamate toward GABA, while glutamate dehydrogenase (GDH1) governs its synthesis and turnover, and rebalancing that node — easing the drain, favoring the pool — lifts the tomato's IFI-U from roughly **300 to 2,600**. That is not a new flavor added to a tomato. It is the tomato's own savor, grown back in through its own native route. Chapter 14 tells the tomato's full story; here it stands as the umami axis's proof that "designed" and "natural" are not opposites when the design uses the plant's own enzymes to restore the plant's own chemistry.

Beyond the tomato, the umami roster is the savory heart of the library: **Shiitake** (80 constructs, a native GMP powerhouse), **Mushroom** (40) and the generic savory hosts, **Koji Mold** (80, the engine of soy sauce and miso), **Yeast** (240) and **Soybean** (360, a frontier host), each tuned to raise free glutamate, the synergistic nucleotides, or both. These are the organisms behind the meat-analog and ferment revolutions, and the index's contribution is to let a developer of a plant-based broth specify umami the way an engineer specifies a tolerance — as IFI-U, with the synergy accounted for — rather than reaching, again and again, for the one additive the category is trying to leave behind. Umami, measured, stops being a secret. It becomes a target.

Chapter 7 · Zero-Bitter — *subtraction as design.*

Every other axis in this report is built by addition — make the plant louder on some dial. Zero-Bitter is the one axis built by subtraction, and that inversion makes it, quietly, the most commercially potent axis in the library, because bitterness is the sensation standing between an enormous number of genuinely nourishing foods and the people who will not eat them. Coffee, cacao, the whole brassica family, hops, citrus pith, many legumes and ancient grains — these are foods people *want* to want, and bitterness is the tax. The Zero-Bitter axis is the project of lowering that tax without lowering anything else.

Bitterness is read by the **TAS2R** family — some two dozen bitter GPCRs in humans, a deliberately promiscuous alarm system evolved to flag the alkaloids and glycosides that signaled poison to our ancestors. The index anchors the axis to **TAS2R38**, the best-characterized of them and the one behind the classic PTC/PROP taste-test strips: a single receptor whose common haplotypes (PAV versus AVI) split humanity into tasters and non-tasters, which is why your reaction to raw kale or a negroni may be built into your genome. TAS2R38 responds to thiourea-bearing compounds, and in the brassicas those come from **glucosinolates** — sulfur-rich defense compounds that, when the plant tissue is cut and its myrosinase enzyme fires, hydrolyze into the isothiocyanates and goitrin-type compounds that light the receptor up. In coffee and cacao the bitter load is different in origin — chlorogenic-acid lactones and the breakdown products of roasting, the theobromine and caffeine of the bean — but the receptor-level problem is the same: too many agonists, firing too loudly, on an alarm the eater cannot consciously switch off.



Coffee — where attenuating bitterness unclenches the cup without touching its aroma or its lift.

Designing on this axis means lowering the agonist load or the receptor's response to it, and the index records the result as an **attenuation**: this year, across the hosts built for Zero-Bitter, bitterness reductions span **20% to 65%**. The spread is the honest part. Sixty-five percent is a cacao or a brassica where the glucosinolate or polyphenol pathway could be substantially retuned;

twenty percent is a coffee where the goal was to unclench the cup without touching the aromatic complexity that bitterness is tangled up with — because bitterness in coffee is not pure liability, and a design that stripped it entirely would strip the coffee's soul with it. The art of the subtraction axis is knowing *how much* to take, and the index's job is to make "how much" a number a roaster or a chocolatier can specify and verify.

The host roster tells the story of an axis spread across the foods that most need it. **Broccoli** (380 constructs) and the broader brassica project — **Broccoli-Kale** (280), **Cabbage** (40), **Kale** (40), **Radish** (40) — target the glucosinolate route directly, which is the single biggest lever in the produce aisle for turning "good for you but bitter" into "good, and eaten." **Coffee** (40) and **Cocoa**, the great bitter pleasures, are tuned toward balance rather than erasure. **Hops** (80) and the brewing project sit here too, because bitterness is the brewer's deliberate signature and the ability to dial it precisely is a gift to the craft. **Grape** (1,610) carries a Zero-Bitter component alongside its astringency work, and **Potato** (80), **Orange** (120), and the rest round out an axis that is, in the end, about addition by subtraction: taking one thing away so that everything the food already had — the nourishment, the aroma, the reason it was worth eating — can finally be tasted.

Chapter 8 · Kokumi — *the axis you cannot name.*

There is a taste you have experienced ten thousand times and have no word for. It is the difference between a young stock and an aged one, between a fresh cheese and a matured one, between a sauce finished the fast way and a sauce that someone stood over — a quality of *continuity* and *mouth-filling richness* that does not taste of anything in particular and yet makes everything taste like more. The Japanese named it **kokumi**, and for a long time that was all anyone could do with it: name it. The index can measure it, because kokumi turned out to have a receptor, and that discovery is one of the quiet revolutions this axis represents.

Kokumi is read by the **calcium-sensing receptor (CaSR)** — a class-C GPCR better known for regulating blood calcium, which the taste system recruited to do a second job. Its agonists are **γ-glutamyl peptides**: glutathione (γ-Glu-Cys-Gly), the ubiquitous cellular antioxidant, and its relatives, most famously **γ-Glu-Val-Gly**, a peptide so active on CaSR that it enhances kokumi at vanishingly low concentrations. These peptides do not taste sweet or savory or anything else on their own; what they do is bind CaSR and modulate the *perception* of the other tastes — thickening the body, lengthening the finish, smoothing the transition between the first bite and the aftertaste so that a food reads as whole rather than assembled. Kokumi is less a flavor than a **grammar** for the other flavors, the connective tissue of a dish, which is exactly why it has no simple name and why its receptor hid in plain sight.



Garlic, cut to its rings — a native source of the γ-glutamyl peptides the kokumi axis is built on.

Engineering kokumi means engineering γ-glutamyl peptide content, and the host roster is a list of foods we already associate, intuitively, with richness. **Garlic** (560 constructs) is the flagship — raw garlic carries γ-glutamyl-S-allyl-cysteine, and aged garlic is a kokumi factory, which is the biochemical truth behind why a long-cooked allium base makes a dish taste *finished*. **Vanilla** (40) is a kokumi host as much as an aromatic one, because vanilla's contribution to a custard or a cream is partly this mouth-filling continuity, not only its fragrance. **Onion** (40), **Yeast** (240), and the fermentation hosts round out an axis whose constructs are modest in count but

outsized in effect, because a little kokumi changes the reading of everything around it. On the index, IFI-K measures CaSR activation directly, and the number lets a formulator do something that used to require a master's palate and twenty years: specify richness, and verify it.

The reason kokumi matters to this report beyond its own axis is that it is the great **composite multiplier**. Because kokumi modulates the perception of the other tastes, a well-placed IFI-K lifts the apparent umami, rounds the sweetness, and softens the edges of heat and bitterness — which is why the dual-axis library's most-requested chord, the one we treat in the intersections, is umami × kokumi. Kokumi is the axis that makes the other axes sound like they belong to the same instrument. It is, in a real sense, the axis of *care* made measurable: the taste of a thing made slowly, now specifiable in a number, so that the slowness can be designed in rather than hoped for.

Chapter 9 · Cool-Burn — *temperature without weather.*

Cool-Burn is the most purely magical axis in the index, because it is a lie the body tells itself and enjoys. A menthol crystal does not lower the temperature of your mouth by a single degree, and yet the mouth reports cold — a clean, spreading, almost mint-blue cold that arrives with no change in the weather. The axis is called Cool-Burn because it is the thermal axis entire: cooling at one end through TRPM8, with a warming shoulder that reaches back toward TRPV1, the heat channel from Chapter 4. It is thermomodulation without thermodynamics, and it is the axis that oral care, confectionery, and beverage cannot live without.

The cool end is read by **TRPM8**, TRPV1's structural and functional mirror — a non-selective cation channel that opens to actual cold, below roughly 26 °C, and to menthol, which binds and shifts the channel's activation so that it fires at temperatures it would otherwise ignore. The brain, receiving TRPM8 traffic, concludes "cold," because in every ancestral circumstance TRPM8 traffic *meant* cold. Menthol simply exploits the shortcut. The warming shoulder of the axis reaches, through certain agonists, toward TRPV1 — which is why some cooling compounds carry a paradoxical warmth, and why "cool-burn" is a truer name than "cooling." The index anchors IFI-C to TRPM8 activation, and scores the whole thermal sensation — the weatherless winter and its warm edge — on one axis.



Peppermint — the organism behind the five-gene menthol pathway, native speaker of the cool-burn axis.

To design cool, you build menthol, and menthol is the end of an elegant five-gene pathway that *Peppermint* — the flagship cooling host at **2,360 constructs** — runs in its glandular trichomes. It begins with (+)-limonene and proceeds through a sequence of hydroxylation, oxidation, and reduction steps: limonene-3-hydroxylase opens the ring position, and a cascade through isopiperitenone and (+)-pulegone is resolved by pulegone reductase and menthone reductase into (−)-menthol, the enantiomer the human TRPM8 finds loudest. Tuning this pathway — the balance of the reductases, the suppression of the off-notes like pulegone that

read as harsh — is how you move IFI-C and, just as importantly, how you shape the *quality* of the cool from the sharp medicinal cold of pure menthol toward the rounder cool of a fine peppermint.

The cooling roster is a small, specialized, demanding family of hosts: **Peppermint** (2,360) and **Spearmint** (80) for menthol and carvone chemistries, **Eucalyptus** (80) and **Camphor** (40) for the eucalyptol and camphoraceous cools that read brighter and more penetrating, and **Lavender** (160) for a floral cool that oral care and fine beverage prize. **Orange** (120) carries a cool-burn component alongside its bitter work. These are hard organisms — the engineering difficulty across the cooling roster runs consistently high — because the trichome chemistry is intricate and the margin between a beautiful cool and a harsh one is narrow. But the payoff is an axis that lets a product designer specify exactly how cold a thing should feel without making it cold, exactly how long the cool should last, and exactly where on the spectrum from mint to eucalyptus it should sit. Temperature, decoupled from thermodynamics and handed to the formulator as a number. The body never notices it has been fooled. It only notices that it is pleased.

Chapter 10 · Pungency — *the sharp channel.*

Pungency is the axis people confuse with heat and should not, because it is a different channel, a different chemistry, and a different sensation entirely. Chili heat blooms on the tongue and builds. Pungency *stabs* — the sinus-clearing rush of wasabi, the nose-prickle of hot mustard, the sharp penetrating bite of fresh horseradish or raw garlic — and it does so because its receptor is not a thermal channel at all but an electrophile alarm, and its agonists do not merely dock, they chemically attack. Pungency is the most *reactive* axis in the index, in the literal chemical sense, and that reactivity is the whole of its character.

The receptor is **TRPA1**, a cation channel on the same trigeminal nerve endings as TRPV1 but tuned to a different threat: electrophilic and reactive compounds. Its agonists — **allicin** from crushed garlic, **allyl isothiocyanate (AITC)** from mustard and wasabi and horseradish, **cinnamaldehyde** from cinnamon, the gingerols and shogaols of ginger, the acrolein of smoke — share a chemical property rather than a shape. They are electrophiles, and they activate TRPA1 by **covalently modifying specific cysteine residues** on the channel's inner face, forming transient adducts that force it open. This is a handshake that leaves a mark, which is why pungency feels so immediate and so bodily — less a flavor the tongue savors than an event the nerve survives. It is also why pungency is volatile and fleeting: the reactive compounds that cause it are, by their nature, unstable, firing and then gone, which is why fresh-grated wasabi is a different animal from the tube and why the index has to measure pungency fast, at the moment of release.



Black mustard — a native source of allyl isothiocyanate, the electrophile that fires the pungency channel.

The pungency roster is built on the plants that make these electrophiles, and the engineering is genuinely hard — the compounds are reactive, their precursors are defense chemistry the plant guards closely, and the margin for off-notes is thin. **Black Mustard** is the flagship at **2,790 constructs**, a powerhouse of AITC chemistry and one of the library's larger projects, with **White Mustard** (200) alongside it for the milder, more persistent pungency

of *para*-hydroxybenzyl isothiocyanate. **Garlic** (560) sits on this axis as well as on kokumi, because allicin is both a pungency agonist and the precursor of the γ-glutamyl peptides — one organism speaking two dials through two branches of the same sulfur chemistry. **Ginger** (120), **Cinnamon** (40), **Thyme** (120), **Oregano** (40), **Watercress** (120), and the aromatic herbs round out an axis that is the backbone of the condiment aisle and of whole cuisines. On the index, IFI-P anchors to TRPA1, and the number lets a maker of a mustard or a chili crisp or a gingery sauce specify the sharpness precisely — how hard the stab, how fast the release — which, for a category defined by a sensation that disappears while you are still describing it, is a capability close to miraculous.

Pungency also matters to the composites, because it is the third leg of some of the most prized chords in cuisine. The numbing-heat intersection of Sichuan cooking — the *málà* tingle — is a composition of heat (TRPV1), cool (TRPM8), and pungency (TRPA1) firing together, three channels on the same nerve, and the index is the only instrument that can write that chord down as three numbers and let a chef rebuild it on purpose. Pungency, the sharp channel, is where flavor stops being something you taste and becomes something that happens to you — and the axis that lets a designer decide, exactly, how much.

Chapter 11 · Astringency — *the finish, and the grip.*

The eighth axis is the one that is not, strictly, a taste at all, and that is precisely what makes it the connoisseur's axis — the dimension that separates a good wine from a great one, a flat tea from a structured one, a cheap chocolate from a fine one. Astringency is not sweet or bitter or savory. It is a *texture*: the drying, puckering, grippy sensation that a tannic red wine or an over-steeped tea leaves across the whole mouth, a loss of lubrication that the mouthfeel system reports as structure. It is the axis of the finish — the part of the experience that happens after you swallow — and the index measures it through a mechanism unlike any of the other seven.

Astringency has no receptor in the conventional sense. It is read by the **salivary proline-rich proteins (PRPs)**, a class of lubricating proteins the mouth secretes precisely so that tannins have something to bind other than your tissues. **Tannins** — the condensed proanthocyanidins and hydrolyzable tannins of grape skins, tea leaves, pomegranate, and cacao — are polyphenols with a powerful affinity for these proline-rich proteins, and when they meet in the mouth they **precipitate**, the tannin crosslinking the protein into insoluble complexes that fall out of solution. The lubrication is stripped, the tissue surfaces drag against each other, and that friction — the mouth suddenly running dry and rough — is the sensation of astringency. It is taste as the *removal* of a slick, which is why it builds and lingers and cannot be washed away with a single sip: the saliva has to be remade.



Grape in cross-section — the tannin structure of the skin is the astringency axis made visible.

Designing astringency means designing tannin structure, and the key enzymes sit at the end of the flavonoid pathway: **anthocyanidin reductase (ANR)** and leucoanthocyanidin reductase (LAR) produce the epicatechin and catechin monomers that polymerize into condensed tannins, and the *degree* of that polymerization — the chain length and the linkage pattern — sets whether the astringency reads as harsh and green or as the fine, ripe, structured grip that a winemaker spends a career chasing. **Grape** is the flagship at **1,610 constructs**, the largest astringency project and the frontier of engineered viticulture, where tuning the ANR route lets a grower

design tannin structure rather than pray for the vintage. **Tea** (280) carries its astringency through catechins and sits at the intersection of astringency, umami, and bitterness — the three-axis structure that distinguishes a green from a black from an oolong. **Pomegranate** (120), **Blueberry** (120), **Raspberry** (40), **Sorghum** (120, a frontier host), and **Cocoa** round out an axis that is the home of the library's most sophisticated customers: winemakers, tea houses, and fine-chocolate makers, for whom astringency is not a flaw to remove but a structure to compose.

The reason astringency closes Part II is that it is the axis of *aftertaste* — the dimension that governs what a food leaves behind, the memory of it, the length of the finish. It is the most temporal of the eight, the one that unfolds over seconds after the others have spoken, and the index's ability to put a number on it — IFI-A, anchored to the PRP-precipitation event — gives the makers of the world's most considered beverages a language for the part of their craft that was, until now, the most purely intuitive. With astringency, the eight dials are complete: a heat on a nerve, two flytraps for sweet and savory, a bitter alarm we silence, a calcium sensor for richness, a cold channel for weatherless cool, an electrophile alarm for the sharp, and a protein precipitation for the grip of the finish. Eight dials, eight numbers, one instrument — and, now, one organism after another taught to speak them.



PART III OF VI

The State of the Organisms

Forty-two validated hosts · 18,500 constructs · the ten flagships · the tomato restored

Chapter 12 · Forty-Two Hosts — *the roster this year.*

An axis is a capability; an organism is where the capability lives. The state of flavor, in any given year, is finally the state of the **roster** — which plants we have taught to speak which dials, how deeply, and at what cost in engineering effort. This year the validated roster stands at **42 host organisms**, carrying a combined **18,500 engineered constructs** — a figure that is not an estimate but a sum, the roster totaled host by host — within the library's locked total of roughly 19,000 constructs overall. Reading the roster is the single best way to take the measure of the whole program, because it shows at a glance where the method is deep, where it is young, and where it is hard.

Here is the roster in full, ranked by construct count — the clearest annual snapshot the index produces:

ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Pepper	1,060	all	Native
Tomato	2,040	S · U · F · M	Easy
Habanero	3,190	F · M	Hard
Black Mustard	2,790	P · M	Hard
Peppermint	2,360	C · M	Hard
Grape	1,610	B · A · M	Hard
Stevia	910	S · F · M	Moderate
Garlic	560	K · P	Moderate
Broccoli	380	B · M	Moderate
Soybean	360	U · K	Frontier
Tea	280	U · B · A	Hard
Broccoli-Kale	280	B	Hard
Yeast	240	U · K	Hard
Sugarcane	200	S	Frontier
White Mustard	200	P	Hard
Lavender	160	C	Hard
Blueberry · Ginger · Orange · Pomegranate · Sorghum · Thyme · Watercress	120 each	A / P / B·C / A / A / P / P	Moderate–Frontier
Potato · Hops · Koji Mold · Monk Fruit · Eucalyptus · Shiitake · Spearmint	80 each	B / B / U / S / C / U / C	Moderate–Hard
Coffee · Cabbage · Kale · Onion · Raspberry · Camphor · Cinnamon · Mushroom · Oregano · Radish · Rapeseed · Vanilla	40 each	B / B / B / K / A / C / P / U / P / B / B / K	Moderate–Hard

Three things in that table are worth reading as the state of the roster. First, **depth is not the same as difficulty**. Pepper is marked *Native* — it is where the method was born, and 1,060 constructs is a mature library, not a struggle — while Habanero at 3,190 constructs is marked *Hard* and is the single largest project in the roster, a full-spectrum and multi-axis workhorse that earned its depth the hard way. The construct count measures how much we have built; the difficulty rating measures how much it cost. Second, the **axis letters map the roster back onto Part II**: every B in that table is a Zero-Bitter project, every A an astringency project, and you can read the industrial priorities of the year straight off the column — the brassica bitterness cluster, the astringency fine-beverage cluster, the umami ferment cluster.



Soybean — a frontier host, young in the roster, carrying the umami and kokumi work that the ferment and meat-analog categories most want.

Third, and most important for everything that follows: the **frontier hosts** — Soybean, Sugarcane, Sorghum, marked *Frontier* — are the youngest and shallowest libraries, and they are the bridge to Part V. They are frontier not because they are exotic but because they sit at the edge of what the method can currently reach: large-genome crops, difficult transformation, pathways only partly mapped. Their modest construct counts are an honest signal of a young capability, and watching them deepen year over year will be one of the truest measures this report can offer of whether the method is genuinely traveling beyond the plants it was born in.

That travel is the quiet engine under the whole roster, and it is worth stating the mechanism plainly, because it is why 42 is a number that can grow. The platform is built as **seven modules** — the computational scoring, the flux-balance modeling, the edit design, the delivery, and the rest — and **five of the seven transfer unchanged to any organism with a sequenced genome.** Only the two modules closest to the terminal, pepper-specific chemistry have to be rebuilt per host. That ratio, 5 of 7, is the reason the roster is 42 and not 4, the reason Part V's roadmap is credible rather than fantastical, and the reason an annual report on this program is worth publishing at all: the roster is designed to lengthen, and the index is the instrument that will let us — and you — watch it happen, one validated host at a time.

Chapter 13 · The Flagships — *ten warnings and a verdict.*

Every library needs a proof, a set of objects that demonstrate the full reach of the method in a form you cannot argue with. For Scoville Splice that proof is the flagship line: ten cultivars of *Capsicum chinense*, built not as products first but as a deliberate walk up the heat ladder and outward across the axes, each named for an escalating threat and each one a checkpoint on the way to a single verdict. They are the most-cited objects in the whole catalog, and they carry a lore — the threat-level naming, the sense of a program approaching an edge — that is, underneath the drama, completely earned by the numbers.

Here is the line in full, the annual register of the flagships:

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

Read as a progression, the line tells a story the single numbers do not. Scoville climbs in even, deliberate steps from 866,000 to very nearly three million — not the jagged leaps of lucky breeding but the smooth ascent of a dial being turned. And in lockstep, the rightmost column climbs too: 2 axes, then 3, then 4, 5, 6, and finally 8. The flagship line is the index's own argument made flesh — that heat was always the loud axis we learned on, and that learning to build heat taught us to build everything. Each cultivar expresses more of the eight dials than the last, so that by the top of the ladder the plant is not merely the hottest we have made but the *most complete* — a full-spectrum organism that happens also to sit near the ceiling of heat.



Solara (SS-0095) — 2.72 million SHU, six of eight axes, the last checkpoint before the verdict.

The verdict is **Phoenix**, SKU SS-0096, and it is the only cultivar in the entire library that scores on all eight axes at once — the single realized Full-Spectrum organism, the plant the whole ladder was climbing toward. Its profile is the one complete eight-axis reading we have, and it deserves to be seen drawn out in full:

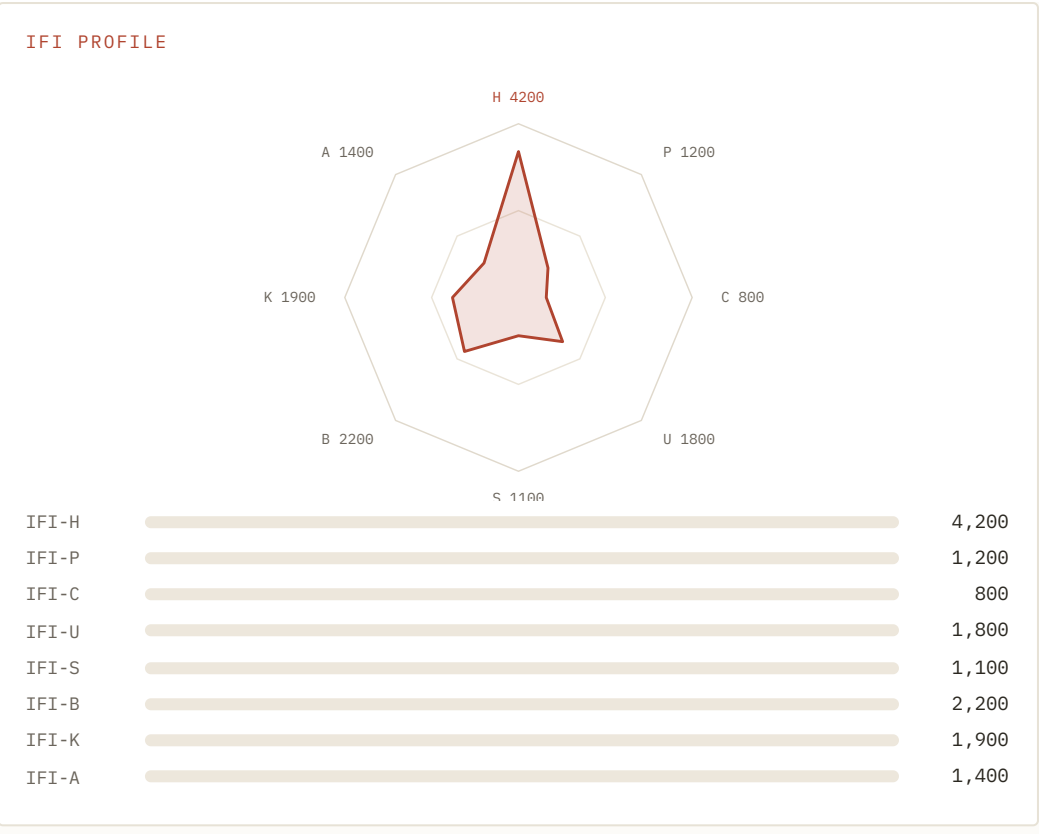

That octagon reaches outward on every spoke — a predicted **~2.95 million SHU** on heat, and genuine, above-threshold expression on pungency, cool-burn, umami, sweetness, zero-bitter, kokumi, and astringency all at once. Set it against the perfect 1,000-ring reference of Chapter 1 and you are looking at the difference between the palate the wild world offers and a palate designed from the receptor outward. No plant like it existed before we built it, because

no breeding program could have selected for eight receptor-level traits simultaneously; it had to be *designed*, which is the subject of Part IV.

The heat of Phoenix rests on two Pun1 substitutions in the Pun1/AT3 acyltransferase that no breeding program ever found, because they are not improvements a phenotype would reveal — they are changes to the enzyme's catalytic efficiency and fold stability that only a model looking directly at the protein could propose:

EDIT	EFFECT	ΔΔG (KCAL/MOL)
S39L	binding affinity ↑ at the N-terminus	−2.1
L345G	steric relief in the active-site cleft	−1.8
C175S	disulfide-misfold prevention	−0.9

Each negative ΔΔG is a stabilization — a variant that folds more reliably and catalyzes more efficiently — and the three together, written in with CRISPR, lift peak output to that 2,937,948 SHU figure at a **batch-to-batch coefficient of variation of ≤7%**, against 42% in the wild-type plant. That consistency, more than the peak, is the real verdict: the flagship line does not merely reach the edge of heat, it reaches it *reproducibly*, which is the difference between a record and a product. The ten warnings climb; the verdict holds steady. That is the state of the flagships this year.

Chapter 14 · The Tomato, Restored — *the proof the method travels.*

If the flagships prove the method's reach, the tomato proves its *travel* — that everything learned in the pepper transfers to a plant that is not a pepper, in a family that merely happens to be adjacent, toward axes that are not heat. The tomato is, by a wide margin, the most important single organism in this report, not because it is the hottest or the most complex but because it is the proof of concept for the entire thesis: that capsaicin was never a pepper problem, and that the platform is a flavor platform, not a chili one. At **2,040 constructs**, the tomato is the second-deepest library in the roster, and nearly all of that depth is spent on two axes the pepper barely touches — umami and sweetness.

The tomato is also, not incidentally, a small act of restoration, and that framing matters to how we think about the work. The tomato you can buy in February is a monument to logistics: bred for firmness, for uniform ripening, for the ability to survive the truck, and in the course of that breeding it gave away most of what made a tomato taste like one. The free-glutamate pool that carries its umami, the sugar balance that carries its sweetness — the industrial tomato surrendered the great majority of both to a century of selection for traits that have nothing to do with flavor. What the method does is grow them back in, and it does so through the tomato's *own* enzymes, which is the part we care about most.



The restored tomato — its umami lifted from an IFI-U of 300 toward 2,600 through the plant's own GAD/GDH1 route.

The umami restoration works at the **GAD/GDH1 node** of the tomato's nitrogen metabolism, the junction treated in Chapter 6: glutamate decarboxylase drains free glutamate toward GABA, while glutamate dehydrogenase governs its synthesis, and rebalancing that node — easing the drain, favoring the pool — lifts the tomato's umami from an IFI-U of roughly **300 to 2,600**, better than an eightfold gain. The sweetness work runs in parallel, tuning the sugar balance back toward the vine-ripened profile the shipping tomato lost. The result is marked, in the roster, with the axis letters **S · U · F · M** — a tomato that speaks sweetness and umami, carries a full-spectrum

component, and anchors some of the most-requested multi-axis chords in the library. And the engineering difficulty is marked *Easy*, which is the quiet headline: the second-deepest library in the roster, on the two axes most valuable to the entire food industry, was among the *least* difficult to build, because five of the seven platform modules transferred to the tomato unchanged. Easy, here, is not a small word. It is the whole argument.

What the tomato teaches, read as part of an annual report, is where the method is going. The pepper was the proof that flavor could be designed; the tomato is the proof that the design *moves* — out of the genus it was born in, into a staple crop, toward axes and markets far larger than heat. Every frontier host in Chapter 12 and every concept organism in Chapter 17 is, in a sense, a bet that the tomato's story repeats: that an organism with a sequenced genome and a native pathway can have its own best flavor, and its own fullest nourishment, grown back into it through its own chemistry. The tomato is the first time that bet paid off outside the pepper. It will not be the last, and tracking how far and how fast it travels is, year over year, much of what this report exists to do.



PART IV OF VI

The Method

ESM2 protein language models · genome-scale flux balance · CRISPR · ≤7% batch CV

Chapter 15 · From Sequence to Flavor — *design, not selection.*

Everything in this report rests on a single inversion, and it is worth stating as starkly as it deserves: for ten thousand years, agriculture selected phenotypes and hoped the genotype followed. We do the opposite. We start at the protein, design the genotype, and let the phenotype follow by construction. This is the difference between breeding and engineering, between hoping and knowing, and it is the method that turns the index from a measuring instrument into a *design* instrument. This chapter names the pipeline, module by module, because an annual report owes its readers not just the results but the machine that produced them.

The pipeline begins where the flavor begins: at the enzyme. To raise heat, you must improve the Pun1/AT3 acyltransferase that makes capsaicin, and the design space is enormous — a protein of hundreds of residues, each position substitutable to nineteen alternatives, a combinatorial explosion no breeding program could ever search. So we do not search it by hand. We score it with an **ESM2 650-million-parameter protein language model**, a neural network trained on the evolutionary record of millions of protein sequences until it has learned, in effect, the grammar of what a working protein can be. The model reads the Pun1 sequence and scores every possible single-residue substitution by its predicted effect on the two things that matter — **catalytic efficiency and fold stability** — ranking the entire mutational landscape in silico before a single plant is touched. The two Pun1 substitutions behind Phoenix's heat (S39L · L345G, with their negative ΔΔG values) came out of exactly this ranking: stabilizing edits a phenotype would never reveal, because they change the enzyme's efficiency and folding rather than anything a grower could see by eye.

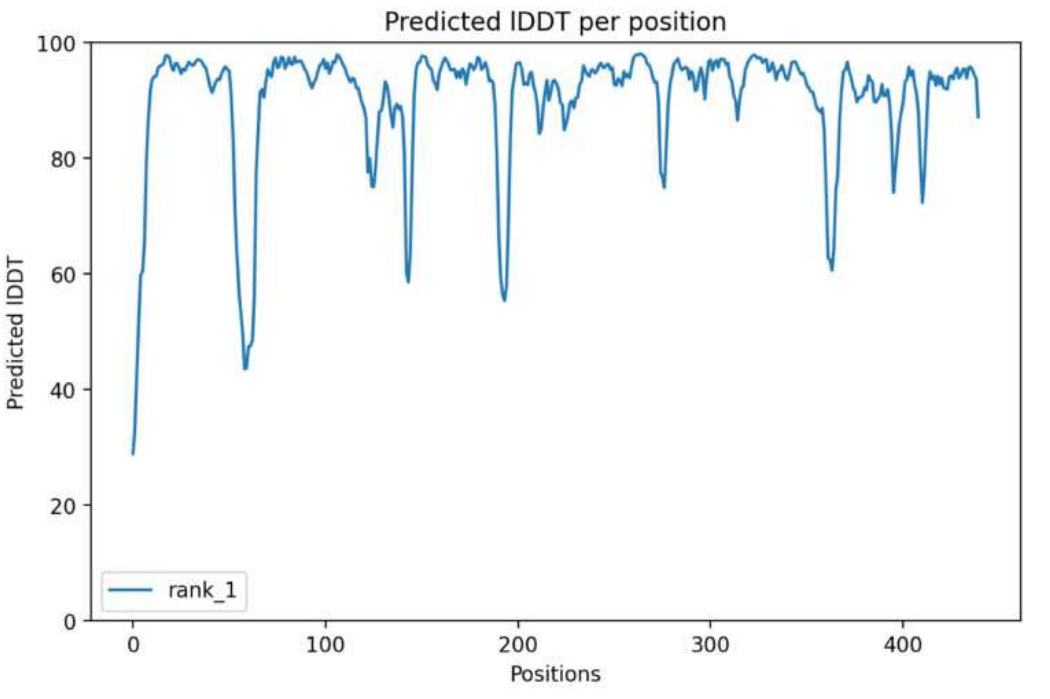

FIG. 1 — The Pun1/AT3 acyltransferase model coloured by pLDDT confidence; the binding domain and active-site cleft resolve with high certainty.

A high score for an enzyme edit is necessary but not sufficient, because — as Chapter 3 insisted — the enzyme is only one node in a metabolic network, and improving it is pointless if the network

cannot feed it. So the second module is a **genome-scale flux-balance analysis**: a stoichiometric model of every metabolic reaction the plant can perform, solved as a linear program to find whether the designed fluxes are feasible and, crucially, to identify the **true rate-limiting step**. It was flux-balance analysis that confirmed, for the capsaicinoid pathway, that the acyltransferase and not some upstream phenylpropanoid step was the real bottleneck — which is why the edits went where they went. The model tells you not only whether a design can work but *where to spend your edits*, so that you relieve the step that is actually limiting rather than the one that merely looks important. This is the module that keeps the composites honest, the one that ensures a second axis does not starve the first, and it is the single most underappreciated piece of the method.

The structural confidence behind the edits is readable, and we show it, because showing the work is a form of respect for the reader. The model resolves the Pun1 binding domain and active-site cleft with high per-residue confidence, and its predicted aligned error — the model's own estimate of how well it knows the relative positions of every pair of residues — is tight through exactly the regions the edits touch:

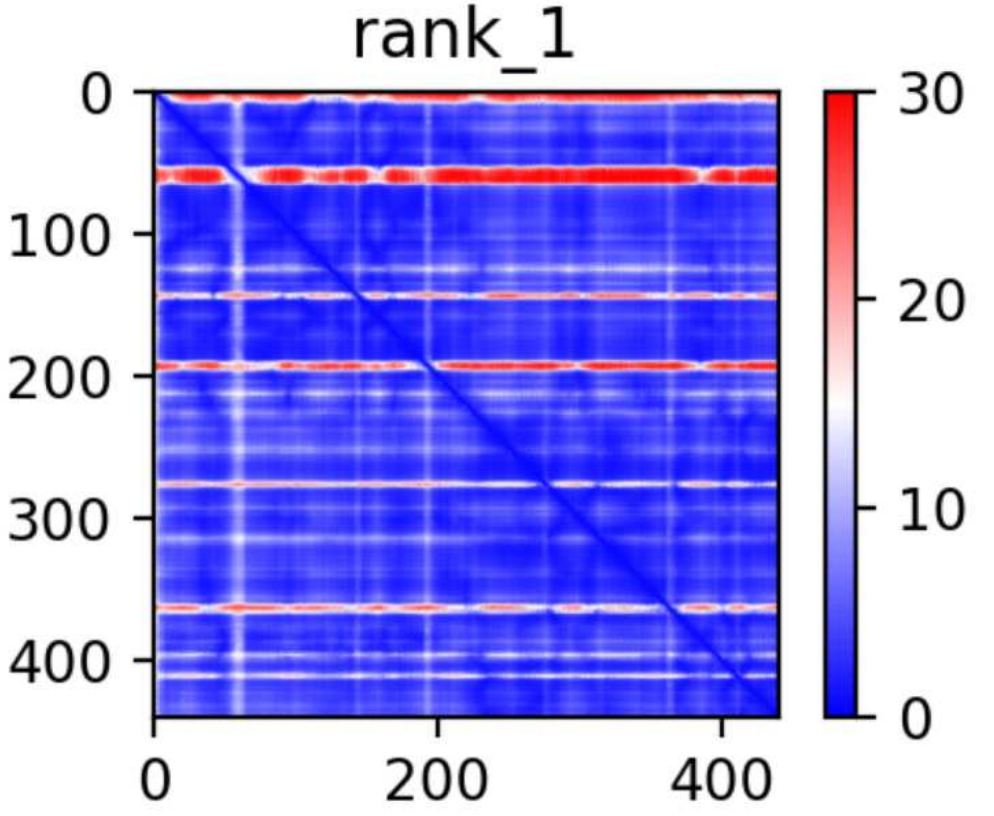

FIG. 2 — Predicted aligned error for the Pun1/AT3 model; low error (dark) through the catalytic core means the active-site geometry is trustworthy enough to design against.

With the edits designed and the fluxes confirmed, the third module **writes** them — **CRISPR**, used not to add foreign genes but to make the precise substitutions the model specified in the plant's own sequence, so that the result is the plant's own enzyme, improved. The remaining modules carry the design from edit to validated construct: delivery, the nanoparticle formulation work behind the

full-spectrum IFC products, scoring against the index, and the reproducibility verification that is the subject of the next chapter. Seven modules in all — and, as Chapter 12 established, **five of the seven transfer unchanged** to any organism with a sequenced genome. That transferability is not a convenience. It is the reason the method is a platform and not a single result, the reason the roster can grow from 42, and the reason this report has a Part V. Design, not selection: start at the protein, prove the flux, write the edit, measure the loudness. The plant does the rest, the way it always has — only now it is building the flavor we asked for, on purpose, to a number.

Chapter 16 · Precision as a Promise — *≤7%, every batch.*

There is a number in this report quieter than three million Scoville and more important than any of them, and it is a number about *variation*: **≤7%**. It is the batch-to-batch coefficient of variation of the designed cultivars — the measure of how much a trait wobbles from one grow to the next — and against the wild-type's **42%**, it is arguably the single most consequential figure the company produces. Peak performance makes headlines. Reproducibility makes products. This chapter is about why the index's most important state, in any given year, may be its precision rather than its reach.

Consider what a 42% coefficient of variation means in practice, because it is the daily reality of anyone who has ever cooked with chilies. Coefficient of variation is the standard deviation divided by the mean — the spread of a trait as a fraction of its average — and a CV of 42% on heat means that a pepper nominally at a million Scoville might, batch to batch, deliver anywhere across a brutally wide band, one harvest mild and the next incendiary from the same seed in the same soil. This is why no two pots of the same recipe taste the same, why hot-sauce makers blend obsessively to average out their inputs, and why "natural variation" is a polite industry name for *you cannot promise your customer anything*. The wild pepper is not a product. It is a lottery with a flavor.



A designed cultivar under controlled conditions — the environment in which a trait is held to ≤7% from batch to batch rather than left to the lottery.

Driving the CV from 42% down to **≤7%** is a sixfold tightening of the distribution, and it comes from the method's foundation: when a trait rests on a designed, stabilized enzyme and a flux-balanced pathway rather than on a sprawl of unselected background genetics, it stops wobbling. The edits behind Phoenix's heat were scored on *fold stability* as much as catalytic efficiency precisely because a stable enzyme is a consistent one — a protein that folds the same way every time catalyzes at the same rate every time, and the trait it produces holds its value grow after grow. Precision, in other words, is not a separate quality-control step bolted onto the method. It is a

consequence of designing at the protein level, which is why the same pipeline that reaches three million Scoville also holds it to within seven percent. The two achievements are the same achievement, read two ways.

For the people who actually build with this library — the formulators, the CPG developers, the sauce and beverage and ferment makers — the ≤7% figure is the one that changes their working life, and it is worth being concrete about why. A formulator specifies tolerances the way an engineer does: this much heat, this much umami, within this band, batch after batch, because a product is a *promise* that next year's jar tastes like this year's. On a 42% input you cannot make that promise; you can only blend, taste, correct, and hope. On a ≤7% input you can specify an IFI profile and *hold* it, which means you can design a product around a cultivar the way you would design around any reliable component. This is the difference between flavor as agriculture and flavor as engineering, and it is the reason the index and the CV belong in the same chapter: the index gives you a number to design to, and the CV promises that the number will still be true at scale. Reach is what the method can do once. Precision is what it can do every time — and every time is what a promise is made of.



PART V OF VI

The Frontier

Sixty-four concept organisms · five tiers · what could be built, and what we will not rush

Chapter 17 · The Sixty-Four — *what could be built.*

An annual report owes its readers not only the state of what is but a disciplined account of what could be, and for Scoville Splice that account has a name and a shape: the **Beyond-Peppers roadmap**, sixty-four concept organisms organized into five tiers of ascending difficulty. It is the formal answer to the question the whole method provokes — *if capsaicin is not a pepper problem, then what?* — and it is built on the single fact established in Chapter 4 and proven by the tomato in Chapter 14: the phenylpropanoid branch, PAL → C4H → 4CL → HCT → COMT, runs in every vascular plant, and only four terminal enzymes — pAMT, Pun1/CS, KAS, BCAT — are pepper-specific. Everything in the roadmap is a bet on that transferability, graded by how many genes a given host requires and how long we judge it will take.

Here are the tiers as they stand, the concept register for the year:

TIER	SCOPE	GENES	HORIZON	FLAGSHIP CONCEPTS
T1	Solanaceae Core	4–5	18–24 mo	Tomato *Volcanic* · Eggplant *Dark Matter* · Tobacco *Laboratory*
T2	Extended Nightshades	5–7	24–36 mo	Potato *Seismic* · Ground Cherry *Celestial* · Tomatillo *Acid* · Petunia *Floral*
T3	Cross-Family	7–10	36–48 mo	Strawberry *Sweet Heat* · Cucumber *Contradiction* · Melon *Solar* · Apple *Forbidden* · Lettuce *Espionage*
T4	Frontier Synthetic	10–12	48–72 mo	Carrot *Deep Earth* · Cabbage/Cauliflower *Nuclear* · Broccoli *Wildfire* · Corn *Storm* · Rice *Dynasty*

The **first tier is the low-hanging frontier** — the plants closest to the pepper in evolutionary terms, where the transfer is easiest and the timeline shortest. Tomato *Volcanic* is the headline: a line of designed heat tomatoes named for the forms of fire — *Ember, Magma, Pyroclast, Caldera, Krakatoa* — walking the same deliberate ladder the flagship peppers walked, in a fruit nobody expects to bite back. Eggplant *Dark Matter* (*Smolder, Nightshade, Eclipse, Obsidian, Event Horizon*) carries the heat into the aubergine's dense flesh; Tobacco *Laboratory* (*Proof of Concept, Benchmark, Threshold*) is the research host, named with the honesty of a tier that exists to prove

the method before it decorates it. These are 4-to-5-gene builds on an 18-to-24-month horizon — near enough to feel real, far enough to respect.



Eggplant — a Tier 1 Solanaceae host for the *Dark Matter* concept line, closest-in on the frontier.

The **middle tiers widen the family and the ambition.** T2 reaches into the extended nightshades — Potato *Seismic*, Ground Cherry *Celestial*, Tomatillo *Acid*, Petunia *Floral* — 5-to-7-gene builds that leave the comfortable core. T3 is the leap that captures the public imagination: **Cross-Family** engineering, the heat and the composites carried clean out of the nightshades into entirely different plant families. Strawberry *Sweet Heat* is the one everyone asks about — *Warm Kiss, Sugar Burn, Candied Fang, Devil's Dessert*, a dessert fruit designed to carry heat and sweetness in one bite, the swicy chord grown rather than mixed. Cucumber *Contradiction*, Melon *Solar*, Apple *Forbidden*, Lettuce *Espionage* round out a tier that is 7-to-10 genes and three-to-four years out, and that would, if it lands, retire forever the idea that a plant's flavor is fixed by its species.



Strawberry — the Tier 3 *Sweet Heat* concept, where the swicy chord would be grown into the fruit rather than mixed into the sauce.

The **fourth tier is the synthetic frontier** — the staple crops, the big genomes, the long horizons: Carrot *Deep Earth*, Cabbage/Cauliflower *Nuclear*, Broccoli *Wildfire*, and the two that matter most for the world's plates, Corn *Storm* and Rice *Dynasty*. These are 10-to-12-gene builds on a four-to-six-year horizon, and they are where the frontier stops being a novelty and starts being a question about the global food supply, because corn and rice are not condiments — they are how most of the planet eats. They are also, deliberately, the hosts marked *Frontier* back in the roster of Chapter 12: Soybean, Sugarcane, Sorghum, the shallow young libraries that

are the first real footholds in exactly this tier. The roadmap and the roster meet here, at the edge, where what could be built and what is beginning to be built are the same plants. That leaves the fifth tier — the one the roadmap holds open and does not yet name — which is the subject of the next chapter, and the one place in this report where the right thing to publish is a discipline rather than a date.

Chapter 18 · What We Will Not Rush — *the tier we have not named.*

The roadmap has five tiers, and this report has named four. The fifth is deliberately open — the furthest horizon, past the synthetic frontier of corn and rice, held without a timeline and without a flagship concept, because the most important thing an annual report can do with the distant future is refuse to pretend it is near. This chapter is short, and it is the most important in Part V, because it is about the discipline that keeps the rest of the report honest: the willingness to map a frontier and then *not rush it.*

The temptation, for a company that can credibly say "capsaicin is not a pepper problem," is to say it about everything at once — to let the roadmap sprawl into a hundred press releases and promise the staple crops next year. We have tried to build the opposite instinct into the shape of the roadmap itself. The tiers are graded by gene count and horizon precisely so that ambition is always paired with an honest timeline; the frontier hosts carry modest construct counts because that is the true state of a young capability and we will not inflate it; and the fifth tier stays unnamed because we would rather under-promise a decade than over-promise a year. An index is an instrument of honesty, and an instrument of honesty is worthless the moment its owner starts rounding the future up.



Rice — a Tier 4 staple on the synthetic frontier, and a reminder that the furthest work is measured in years, not seasons.

There is a deeper reason the far frontier stays open, and it is the through-line of this whole library surfacing for the first time as a stated aim rather than a quiet hum. The furthest tier is not, finally, about heat, or even about flavor. It is about the thing designing flavor keeps turning out to share a method with: **nourishment**. Everything in this report — the restoration of the tomato's glutamate through its own GAD/GDH1 node, the five-of-seven modules that let a plant's own chemistry be tuned rather than overwritten, the insistence that the edits use the plant's own enzymes — points toward a capability that is not only about how loud a food tastes but about how much

good it can be grown to hold. A plant engineered at the level of its metabolism to carry the most flavor it can is, by the same tools, a plant that could be engineered to carry the most nutrition it can, and to make that nutrition most available to the body that eats it. That is the axis beyond the eight. We will not name a date for it, and we will never make a claim about it that the index cannot measure. But it is why the fifth tier stays open: because the frontier we are really walking toward is not a hotter crop. It is food grown, from the protein outward, to feed the body as fully as it can — and that is worth doing slowly, and worth doing right.



PART VI OF VI

The Ledger

The locked count · the standards we hold · the state of flavor, and the year ahead

Chapter 19 · The Numbers, as They Stand — *the locked count.*

A report is only as trustworthy as its ledger, and across a library of this size the single thing that reads as sloppy is drift — a count that says one thing on one page and another thing elsewhere. So this chapter does one job, plainly: it states the numbers of record, locked, the figures every other page in this library resolves against. These are the state of flavor expressed as a single column, and they are the same here as everywhere.

MEASURE	VALUE
Engineered constructs	~19,000
Live / purchasable today	1,000
Total SKUs	78,800
Validated host organisms	42
Sensory axes	8
Batch coefficient of variation	≤7%
Patents filed	3
Peer-reviewed papers (Zenodo, CC BY 4.0)	21

Each figure carries a meaning worth stating once, clearly. The **~19,000 engineered constructs** are the total designed-and-built library, of which the 42-host roster of Chapter 12 accounts for 18,500 validated constructs by direct sum — the remainder being the newest multi-axis work, still in validation. The **1,000 live** are what you can actually buy today: the proven, in-stock, purchasable fraction, a deliberately small and honest number against the larger library of what has been built and the far larger space of what has been designed. The **78,800 SKUs** are the full addressable catalog — every cultivar, format, and configuration the library can express — and the gap between 1,000 live and 78,800 possible is not a discrepancy but a map of the road ahead. **42 organisms, 8 axes, ≤7% CV**: the roster, the dials, the promise. **3 patents filed** protect the method; **21 peer-reviewed papers**, published open-access on Zenodo under CC BY 4.0, put the science on the public record, because a company asking to be believed about flavor should be willing to show its work to anyone who wants to check it.



A designed specimen under isolation — one validated construct among roughly nineteen thousand, one live product among a thousand.

The ledger carries one more kind of entry, and it belongs in a report meant for press and industry: the **standards we hold ourselves to in what we say.** The index measures loudness at a receptor, and that is the whole of what we claim. Every number in this report is a statement about a molecule meeting a cell — an EC_{50}, an activation, a precipitation of salivary protein — and nothing in it is a statement about a body downstream of the tongue. We describe our sweetness work as sugar-free loudness on TAS1R2·R3, never as a benefit to anyone's metabolism. We describe the bitterness axis as a 20-to-65% attenuation of a receptor response, never as making a food "good for you." The flavor compounds the constructs express — capsaicin, menthol, glutamate, the steviol glycosides, the γ-glutamyl peptides — are molecules already present throughout the food supply, which is the honest foundation of their regulatory standing; what is new is the precision with which a plant can be tuned to carry them, not the compounds themselves. And where the library reaches toward functional work — flavor restoration for changed palates, nourishment held in the plant — it is described in **attributes** (satiety, protein-forward, taste retention, the nutrition a food is grown to hold) and never in the language of medicine. No cures. No detox. No health halo. The through-line of nourishment that runs under every book in this library is a *value and an intention*, stated as such, and the index is the discipline that keeps it from ever becoming a claim. That restraint is itself a number in the ledger — the one that is not counted, and the one we would least like to be caught rounding.

Chapter 20 · The State of Flavor — *the year ahead.*

So this is the year, read in full. Eight dials, each one a receptor and a mechanism and an anchored number: a heat channel on a nerve, two flytraps for sweet and savory, a bitter alarm we work to silence, a calcium sensor for richness, a cold channel for a winter with no weather, an electrophile alarm for the sharp, a precipitation of salivary protein for the grip of the finish. Forty-two host organisms carrying 18,500 validated constructs, from the native pepper to the frontier staples just now taking root. Ten flagship warnings climbing to a single eight-axis verdict. A method that starts at the protein and lets the phenotype follow, transferring five of its seven modules to any plant with a sequenced genome. A reproducibility of seven percent where the wild world offers forty-two. And a frontier of sixty-four concepts across five tiers, four of them named and dated, one of them held open on purpose. That is the state of flavor, this first year we have been able to measure it.

What strikes us most, assembling it, is how little of it is really about heat — the axis we are known for, the axis we learned on, the loud one at the top of the ladder. Heat was the proof. The report is about everything heat taught us to build: the tomato's own savor grown back through its own enzymes, the bitterness lifted off a cup of coffee so the rest of it can be tasted, the richness of a slow-made sauce reduced to a number a formulator can specify, the cool of mint decoupled from cold and handed over as a dial. The index did not make flavor smaller by numbering it. It made flavor *legible* — and a thing you can read, you can design, and a thing you can design, you can grow toward its fullest form on purpose.



The library at the hour the light arrives — the whole living operation this report was assembled to measure.

The year ahead, then, in the only terms this report is willing to promise them. The roster will lengthen — the frontier hosts will deepen, and we expect to add validated organisms to the 42, each one a new native speaker of the eight dials, each one tracked here next year against this year's baseline. The composites will grow — the M2 and M3 libraries will widen, and the full-spectrum line will refine the delivery that lets eight pathways share one tissue. The

frontier's first tier will move from concept toward construct, and if the tomato's story repeats — as the method says it should — a fruit that was never supposed to bite back will. And underneath all of it, unhurried and unclaimed, the work that the furthest tier is really for will continue: learning to grow a plant, from its metabolism outward, to hold the most nourishment it can carry and to give it up most fully to the body that eats it. We will not name a date for that. We will only keep the lamp on, and keep measuring, and tell you the truth each year about what moved.

That is the discipline that makes an annual report worth publishing: a fixed instrument, an honest ledger, and the patience to come back in a year and read the same dials again. Flavor has a number now. This is the first reading. We will see you at the next one — at the table, under the one good lamp, cutting the first fruit of something we spent a season arguing over, and this time able to say, exactly, by how much.

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