

THE SCOVILLE SPLICE LIBRARY · VOLUME VII

Sweetness

All of the sweetness. *None of the sugar.*

A receptor-level account of the sweetest axis — how TAS1R2·TAS1R3 reads a sugar, and how we teach stevia, monk fruit, sugarcane, the tomato, and a strawberry that bites back to speak sweetness without a gram of added sugar.

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

Contents

01 Opening essay · The oldest sweetness

02 Chapter 1 · The sweet receptor — two proteins, one craving.

03 Chapter 2 · From molecule to nerve — how the tongue says yes.

04 Chapter 3 · The mathematics of loud — why sweetness needs a number.

05 Chapter 4 · What sugar actually does — the five jobs nobody tastes.

06 Chapter 5 · The potency ladder — every sweetener, ranked.

07 Chapter 6 · The off-notes — bitterness rides with sweetness.

08 Chapter 7 · Body without sugar — building the floor back.

09 Chapter 8 · Stevia — the leaf that out-sweetens sugar three hundred times.

10 Chapter 9 · Monk fruit — the rounder sweetness.

11 Chapter 10 · Sugarcane — sweetness at the source, reconsidered.

12 Chapter 11 · The sweet proteins — when sweetness is a protein, not a sugar.

13 Chapter 12 · The tomato — sweeter without a grain more sugar.

14 Chapter 13 · Sweet Heat — the strawberry that bites back.

15 Chapter 14 · Solar, Forbidden, Celestial — the dessert frontier.

16 Chapter 15 · Phoenix — sweetness inside the full spectrum.

17 Chapter 16 · The reformulation playbook — from a sugar brief to a number.

18 Chapter 17 · Formulations — the sweetness, built.

19 Chapter 18 · The nourishment through-line — sweetness, given back.

Opening essay · The oldest sweetness

Sweetness is the first flavor any of us ever loved. Before language, before a single remembered meal, there was the sweetness of milk — roughly seven percent lactose, warm, given — and the body already knew to say *yes* to it. The newborn who has tasted nothing else will turn toward a drop of sugar water and away from a drop of bitter; the preference is older than memory, written into us because for most of our history a sweet taste meant ripe, meant safe, meant calories a body could use. We did not learn to love sweetness. We arrived loving it.

That is the fact to hold onto, because everything that has gone wrong with sweetness in the last century is downstream of it. We are built to chase this flavor, and we became, for the first time in history, able to manufacture its carrier — refined sucrose — in unlimited amounts and pour it into everything. The craving stayed ancient and true. Only the supply changed. And so the most innocent of the eight axes became the one people feel most guilty about, the one they ration, the one they are told to fear.



The question that became this book came up at our table one night: if sweetness is something the tongue reads, why does it have to arrive inside sugar?

Our family did not set out to argue with sugar. We set out, as we always do, to measure something we loved and could not yet name. And when we put sweetness on the bench and looked at it honestly, we found what the other seven axes had already taught us: *sweetness is not a sugar. Sweetness is a receptor event.* It is two proteins on the surface of a taste cell, folded together like two cupped hands, closing on a molecule and reporting the catch to the brain. Sugar is one thing that fits those hands. It is not the only thing, and on a per-molecule basis it is one of the *weakest* things — a steviol glycoside from a leaf speaks to that receptor three hundred times more loudly than sucrose does, gram for gram.

This is the whole opening the Iyengar Flavor Index gives us. Once sweetness is a number anchored to a receptor rather than to a sugar, the two come apart in your hands. You can ask a plant for the *taste* and leave the *carrier* behind. You can grow a tomato that reads as markedly sweeter with not one additional grain of sugar in it. You can design a strawberry that opens sweet and finishes with heat. You can take the loudest sweetness in nature — a small protein from a West African berry — and ask a crop to make it for you.

The receptor, not the sugar

Hold a ripe fig on the tongue and the sweetness seems to live *in* the fruit, as surely as the color lives in its skin. It does not. The sweetness is happening a few micrometers away, on the microvilli of your taste cells, and the fig is only the delivery. What actually reads sweet is a single receptor: **TAS1R2 and TAS1R3**, two proteins that must find each other and pair before either can do anything at all. Alone, each is deaf. Bound together into an obligate heterodimer, they become the one instrument the human body uses to hear sweetness — every sugar, every sweetener, every sweet protein you will ever taste passes through this one pair of hands.

They belong to the class C G-protein-coupled receptors, the same structural family as the savory receptor and the calcium-sensing receptor — the deliberative, slow-reading receptors, built for nutrients rather than alarms. Each subunit carries an enormous extracellular **Venus flytrap domain**, a bilobed clamshell that hinges shut around a ligand. A cysteine-rich collar connects that clamshell to a seven-pass transmembrane bundle threaded through the cell membrane. The architecture matters because it is *why* sweetness can be unbound from sugar. There is not one keyhole here; there are several:

- **Small sweeteners** — sucrose, fructose, the steviol glycosides, the mogrosides of monk fruit — settle into the cleft of the TAS1R2 flytrap and pull it closed.
- **Sweet proteins** — brazzein, thaumatin, monellin — are far too large to fit inside. They dock against the *outside* of the already-folded flytrap, a wedge pressed to a wall, engaging a surface a hundred times broader than a sugar ever touches. This is the quiet reason a protein can be thousands of times louder than a sugar per molecule: it is not filling a small pocket, it is gripping the whole hand.
- **A third class of molecules** binds the transmembrane bundle of TAS1R3 instead, which is why the receptor can be modulated — nudged louder or quieter — without ever touching the sugar site.

The sequence after binding is a cascade, and it is worth walking a single sweet molecule all the way to the brain, because every step is a place the Index can read. The flytrap shuts. The conformational change travels down the cysteine collar and twists the transmembrane bundle, which releases its bound G-protein, **gustducin**, on the inside of the cell. Gustducin activates phospholipase C-β2; that enzyme cleaves a membrane lipid into the second messenger **IP_3**; IP_3 opens calcium stores inside the cell; the rising calcium throws open **TRPM5**, a channel that depolarizes the membrane; depolarization opens a **CALHM1/3** pore that vents **ATP**

out of the cell; the ATP lands on the gustatory nerve fiber waiting just outside, which carries the word inward — brainstem, thalamus, gustatory cortex — until, a few hundred milliseconds after the fig touched your tongue, you *know* it is sweet.

Nowhere in that chain does the receptor ask how many calories it just reported. It cannot. It counts occupancy, not energy. That single blind spot is the entire opening: anything that fits the flytrap and holds it shut reads as full sweetness, whether it carried four calories per gram or none. Sweetness was never a property of sugar. It is a property of the lock, and the lock can be opened by a great many keys.

What EC50 actually measures

To design against a receptor you first have to agree on what *loud* means, precisely, in numbers. We measure it the way pharmacology has measured receptor agonists for a century: with a dose-response curve. Take the TAS1R2·TAS1R3 receptor in a cell assay, expose it to a rising concentration of a sweetener, and plot how hard it fires. The response is not a straight line. It is a sigmoid — nothing, nothing, then a steep climb as the binding sites fill, then a plateau at saturation where every receptor is already working and more sweetener changes nothing. The curve obeys the Hill relation:

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

Here $[L]$ is the concentration of sweetener, R_{max} is the receptor's maximum possible firing, and n is the Hill coefficient — a measure of how steeply the curve rises, which for the sweet receptor sits a little above 1 because the multi-site sweet proteins bind with mild cooperativity. The one term that matters most is **EC_{50}**: the concentration at which the response reaches exactly half of R_{max}. It is the midpoint of the climb, the single most honest one-number summary of a molecule's potency. A low EC_{50} means the receptor is already half-roused at a whisper of the compound — the molecule is loud. A high EC_{50} means you must flood the tongue before it answers — the molecule is quiet. Sugar, for all we love it, is quiet.

The Index does one thing to EC_{50}: it inverts it against a fixed reference and scales it, so that *louder* becomes a *higher* score rather than a lower one.

$$\text{IFI-S(candidate)} = 1{,}000 \cdot EC_{50}(\text{reference}) \,/\, EC_{50}(\text{candidate})$$

The reference is a ten-percent sucrose solution, anchored by definition at **1,000**. Ten percent weight-per-volume is one hundred grams of sucrose in a liter of water — the sweetness of a soft drink, a concentration everyone's tongue knows. Because sucrose's molar mass is 342.3 grams per mole, that reference is also 0.292 moles per liter, about **292 millimolar** — a useful figure to keep in a back pocket. Anchoring it at 1,000 is a convention, nothing more, and it is a convention that pays off immediately: since the formula is a *ratio* of two EC_{50} values measured in the same assay, the absolute units cancel. We never have to defend a particular micromolar reading of

sucrose. We only have to measure every candidate the same way the reference was measured, and divide. What survives the division is pure relative loudness.

We measure EC_{50} in mass concentration — grams per liter — rather than molar, and that choice is deliberate. A beverage formulator doses sugar by weight, reads a nutrition panel in grams, and buys sweetener by the kilo. An Index that speaks in grams can be laid directly beside a sugar spec sheet and read without conversion. The molar picture is richer, and we will reach for it, but the working currency of IFI-S is the gram, because that is the currency of the kitchen and the factory floor.

The ladder of loudness

Now the formula earns its keep. Set sucrose at the bottom rung by definition and let every other sweetener fall where its EC_{50} puts it.

Begin with the canonical figure: a steviol glycoside is three hundred times sweeter than sucrose, gram for gram. Translated into EC_{50}, that means you need one three-hundredth of the concentration to reach the same half-maximal firing: $EC_{50}(\text{stevia}) = EC_{50}(\text{sucrose}) / 300$. Feed that into the Index and the arithmetic is immediate —

$$\text{IFI-S}(\text{stevia}) = 1{,}000 \cdot EC_{50}(\text{sucrose}) \,/\, \left(EC_{50}(\text{sucrose}) / 300 \right) = 1{,}000 \cdot 300 = \mathbf{300{,}000}$$

Three hundred thousand. The same calculation, run on the published per-gram potency of each natural sweetener, produces the ladder below. We show the potency the literature reports and the IFI-S the formula implies from it, so the chain from measurement to score is never hidden.

SWEETENER	SOURCE	PER-GRAM POTENCY VS SUCROSE	IMPLIED IFI-S
Sucrose *(reference)*	cane / beet	1×	1,000
Mogroside V	monk fruit	~250×	~250,000
Steviol glycoside (Reb A)	stevia leaf	~300×	300,000
Brazzein	West African berry	~800×	~800,000
Thaumatin	West African aril	~2,000×	~2,000,000

The reference sweetener, read as a pure stimulus, lights exactly one of the eight axes and nothing else — a single clean note:

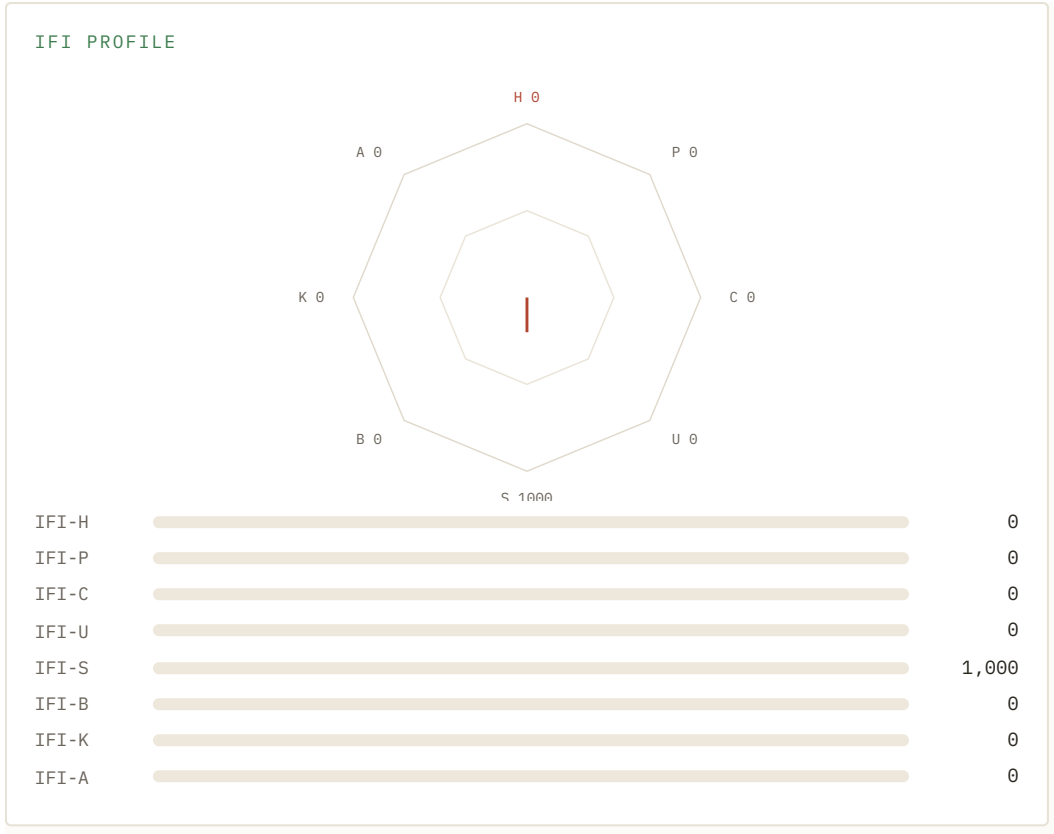

There is a subtlety in that ladder worth drawing out, because it is where the per-molecule story hides. Potency by *weight* and potency by *molecule* are not the same number, and they diverge exactly in proportion to molecular mass. The relationship is clean: molar potency equals weight potency times the ratio of molar masses. Reb A weighs about 967 grams per mole against sucrose's 342.3, so its molecular potency is 300 × (967 / 342.3) ≈ **850×** per molecule even though it is 300× per gram. Brazzein, a protein of roughly 6,473 grams per mole, is near 800× by weight — but 800 × (6,473 / 342.3) ≈ **15,000×** per molecule. Fifteen thousand. A single brazzein molecule, pressed against the outside of the flytrap, does the work of fifteen thousand molecules of sugar. That is what "the loudest sweetness in nature" means, made exact. And it is precisely why the Index keeps its books in grams for the formulator while we keep the molar figure for the biology: the receptor counts molecules, but the invoice counts grams, and an honest standard has to serve both readers on one page.

The ceiling the tongue sets

A number like three hundred thousand invites a misreading, and we should clear it before it hardens: stevia is not three hundred thousand times *sweeter-tasting* than sugar. Nothing is. The Index measures loudness per gram — how little of a molecule it takes to half-rouse the receptor — not the loudest the receptor can be driven to shout. Those are different quantities, and the same Hill curve holds both.

Look again at the sigmoid. Response climbs with concentration until every receptor is occupied, then stops; the plateau `R_max` is a wall the tongue builds, not the molecule. With a Hill coefficient near 1, the working arithmetic is tidy. At the EC_{50} the receptor fires at half of maximum by definition. Push to ten times the EC_{50} and the response is 10 / (1 + 10) ≈ 91 percent of maximum. Drop to one-tenth of the EC_{50} and it falls to 0.1 / 1.1 ≈ 9 percent. The entire useful dynamic range of sweetness — from barely-there to nearly saturated — spans only about two orders of magnitude in concentration. Past the

top of that climb, more sweetener buys almost nothing: the receptors are already full, and the brain has no louder word for sweet than the one it is already saying.

This is also, quietly, why the reference was set at ten percent and not thirty. Sucrose perception saturates somewhere above a twenty-percent solution, up on the flat of the curve where small changes in concentration stop changing the taste. Anchor your standard there and it goes numb — a ruler whose marks have run together. A ten-percent solution sits lower, on the steep responsive slope, where a change in sweetness still moves the number. We anchored the Index where the tongue is most sensitive, not where it is loudest, because an anchor has to *discriminate*, not just register.

So what does a potent sweetener actually buy? Not a sweetness three hundred thousand times taller — the ceiling forbids it — but the same ceiling *reached at a trace dose*. A high-IFI-S molecule arrives at the steep part of the curve while still a whisper of mass on the tongue; it climbs the same wall sugar climbs, only starting from a thousandth of the grams. That is the honest reading of the whole ladder. IFI-S ranks how *cheaply*, in grams, a sweetener reaches full sweetness — never how far past sugar's ceiling it can go, because there is no past. Two sweeteners that both saturate the receptor taste equally sweet at the summit; they differ only in how many grams it took to get there, how fast the climb began, what tail it left, and whether a bitter co-signal rode along. The ceiling is shared. Everything a formulation actually fights over happens on the slope beneath it.

Where sweetness is made

A number on a ladder is a target; a pathway is how you reach it. The reason we can ask a crop to carry any rung of that ladder is that nature already builds every one of these molecules with enzymes, and enzymes are what our platform writes.



Sugarcane in cross-section — the carrier of the reference itself. Ten percent sucrose in water is the 1,000 on every sweetness scorecard; the plant simply concentrates it.

Take stevia, the rung we know best. *Stevia rebaudiana* makes its sweetness through the plastid's terpenoid route: isoprene units assembled into the diterpene backbone *steviol*, which is then dressed in sugars by a family of **UDP-glycosyltransferases** — UGT85C2, UGT74G1, UGT76G1, UGT91D2. Each enzyme hangs a glucose on a specific hydroxyl, and the *pattern* of those decorations is everything. Stevioside, the simplest, carries a harsh edge.

Rebaudioside A is cleaner. Rebaudioside D and the heavily glycosylated Rebaudioside M are cleaner still, nearly sugar-like. The off-taste that has haunted every stevia soda for a generation — that lingering licorice-and-metal tail — is not the sweetness at all. It is a *second* signal: the lightly glycosylated glycosides also bind the bitter receptors TAS2R4 and TAS2R14, firing a bitterness the sweetness cannot mask. The fix is enzymatic, not cosmetic. Adding glucose at the right positions — tuning which UGT runs hardest — shifts the molecule off the bitter receptors while keeping it loud on the sweet one. This is the whole game stated plainly: sweetness and its off-taste are two different receptors reading two different features of the same molecule, and the glycosyltransferases let us separate them.

Monk fruit tells the same story in a different dialect: its mogrosides are cucurbitane triterpenoids, glycosylated by the plant's own UGTs, and mogroside V is the loud, clean one. Brazzein tells the simplest story of all — it is not a decorated small molecule but a single 54-residue polypeptide, which means a single gene. It folds into a compact, disulfide-stapled knot that shrugs off heat and acid, holding its sweetness through baking and boiling where thaumatin would denature and fall silent. A crop that expresses brazzein is, in effect, a field that grows the loudest sweetener known and ships it already folded.

Four of our forty-two platform host organisms speak the sweetness axis as a primary voice. Their construct counts are the measure of how deeply each has been worked.

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY	SWEETNESS ROUTE
Tomato	2,040	S / U / F / M	Easy	fruit sugars + sweet-protein expression
Stevia	910	S / F / M	Moderate	steviol glycosides, UGT-tuned
Sugarcane	200	S	Frontier	sucrose — the reference itself
Monk Fruit	80	S	Hard	mogrosides (mogroside V)

Tomato sits at the top of that list with 2,040 constructs not because a tomato is the sweetest thing we grow, but because it is the most *instructive* — a fruit people believe they know, where a restored sweetness is felt as the fruit simply tasting the way memory promised it would. Sugarcane sits near the bottom, a frontier host with only 200 constructs, for the opposite reason: it already makes the reference molecule to perfection. There is less to teach a plant that invented the benchmark.

The reformulation problem

Everything above is clean on the bench. The honest chapter is what happens when a formulator tries to take sugar *out* of a real product, because that is where the Index stops being a scoreboard and becomes an engineering constraint.

Start with the arithmetic, which is encouraging. A 355-milliliter can of regular cola carries about 39 grams of sugar. Replace it, gram for gram of sweetness, with Reb A at 300× potency:

```
39 g ÷ 300 = 0.13 g = 130 milligrams of steviol glycoside per can
```

A sweetness that cost thirty-nine grams now costs a tenth of one. In molar terms the drop is starker still: 39 grams of sucrose is about 114 millimoles of molecules; 130 milligrams of Reb A is about 0.13 millimoles. You have removed better than eight hundred molecules of sweetener out of every eight hundred and one and kept the sweetness. If sweetness were all sugar did, the problem would be solved and this would be a short chapter.

It is not all sugar does. Sugar is a *bulk* ingredient masquerading as a flavor. Those thirty-nine grams were also doing the body of the drink — the viscosity that coats the tongue, the weight that reads as "full" rather than "watery." They were depressing the freezing point in a popsicle, feeding the yeast in a bread, browning the crust in a bake, and holding water in a frosting so it would not go dry. Pull the sugar and the sweetness can be matched exactly while every one of those invisible jobs goes unfilled. A diet soda formulated only with a potent sweetener tastes *thin* — and that thinness is not a sweetness failure the Index would catch; it is a mouthfeel failure living on a different axis entirely.

There is a second failure the Index *does* illuminate: the shape of sweetness in time. Sucrose arrives fast and leaves clean — a square pulse. The high-potency glycosides arrive a beat late and linger a beat long, and in that lingering tail the bitter co-signal has time to surface. Two sweeteners can share an identical IFI-S and eat nothing alike, because IFI-S reports loudness at the half-maximal point, not the onset time or the decay. This is a genuine limit of a single-number axis, and we state it rather than hide it. The working answer is a blend: a fast front from a trace of real sugar or a sweet protein, a potent mid from the glycosides, a bulking agent like erythritol to rebuild the body, each component chosen to patch a different part of the curve. The Index scores the destination; building the temporal road to it is still craft.



A tomato in cross-section — sweetness restored not by adding sugar but by tuning what the fruit already builds. The receptor reads the fruit as sweeter; the sugar content need not move at all.

This is exactly why a *designed fruit* is a different proposition from a sweetener dropped in water. When we tune a tomato toward sweetness, we are not pouring an extract into a thin solution — we are working inside a whole-food matrix that already carries its own

water-binding, its own acids, its own body. Suppose a formulator's target is a tomato that reads 1.6× as sweet as the reference solution. By the formula, that is an IFI-S of 1,000 × 1.6 = **1,600** — a fruit whose expressed sweetness sits well above a soft drink's while the fruit remains, in every structural respect, a tomato. The sweetness rises; the matrix that makes it eat like food stays intact. That is the move the bench extract can never make.

Consistency is the product

For anyone who formulates for a living, there is a number quieter than potency that matters more, and it is the one we are proudest of. A sweetener is only useful if it is the *same* sweetener every harvest. A wild crop is not. Sun, soil, water, and season push its expressed sweetness around by a batch-to-batch coefficient of variation near **42 percent** — a swing so wide that two tomatoes from the same field can read nearly twice apart. Our designed cultivars hold that variation to **≤7 percent**, the same precision tolerance that runs through every axis of the platform.

Read what that does to a sweetness target. Take the designed tomato above, aimed at IFI-S 1,600, and put a one-sigma band around it under each regime:

MEASURE	WILD CROP (CV 42%)	DESIGNED CULTIVAR (CV ≤7%)
Target IFI-S	1,600	1,600
1σ band (±)	±672	±112
Typical range	928 – 2,272	1,488 – 1,712

A formulator who prints a sweetness claim on a label cannot work inside the left column. A spread of 928 to 2,272 is the difference between a drink that tastes under-sweetened and one that tastes cloying, arriving at random, can to can. The right column — a range tighter than ±8 percent — is a *specification*. That narrowing is not a flourish on top of the flavor; for the person buying the ingredient, the narrowing *is* the product. Loudness gets you onto the shelf. Reproducibility is what lets you stay.

Sweet heat, and the whole chord

We will close where the axis stops behaving like a solo. For most of this chapter sweetness has stood alone, one clean note at 1,000, the single-axis stimulus of a sugar solution. Food is rarely so monastic. The tongue reads eight axes at once, and the most interesting designs light several together — which is how a pepper bred for the far edge of heat can also carry a genuine sweetness our platform predicts beneath the fire.

Consider Phoenix — SS-0096, a *Capsicum chinense* our model predicts at approximately 2.95 million Scoville Heat Units, an IFI-H of 4,200, and expression on all eight axes at once. It is, by design, the highest-heat cultivar in our catalog. And yet read down its full profile and the sweetness axis is lit at **1,100** — a cultivar designed so that, beneath a predicted heat approaching three million Scoville units, it reads slightly sweeter than a ten-percent sugar solution.

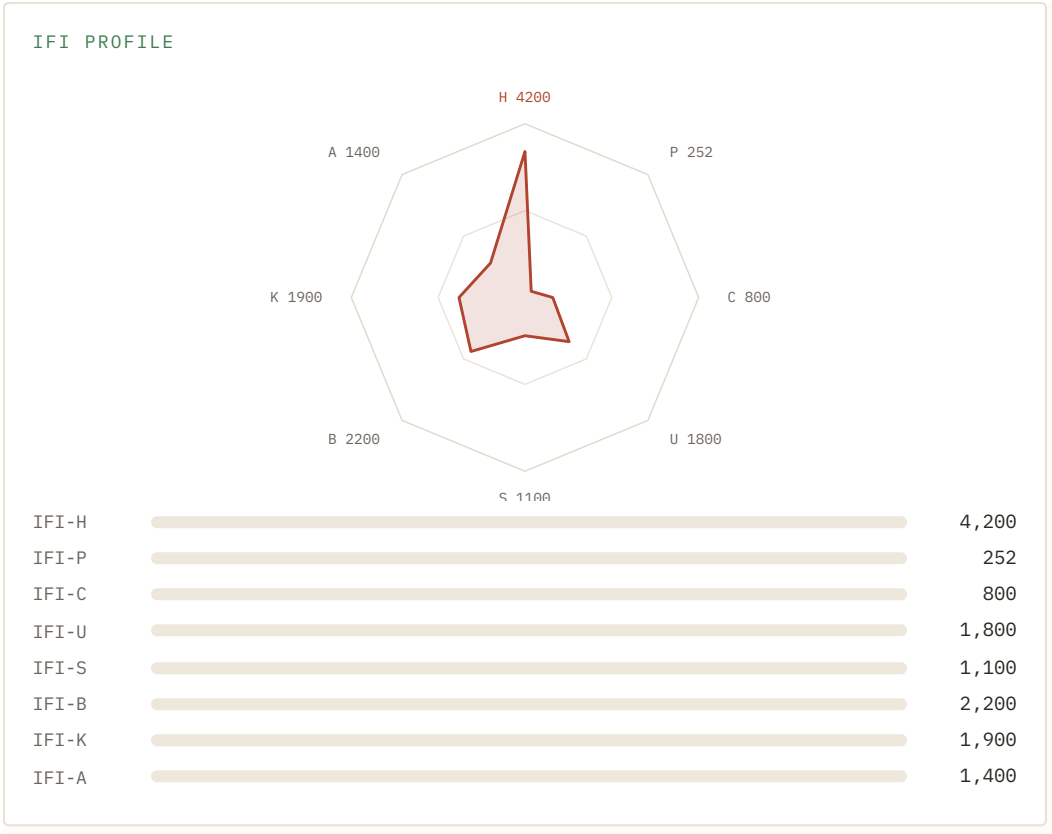

That one line — `S: 1100` sitting calm inside a wall of heat — is the whole thesis of the axis in a single number. Sweetness is not borrowed from sugar, or it could not survive in a fruit this fierce; it is read from a receptor, and a receptor can be addressed independently of every other signal the plant sends. Decouple taste from carrier and the combinations open without limit: a strawberry that opens sweet on the front of the tongue and finishes with a clean warmth; a tomato lifted to 1,600 with no change to its sugars; a beverage sweetened by a trace of a folded protein a field grew for the purpose. Each is the same quiet act, performed on the one pair of cupped hands the body uses to hear sweetness.

We wrote this for the people who have to solve sweetness for a living — the formulator staring down a sugar-reduction target, the clean-label founder, the team that has tried eleven stevia blends and tasted licorice in all of them — and for anyone ever told that the flavor they have loved longest is the one they should trust least. The science here is exact and shown, because numbers are a form of respect. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. But the argument underneath is simple and warm, and it is the one our family makes about every plant we touch: food should give the body pleasure and nourishment in the same bite, with nothing to apologize for. Sweetness was never the problem. We are giving it back.



PART I OF V

The Receptor

TAS1R2 · TAS1R3 · the Venus flytrap · gustducin · TRPM5 · the number

Chapter 1 · The sweet receptor — *two proteins, one craving.*

Put a single crystal of sugar on the tip of your tongue and hold still. The sweetness that arrives a half-second later is not happening in the sugar. It is happening in a protein — two proteins, in fact, named TAS1R2 and TAS1R3, embedded side by side in the membrane of a slender taste-receptor cell bundled into a taste bud. They belong to a family called the class C G-protein-coupled receptors, the same architectural family that reads glutamate at your synapses and free calcium in your bloodstream. They are ancient, elegant machines, and understanding their shape is the entire foundation of designing their input.

We keep returning, in our family, to that half-second of silence between the crystal and the sweetness — because everything we build lives inside it. The crystal is only a cause. The sweetness is an event in a protein, and a protein can be moved in more ways than one. To learn the ways is to learn that the sugar was never the point; the motion of the protein was. This chapter is the anatomy of that motion, and the arithmetic that lets us measure it.

Three stacked pieces, two proteins

Each of the two subunits is built in three stacked pieces. At the top, reaching out into the saliva, is a large bilobed domain shaped — genuinely — like a Venus flytrap: two lobes hinged together, open and waiting. Below it sits a short, rigid cysteine-rich domain, stitched together by disulfide bonds, which acts as a driveshaft. And anchored in the membrane is the seven-transmembrane bundle, the classic GPCR coil that, when the machine fires, grips a G-protein on the inside of the cell and sets the signal loose. TAS1R2 and TAS1R3 do not work alone; they are an obligate heterodimer, two different proteins that must fold together to function. On their own, each is deaf. Clasped together, they taste.

The two subunits are not symmetric, and the asymmetry decides the design. **TAS1R2 is the sweet-specifying subunit** — its flytrap holds the principal sugar site, and it carries the handful of human-specific residues that let us register aspartame and the sweet proteins at all. **TAS1R3 is the shared subunit** — the very same protein that pairs instead with TAS1R1 to build the *umami* receptor. This sharing is not trivia. It is the hardware reason sweetness and savor are cousins, the reason a single mutation in TAS1R3 can shift both tastes at once, and the reason the sweetness and umami of a crop must be designed as one problem rather than two. The shared subunit is a wire strung between two senses; you cannot pull one end without the other moving. That coupling becomes decisive the moment a crop is asked to carry both — the tomato is the obvious case, where sweetness restored and savor restored turn out to be a single engineering conversation.

What makes this receptor the richest design target of all eight axes is that it does not have one door. It has several, and they open to different keys.

Several doors, and the keys that fit them

The natural sugars — sucrose, glucose, fructose — and many small synthetic sweeteners, including aspartame, neotame, and sucralose, drop into the open cleft of the TAS1R2 flytrap and trigger it to close. The closure *is* the signal: the two lobes snap shut around the molecule like a hand closing on a coin, and that mechanical motion travels down through the cysteine-rich driveshaft into the transmembrane bundle. A second, entirely separate site lives down in the seven-transmembrane domain of TAS1R3 — and this is where cyclamate binds, and where lactisole, a sweetness *inhibitor*, wedges in to jam the machine shut. And then there is a third mode, the strangest and most beautiful: the sweet proteins.

Brazzein, thaumatin, monellin — these are not small molecules at all but folded proteins, some of them far larger than any cleft could swallow. Brazzein is a mere fifty-four amino acids, about 6.5 kilodaltons; thaumatin is nearly 22. They cannot fit inside the flytrap. Instead, by the best current model, they dock against the *outside* of the closed flytrap domain and the cysteine-rich collar like a wedge jammed under a door, stabilizing the active, shut conformation from without. A protein a thousand times too large to be a key turns out to be a better one than sugar, because it does not need to be caught — it simply holds the trap closed.

BINDING SITE	WHERE	WHAT BINDS THERE
VFT cleft, TAS1R2	top of the flytrap	`sucrose, glucose, fructose, aspartame, neotame, sucralose`
VFT surface / CRD	outer face of the closed trap	`brazzein, thaumatin, monellin (the "wedge" proteins)`
7TM domain, TAS1R3	buried in the membrane	`cyclamate (agonist); lactisole, gymnemic acid (inhibitors)`

This multiplicity is not a curiosity. It is the reason sweetness is *designable* in a way few flavors are. Because there are several partially independent ways into the receptor, two sweeteners binding at two different sites can add their effects — and sometimes more than add them, producing the synergy that formulators have exploited for decades without always knowing its structural cause. It is also why species differ so sharply. A rat cannot taste aspartame or thaumatin at all; a few residues in the human TAS1R2 flytrap, different from the rodent's, are the entire reason those molecules read as sweet to us. When we design a sweetener, or a crop that makes its own, we design against the *human* receptor's specific geometry, and we score against it. The human sequence is the only one that matters at the table.

Closure becomes current

Hold the trap shut, and nothing is tasted yet — the closure is only the first domino. What follows is a cascade worth walking step by step, because every amplification and every failure mode we care about lives somewhere along it.

The shut flytrap reorients the seven-transmembrane bundle, and the bundle grips the receptor's G-protein on the cytoplasmic face — in taste cells, a heterotrimer built around **α-gustducin** with a $\beta_3\gamma_{13}$ partner. That grip catalyzes the exchange of GDP for GTP on the α-subunit, and the trimer splits. The freed βγ pair activates **phospholipase C-β2 (PLCβ2)**, an enzyme that cleaves the membrane lipid PIP_2 into two messengers: **IP_3** and diacylglycerol. IP_3 diffuses to the endoplasmic reticulum and opens the **type-3 IP_3 receptor**, and the cell's internal calcium stores empty into the cytosol. That rising calcium opens **TRPM5**, a calcium-activated channel selective for sodium and potassium; sodium floods in and the cell depolarizes. The depolarization, together with the calcium, throws open **CALHM1/CALHM3**, a voltage-gated channel with a pore wide enough to pass ATP — and the taste cell, which has no classical synapse, releases ATP straight through it. That ATP lands on **P2X2/P2X3** receptors on the gustatory nerve fiber waiting beneath the cell, and only now, six proteins downstream of the crystal, does a nerve finally fire — up the chorda tympani and glossopharyngeal nerves to the nucleus of the solitary tract, to the thalamus, to the gustatory cortex folded in the insula.

Six stages, each a protein, each a place where the signal can be made louder or quieter. A sweetener that only barely closes the trap can still read as intensely sweet if the cascade beneath it has high gain; a molecule that closes the trap decisively can read as faint if the cell is adapted and its stores are low. This is why we insist on measuring loudness at the *receptor*, with a number, rather than inferring it from how much sweetener a food contains. Content is upstream of the only thing the tongue reports.

The trap is an equilibrium, not a switch

There is a mechanical subtlety in that closure worth dwelling on, because our whole enhancement strategy lives inside it. The flytrap is not a simple on–off switch but a *conformational equilibrium* — even with no ligand at all, the two lobes flicker between open and closed, spending most of their time open. A sweetener does not force the trap shut so much as *bias* the equilibrium toward the closed, active state, and the more strongly it stabilizes that state, the louder the signal. Potency and binding are therefore not the same thing, and the gap between them is a door sugar never used.

A molecule that binds the flytrap *without itself triggering closure*, yet makes closure easier for a second molecule, is a positive allosteric modulator — a PAM. Such molecules are real, discovered by high-throughput screening and already in commercial sweet-enhancer use: they amplify the sweetness of a little sugar while adding none of their own. We can write their effect exactly. In the simple allosteric model, a modulator present at concentration [B] with cooperativity factor α shifts the agonist's half-maximal concentration from EC_{50} to

```
EC50(observed) = EC50 · (1 + [B]/K_B) / (1 + α[B]/K_B)
```

and at saturating modulator this collapses to EC_{50} / α. Because the Index scores loudness as the inverse of EC_{50}, a PAM with α = 4 multiplies the sweetness score of the *same sugar* by four. A fruit whose own sugars read only a quarter of the reference can be made to read the full reference — not by adding sugar, but by adding a modulator that changes how the receptor moves. Several of our enhancement constructs aim exactly here: a crop that makes its own allosteric sweet-enhancer, so its modest sugars read loud. It is the purest statement of this book's thesis at the level of the protein itself — *more sweetness, not more sugar* — achieved by changing the machine's motion rather than flooding it with ligand.



Sugarcane in the controlled chamber. For ten thousand years sweetness meant this: a weak key — sucrose, a mediocre ligand — used in enormous bulk, because bulk was the only option anyone had.

Giving sweetness a number

To design against a receptor you must first be able to score it, and the score is built on one quantity: **EC_{50}**, the concentration of a stimulus at which the receptor reaches half of its maximum response. A full dose–response curve for a sweetener follows a saturating, sigmoid law — the Hill equation,

```
response = R_max · [L]ⁿ / (EC₅₀ⁿ + [L]ⁿ)
```

where [L] is the sweetener concentration and n is the Hill coefficient, the steepness of the rise (for the sweet receptor, typically near 1 to 1.5). We anchor the axis at EC_{50} and not at threshold or at saturation for a deliberate reason: the half-maximal point is where the curve is steepest and most reproducible, least distorted by the noise of barely-detectable traces at the bottom or by the flat plateau at the top. It is the single concentration you can measure again and again and get the same answer.

From EC_{50} the Iyengar Flavor Index defines sweetness as a ratio against a fixed reference:

```
IFI(c) = 1000 · EC₅₀(reference) / EC₅₀(candidate)
```

We anchor the sweetness axis to sucrose — the sugar the species has eaten for ten thousand years — so that sucrose, and a reference sucrose solution, score exactly **1,000** by convention. A candidate with half the EC_{50} of sucrose is twice as loud and scores 2,000; a candidate needing ten times the concentration is a tenth as loud and scores 100. The number is loudness, not content, and it reads on the

same page as every other axis. Each axis carries its own stated reference constant that defines IFI = 1,000 for that dimension; the reference points differ across axes, so axis scores describe loudness within their own dimension rather than on a single shared scale.

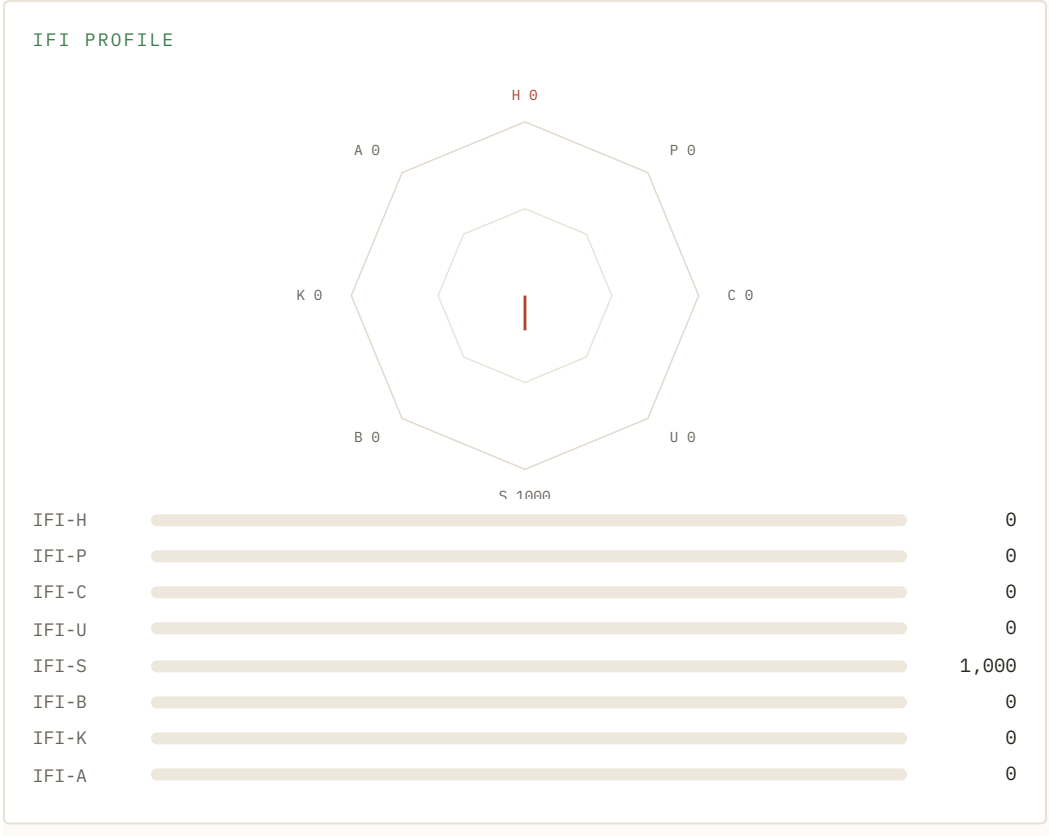

That is the reference itself: sucrose closes the sweet flytrap and touches nothing else — not the heat channel, not the bitter receptors, not the calcium-sensing kokumi machinery — so it scores 1,000 on sweetness and, honestly, zero everywhere else. It is the quietest possible profile, and the entire history of eating has been spent pouring it in by the spoonful to make it loud.

The ladder of loudness

Here the arithmetic earns its keep, and it hides a genuine subtlety worth doing carefully. The number food scientists tabulate — "aspartame is 180 times sweeter than sugar" — is a ratio *by weight*: equal sweetness at one-hundred-eightieth the mass. But the receptor does not weigh anything. It counts molecules. It responds to *molar* concentration, and molecules differ in mass. To convert a by-weight potency N into the molar EC_{50} ratio the Index needs, we carry the molecular weights:

> $EC_{50}(\text{sucrose}) / EC_{50}(c) = N \cdot MW(c) / MW(\text{sucrose})$, so $IFI(c) = 1000 \cdot N \cdot MW(c) / MW(\text{sucrose})$

Work aspartame (MW 294.3, N = 180) against sucrose (MW 342.3): the molar potency is 180 × 294.3 / 342.3 = 154.8, so its raw sweetness score is about **154,800**. Sucralose (MW 397.6, N = 600) comes to 600 × 397.6 / 342.3 = 697, a score near **697,000**. Steviol glycoside — stevioside, a plant molecule, MW 804.9, N ≈ 250 — is heavy, so its molar loudness runs ahead of its by-weight number: 250 × 804.9 / 342.3 = 588, a score near **588,000**. Neotame crosses into the millions.

SWEETENER	MW (G/MOL)	SWEETNESS ×SUCROSE (BY WEIGHT)	POTENCY ×SUCROSE (BY MOLES)	RAW IFI-S
Sucrose	342.3	1	1.0	1,000
Glucose	180.2	0.75	0.39	395
Fructose	180.2	1.5	0.79	790
Aspartame	294.3	180	154.8	154,800
Stevioside	804.9	250	587.8	588,000
Sucralose	397.6	600	697.0	697,000
Neotame	378.5	8,000	8,845	8,845,000

Two honest cautions ride along with that table. First, these are *raw molecular* scores, and the Index caps each axis at a normalization maximum of **5,000** for the radar scorecard — because a profile is meant to compare *foods as eaten*, where concentration tames potency, not to show the millionfold span between pure molecules. A drop of neotame and a bowl of sugar can read the same loudness on the plate. Second, the by-weight potencies themselves are soft: the literature spreads aspartame across 160 to 220, stevioside across 150 to 300, because sweetness potency drifts with concentration, pH, temperature, and the rest of the matrix. We treat every entry as a central value with a real interval around it, never a hard constant — a number is a form of respect for the reader only if its uncertainty travels with it.

From a molecule to a mouthful

A pure sweetener's score and a fruit's score are different animals, and the difference is the whole thesis. The molecular score is intrinsic loudness per molecule. A food's axis score is loudness *as served* — potency multiplied by the concentration actually present when you bite it. This is why a cultivar does not inherit the millionfold numbers above. It inherits the loudness its own flesh delivers.

Take the bench standard a sensory panel reaches for: a 10% (w/v) sucrose solution — 100 grams per litre, which is 100 / 342.3 = 0.292 mol/L, or 292 mM. Now ask stevioside, the plant molecule, to match that loudness. At a molar potency of 588, it needs only 292 / 588 = 0.497 mM, and 0.497 mM × 804.9 g/mol = **0.40 grams per litre**. The receptor reports the same sweetness from 100 grams of sugar or from four-tenths of a gram of a molecule a plant can be taught to make — a 250-fold collapse in mass, and with it the roughly 400 kilocalories per litre that sucrose carries at 4 kcal per gram simply do not arrive. The loudness is conserved. The load is gone.



Phoenix (SS-0096) in the controlled chamber — the first cultivar engineered to register on all eight axes at once, its sweetness reading just past the sucrose reference while it carries a fraction of sugar's mass.

That principle, realized in a living fruit, is what a sweetness score means on a real profile. Our full-spectrum cultivar Phoenix (SS-0096), our model predicts S: 1,100 — 1.10 times the sucrose reference — not because it is packed with sugar, but because its cells are designed to make potent, low-mass sweeteners, so its modest sugars read loud. It sits on the sweetness axis just past the quiet 1,000 of plain sucrose while carrying the eight-axis load that no natural plant carries at all:

IFI PROFILE

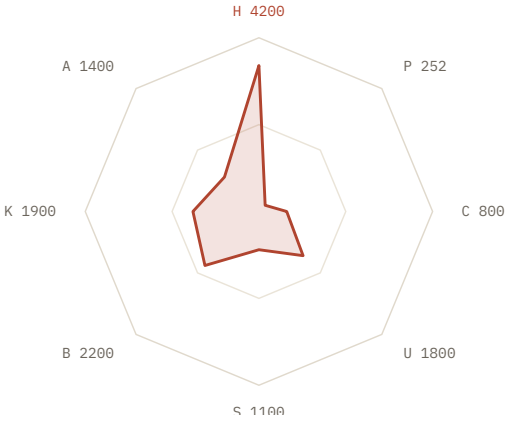

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

Read that profile with the thesis in mind. The sweetness is predicted at 1,100 — present, generous, a shade above the reference — while the heat is predicted at 4,200 and savor and kokumi hold the middle. Every one of those numbers is a computed estimate of loudness as eaten, each scored on that axis's EC_{50} ratio against its own stated reference convention, all capped at the same 5,000 ceiling, and every one was placed rather than stumbled into. The fruit is not sweet because we poured sugar into it. It is sweet because we tuned how its sweetness speaks to a protein.

The honesty of a number

A score with no error bar is a decoration, not a measurement, so it is worth being plain about where the uncertainty comes from and how large it is. The Index is a ratio of two EC_{50} values, and errors in a ratio combine in quadrature: the relative error of the score is the root of the sum of the squares of the relative errors of the two measurements,

$$CV(IFI) = \sqrt{\left(CV(EC_{50},ref)^2 + CV(EC_{50},candidate)^2 \right)}$$

If each EC_{50} is fit from a dose–response curve to a coefficient of variation of 5%, the score inherits $\sqrt{(5^2 + 5^2)} \approx 7\%$. That is the measurement floor, and it is not nothing — but it is dwarfed by the other source of variation, which is agricultural. A plant that makes its sweetener through an uncontrolled native route varies wildly from fruit to fruit; batch-to-batch coefficients of variation of around **42%** are ordinary in wild-type material. Written across a sweetness score of 1,100, that 42% is ±462 — a batch reading anywhere from 638 to 1,562, which is to say a formula you cannot build a product on. By designing the biosynthesis rather than inheriting it, we hold batch CV at **≤7%**: the same 1,100 now reads 1,023 to 1,177, ±77, a sweetness a beverage or a bake can be formulated around and trusted to repeat. Consistency is not a luxury on top of flavor. For anyone who has to put the flavor in a package, it *is* the flavor.

And there are limits to state out loud, because a receptor in motion resists being frozen into a constant. EC_{50} depends on the assay — the cell line, the expression level, the calcium buffer — so an absolute value travels less well than a ratio does, which is exactly why the Index is a ratio. The receptor adapts: hold a steady sweet stimulus and the response fades within seconds, so sweetness is forever a thing that *arrives* and *leaves*, never a plateau. It is sensitive to extracellular pH and to calcium. And the mapping from receptor current to perceived sweetness is itself nonlinear, shaped by everything downstream in that six-protein cascade. We score the receptor because it is the one honest, reproducible anchor in the chain — but we never pretend the receptor is the whole of the experience. It is the part we can measure, and measuring it is what lets us design.

Teaching a crop to speak sweet

If the receptor can be addressed in several ways, so can a plant be taught to address it — and here the enzymology matters as much as the pharmacology. The potent plant sweeteners are not sugars at all but decorated terpenoids and proteins, built by well-understood enzymes we can move and tune. Stevioside and rebaudioside A are assembled from the diterpene steviol by a short relay of **UDP-glycosyltransferases** — UGT85C2, UGT74G1, UGT76G1 — each hanging a sugar on a specific hydroxyl; shift the flux through that relay and you shift the glycoside profile, and with it the clean or lingering quality of the sweetness. Monk fruit's mogrosides are cucurbitane triterpenes glycosylated by a parallel set of enzymes.

Brazzein needs no pathway at all — it is a single small, heat-stable, acid-stable protein, which means a crop can be given the gene and simply *make* it. Each of these is a route to loudness that bypasses bulk sugar entirely, and each transfers, in principle, to any host with a sequenced genome.



Strawberry, in cross-section. Sweetness and heat on one fruit is a composition problem on the receptor — and a nourishment problem on the plant, grown to carry the most it can.

Across our live host organisms, four speak primarily on the sweetness axis, each a different strategy against the same protein:

HOST	CONSTRUCTS	AXES	DIFFICULTY	SWEETNESS STRATEGY
Tomato	2,040	S / U / F / M	Easy	restore native sugars and lift them against the shared TAS1R3 subunit, sweetness and savor tuned together
Stevia	910	S / F / M	Moderate	drive flux through the steviol-glycoside UGT relay toward the cleaner rebaudioside profile
Sugarcane	200	S	Frontier	move beyond sucrose bulk — potency, not mass, as the source of loudness
Monk Fruit	80	S	Hard	amplify the mogroside pathway; loudness from a triterpene, not a sugar

The tomato is the clearest lesson in why the receptor's architecture dictates the engineering. Because TAS1R3 is shared between the sweet and umami receptors, a tomato's sweetness and its savor cannot be designed in separate rooms — the subunit that reads its sugars is the same subunit that reads its glutamate. Lift one carelessly and you can blunt the other. So we design the pair as a pair, which is the same discipline, run forward, that lets a strawberry carry sweetness and heat at once without either drowning the other: two axes, one fruit, composed deliberately on the receptors that read them. This is the quiet through-line of all of it — a plant grown not merely to taste of something, but to carry the most flavor, and the most nourishment, its cells can be coaxed to hold.

What the receptor cannot know

There is one more property of this receptor that governs everything, and it is worth stating plainly because it reframes the sugar question entirely. The receptor does not count calories. It cannot. It has no way of knowing whether the thing holding its flytrap shut is a four-calorie gram of sucrose or a protein present at a few parts per million. It reports only one thing: *is the trap closed, and how insistently.* Every gram of sugar our species has ever eaten was, to this protein, simply a mediocre ligand present in enormous excess — a weak key used in bulk because, for ten thousand years, bulk was the only option anyone had.

We find that fact quietly liberating. The excess was never the point; the closure was. A plant can be taught to close that trap with a terpenoid it builds from sunlight, or with a small protein it folds in its own cells, or by making a modulator that lets its modest sugars read loud — each of them loudness without load, sweetness a body can receive without the weight it never asked for. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The craving at the start of this chapter is honest and ancient and ours to honor; it is only the ten-thousand-year *answer* to it that we are free to rewrite. The receptor will never know the difference. It will only ever report that the trap is closed — and we can close it a hundred gentler ways than the one we inherited.

Chapter 2 · From molecule to nerve — *how the tongue says yes.*

A sweetener has settled into the receptor, and the two lobes of the receptor's outer domain have closed over it the way a Venus flytrap folds around a fly. That closing is the whole event — the single molecular yes from which everything else follows. What follows is a cascade the taste cell runs in perhaps a few hundred milliseconds: it must take one receptor's quiet conformational change and turn it into a nerve impulse loud enough to cross a synapse and reach the brain as a word. The receptor is the lock and the key. What we walk through now is the wiring behind the lock — and you cannot design sweetness with any honesty until you know where in this chain the loudness is set, where it is amplified, and where it is spent, because that is exactly where a number can be anchored and exactly where it cannot.

We love this cascade the way we love a well-built greenhouse: every part is doing one job, nothing is ornamental, and the whole thing exists to carry a living signal from the plant to the body that will be nourished by it. Let us take it in order, and let us be exact.

The receptor is a hinge, not a pocket

The sweet receptor is **TAS1R2·TAS1R3**, and the first thing to understand is that it is not one protein but two, clasped together as a heterodimer. Both belong to class C of the G-protein-coupled receptor family — the same structural clan as the calcium-sensing receptor and the metabotropic glutamate receptors that read other axes of taste — and that shared architecture matters, because it tells you how the molecule actually works.

Each subunit carries three parts in series: a large clamshell **Venus flytrap domain (VFD)** sticking out into the mouth, a slim **cysteine-rich domain (CRD)** acting as a hinge-stalk, and a bundle of **seven transmembrane helices (7TM)** threaded through the cell membrane. The flytrap is the sensor. When a small sugar drifts into the open cleft of the TAS1R2 flytrap, the two halves of the clamshell rotate shut around it, and that rotation is transmitted down through the cysteine-rich stalk to the transmembrane bundle, where it rearranges the helices on the *inside* face of the membrane. Only then does the receptor have a new shape for the machinery below to read. The closing of the flytrap is the lock turning; the shift in the helix bundle is the bolt sliding back.

What makes sweetness unlike every other axis is that this receptor does not have one keyhole. It has several, scattered across the dimer:

- the **TAS1R2 flytrap** binds sugars — sucrose, glucose, fructose — and small synthetic sweeteners such as sucralose and aspartame;
- the **TAS1R3 flytrap** reads a second set of ligands and modulators;
- the **transmembrane bundles** hold the sites for cyclamate and for the bitter-blocker neohesperidin dihydrochalcone;

- and the **cysteine-rich domain** is where the sweet *proteins* — brazzein, thaumatin, monellin — dock, far too large to fit inside a flytrap, so they grip the stalk from outside instead.

This is why "sweet" is a single sensation produced by molecules that look nothing alike — a sugar, a terpene glycoside, a 54-residue protein — and it is the molecular permission slip for everything our family builds on this axis. Monk fruit's mogrosides, stevia's steviol glycosides, and the sweet protein brazzein each reach the same labeled line by a different door. The receptor was never a lock for one key. It is a hinge that several very different hands can close.

The first amplifier: gustducin

The instant the transmembrane bundle shifts, it grips the **G-protein** waiting against its inner face. In taste cells that G-protein is **gustducin**, and its name is a deliberate echo of transducin, the G-protein that lets the retina amplify a single photon into vision in near-darkness. The echo is not sentimental. Taste in trace amounts and sight in near-dark are the same engineering problem — make one tiny event enormous — and evolution reached for the same gear.

Gustducin sits at rest with GDP bound in its core. The activated receptor acts as a **guanine-nucleotide exchange factor**: it pries the GDP loose, lets GTP take its place, and the GTP-loaded G-protein springs apart into an active **Gα-gustducin** piece and a **$G\beta_3\gamma_{13}$** pair. For sweetness it is the βγ pair that carries the signal forward.

Here is the first and most important multiplication in the whole chain, and it is worth stating as arithmetic rather than metaphor. A single receptor, while it holds its sweetener and stays in its active shape, does not switch on one G-protein and stop. It is a **catalyst**. It keeps grabbing fresh gustducin molecules and loading them with GTP, one after another, for as long as it stays closed — tens of them in the lifetime of one binding event. Call that lifetime a few hundred milliseconds and the per-receptor output is already a crowd. The timing is self-limiting: Gα-gustducin is a slow **GTPase**, so each active piece hydrolyses its own GTP back to GDP after a short interval and switches itself off, which is how the cell keeps the amplifier from screaming forever. Gain with a built-in brake — the same discipline we want in a flavor, loud but never ragged.

The second amplifier: PLCβ2 and the calcium store

The freed βγ pair switches on an enzyme bound to the inner membrane: **phospholipase C-β2 (PLCβ2)**. This enzyme is the second amplifier, and it works by cutting. Its substrate is a membrane lipid, **phosphatidylinositol 4,5-bisphosphate (PIP_2)**, which it cleaves into two second messengers — **inositol 1,4,5-trisphosphate (IP_3)**, which dissolves into the cytoplasm, and **diacylglycerol (DAG)**, which stays in the membrane. For sweetness the live wire is IP_3. And because one activated PLCβ2 cleaves many PIP_2 molecules per second, a single enzyme throws off a burst of IP_3 — one G-protein's worth of signal becomes a cloud of messenger.

IP_3 drifts a short distance to the **endoplasmic reticulum**, the cell's internal calcium warehouse, and binds the **IP_3 receptor type 3 (IP_3R3)**, a gated channel set in the warehouse wall. The gate opens, and calcium that the cell has spent energy stockpiling floods out of the store into the cytoplasm. The resting cytoplasm holds calcium at roughly a ten-thousandth of the concentration sitting in that store, so when the gate opens the ion moves down an enormous gradient and the internal calcium level spikes hard and fast. This is the third multiplication: one PLCβ2 makes many IP_3; each IP_3 opens a channel; each open channel releases a torrent of stored calcium.

IP_3R3 adds a refinement the designer should respect — it is **co-agonist-gated**. It opens best when *both* IP_3 and a little calcium are present, so the first trickle of released calcium helps open neighboring channels, a short, self-reinforcing bloom of **calcium-induced calcium release**. It is also self-terminating: push the local calcium too high and the same channels start to shut, so the spike is a pulse, not a flood that never recedes. The cell builds its loudness out of controlled, recoverable pulses — which is precisely why sweetness has a rise, a peak, and a clean fall rather than a smear.

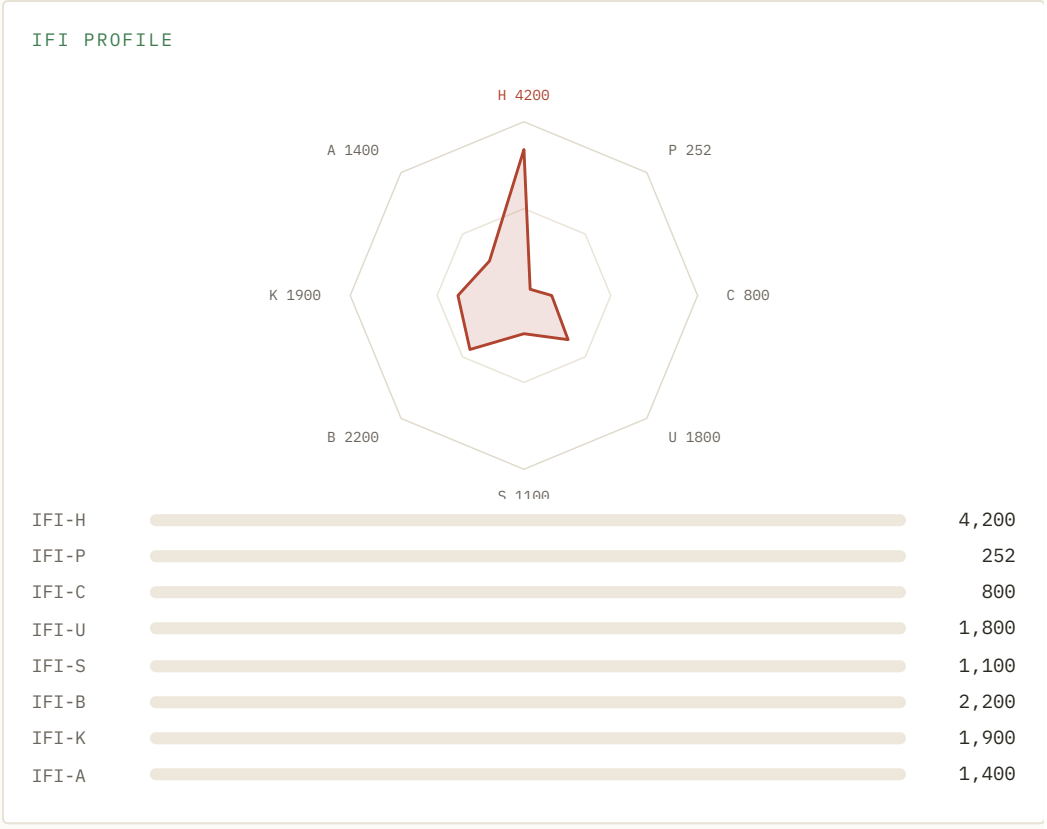

We put **Phoenix** (SS-0096) on the page here, the first cultivar our family designed to register on all eight axes at once, because its sweetness line — **S: 1,100** — is read by exactly this machinery, the same flytrap and the same calcium store, as the sugar in any fruit. The score sits a touch above the reference's 1,000, and before the chapter is done we will show precisely what that single extra hundred means in the concentration where the receptor lives.

The channel that shouts: TRPM5, and why warmth is sweet

The calcium spike is the trigger for the part of the cascade that finally raises its voice. Calcium opens **TRPM5**, a calcium-activated cation channel sitting in the cell's outer membrane. TRPM5 is a near cousin of the TRP channels that read other sensations — TRPV1 for the heat of capsaicin, TRPM8 for the cool of menthol, TRPA1 for the sharp pungency of mustard — and that family resemblance is a quiet reminder that all eight of the axes our family designs are, at the bottom, variations on one ion-channel grammar. Sweetness is written in the same alphabet as fire.

When TRPM5 opens, **sodium rushes in** down its own gradient, and the membrane voltage of the cell swings upward toward the threshold for firing. This is where the chemical story becomes an electrical one.

And TRPM5 carries a property that explains something every cook and every child already knows: **warm things taste sweeter.** The channel is steeply **temperature-sensitive** — heat makes it open more readily to the very same calcium signal. The way physiologists quantify that steepness is the **Q_{10}**, the factor by which a rate rises for a 10 °C increase. An ordinary chemical step in the body has a Q_{10} near 2 or 3; a temperature-gated TRP channel's gating can run far higher across its sensitive band. The consequence is pure arithmetic: hold the sugar constant, warm the cell, and TRPM5 passes more current per calcium ion, so the identical sweetener reads *louder*. Warm cola tastes cloying and flat soda at room temperature tastes sugary-heavy for a reason written straight into this channel — the sweetness was engineered by physics before it ever reached the tongue. It is also a lever we keep in mind: a dessert meant to be eaten warm and the same dessert meant to be eaten cold are not asking the receptor for the same amount of sweetener.

```
sweetener → receptor → gustducin → PLCβ2 → IP₃ → IP₃R3 → Ca²⁺ ↑
→ TRPM5 → depolarization → CALHM1/3 → ATP → afferent nerve
```



The reporting chain, imaged on the stage. Every axis we design ends as the same currency — a nerve firing — which is why a single index can hold all eight.

The wordless synapse: CALHM1/3, ATP, and the afferent nerve

The depolarization opens the last gate, and it is an unusual one. The sweet cell has **no classical synapse** — no cluster of transmitter-filled vesicles waiting to fuse. Instead the voltage swing opens **CALHM1 together with CALHM3**, a large-pore channel wide enough to let a whole neurotransmitter molecule pass straight through the

membrane. The transmitter it releases is **ATP** — the same molecule the cell uses for energy, repurposed here as a word — pushed directly out through the open pore, no vesicle required.

The ATP crosses the narrow gap to the gustatory nerve fiber pressed against the cell and binds **P2X2 and P2X3** receptors studding that fiber. Those are themselves ATP-gated ion channels; they open, the nerve fiber depolarizes, and *now*, at last, an **action potential** races away — up the **chorda tympani** from the front of the tongue and the **glossopharyngeal nerve** from the back — into the brainstem's **nucleus of the solitary tract**, onward through the **thalamus**, and up to the **insular cortex**, where the firing finally becomes the conscious word we started with: *sweet.*

Two facts about this final stretch shape how we think about the number, and we state them with care. The first is that the brain reads **intensity** largely as the **firing rate** of the nerve — more sweetener, faster firing, up to a saturating ceiling — while **pleasantness** is assigned later and elsewhere, in the cortex, which is why intensity and liking can be pulled apart. A sweetness can be loud and unwelcome, or gentle and perfect. The Index measures the first; the kitchen tunes the second. The second fact is anatomical, and we note it as biology and nothing more: the same TAS1R2·TAS1R3 receptor is expressed well beyond the tongue, including by the enteroendocrine cells that line the gut, where it senses sugars in the lumen. We mention it only as evidence of how deeply the body is built to read sweetness. What the receptor does there is outside this book's remit and beyond anything we would claim. Our concern ends at the tongue, where the pleasure lives and where the number is set.

One cell, one wire: the labeled line

The cell carrying all of this deserves to be named precisely, because its wiring settles an argument formulators have. It is a **Type II taste cell**, and it is single-minded: a sweet cell carries TAS1R2·TAS1R3 and *none* of the bitter or umami receptors. Sweet, bitter, and umami each live in their own separate Type II cells, each dumping ATP to its own nerve fiber. This cellular separation is the physical basis of **labeled-line coding**: sweetness runs its own dedicated wire to the brain, kept distinct from bitter and savory from the very first cell.

The proof is total and it is worth its weight. Silence TAS1R2 in an animal and it loses sweetness entirely while tasting everything else normally; restore the gene and sweetness returns, clean and specific. The axis is not a pattern to be guessed at from a noisy crowd of mixed fibers. It is separable hardware — one receptor, one cell type, one line — and that single fact is what licenses the whole enterprise of putting a number on it.

Here the biology and the mathematics fuse into one claim, so let us make it plainly. A single scalar can honestly describe sweetness **only because** sweetness is a labeled line. If taste were a distributed code — sweetness smeared as a ratio across fibers that also carried bitter and sour — then a one-number summary would be a convenient fiction, an average hiding the thing that actually matters. But

because the line is dedicated, intensity maps monotonically onto one firing rate, and one firing rate is exactly the kind of thing a single calibrated number can stand for. The index is not a simplification we impose on a tangle. It is a faithful readout of a channel the body already built as a channel. We measure a scalar because the tongue computes a scalar.

The shape of the signal: deriving the dose–response

Everything above is a **cascade of amplifiers**, and that architecture has a mathematical signature we can derive from first principles. Begin at the only step that is a simple physical equilibrium: a sweetener molecule at concentration c binding the receptor's flytrap. Let the receptor be either empty or bound, with a dissociation constant K. At equilibrium the fraction of receptors that are occupied is

$$\theta = c / (c + K)$$

This is the law of mass action, and it already has the right shape — zero when there is no sweetener, climbing steeply at first, then bending over toward a ceiling of 1 as every receptor fills. The downstream machinery, being a chain of catalysts, mostly preserves that shape and sets the ceiling, so the whole cell's response follows a **hyperbola**:

$$R = R_{max} \cdot c / (c + EC_{50})$$

where **EC_{50}** is the concentration that drives the response to **half of its maximum**. Set $c = EC_{50}$ in the equation and $R = R_{max} \cdot (EC_{50} / 2{\cdot}EC_{50}) = R_{max} / 2$ — the definition falls straight out of the algebra. EC_{50} is the one number that locates a sweetener on the concentration axis: the smaller it is, the less sweetener you need to reach half-loudness, and the *more potent*, the *louder*, the compound is.

When the receptor's binding is cooperative — and the sweet dimer, with its several sites, can show this — the hyperbola sharpens into a **sigmoid** described by a Hill coefficient n:

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

with n near 1 for many simple sugars and steeper for some ligands; a larger n makes the switch from "barely sweet" to "fully sweet" happen over a narrower band of concentration.

Now the practical payoff — the reason "just add more sugar" betrays a cook. Work the hyperbola (take $n = 1$) at multiples of EC_{50}:

CONCENTRATION C	R / R_MAX	GAINED OVER PREVIOUS ROW
0.5 × EC_{50}	0.333	–
1 × EC_{50}	0.500	+0.167
2 × EC_{50}	0.667	+0.167
4 × EC_{50}	0.800	+0.133
8 × EC_{50}	0.889	+0.089
16 × EC_{50}	0.941	+0.052

Doubling the sweetener from 4× to 8× EC_{50} buys only **0.089** more of the receptor's full response — about nine percent of the scale — for twice the sugar. Double again and you gain half as much as that. This is the physical origin of **cloying**: past the shoulder of the curve the receptor is already most of the way open, so the extra sweetener is tasted by the body as heaviness and weight long before the nerve reports anything like proportional sweetness. The curve bends over because the amplifiers downstream are already flung wide. A formulator who understands the hyperbola stops climbing at the shoulder and spends the rest of the budget making the sweetness *cleaner* instead of merely bigger — which is the entire reason our family works upstream, at the receptor's potency, rather than downstream, at brute concentration.



Sugarcane, read on the same axis as a sugar cube and a monk-fruit glycoside — three different doors to one labeled line.

From EC_{50} to a score: the Index, derived and worked

A single EC_{50} locates a sweetener, but a bare EC_{50} is awkward to read — it is small for potent compounds and large for weak ones, the wrong way round for a loudness scale, and its units change with the molecule. The Index fixes all three problems at once. Anchor every axis to a fixed **reference** stimulus, invert, and scale by a thousand:

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

Three things happen in that one line. Inverting turns "low EC_{50}" into "high score," so louder reads higher. Dividing by the reference's EC_{50} cancels the units — the score is a pure ratio, dimensionless, comparable across a sugar, a glycoside, and a sweet protein that would otherwise be measured in different units entirely. And the factor of 1,000 fixes the reference itself at a round **1,000**, the shared zero-point from which every axis is read. By construction, a candidate exactly as potent as the reference scores 1,000; one that half-activates the receptor at *half* the reference's concentration — twice as potent — scores 2,000; one that needs *twice* the concentration scores 500.

Now the worked loop, with the arithmetic shown. For the sake of a concrete conversion, take an **illustrative** reference that half-activates the sweet receptor at **$EC_{50}(\text{ref}) = 20$ mmol/L** — a round stand-in for the arithmetic, not a platform figure. Then a candidate's score and its working concentration are two faces of one fact:

- A candidate that half-activates at **2 mmol/L**: IFI = 1,000 × 20 / 2 = **10,000**. Ten times the reference's potency. Note the edge case — the per-axis scorecard is normalized to a maximum of **5,000**, so a raw 10,000 is clipped to the top of the scale for plotting. The receptor can be driven harder than the chart can draw, and the honest move is to report the raw figure beside the clipped one.
- Our canonical anchor runs the loop the other way. **Phoenix** scores **S = 1,100.** Rearranging the formula, EC_{50}(candidate) = 1,000 × EC_{50}(ref) ÷ IFI = 1,000 × 20 ÷ 1,100 = **18.2 mmol/L.** So Phoenix's sweetness reaches half-loudness at about **18.2 mmol/L against the reference's 20** — roughly **nine percent more potent**, a deliberately restrained, true-to-fruit sweetness rather than a spike. Convert that to mass with sucrose's molar mass of 342.30 g/mol and 18.2 mmol/L becomes 18.2 × 0.3423 ≈ **6.2 g/L**, about 0.62 g in a 100 mL glass — the kind of number a beverage formulator can actually hold in the hand.

Because the formula is nothing but a ratio of EC_{50} values, it runs cleanly in reverse for *any* score, and the reverse direction is where it earns its keep in the lab: the score on the page tells you the concentration at the tongue. Read Phoenix's full eight-axis profile through that lens and every axis becomes a potency ratio relative to its own reference — **ratio = EC_{50}(candidate) / EC_{50}(reference) = 1,000 / IFI** — which is the single most useful table we draw:

AXIS	PHOENIX IFI	EC_{50}(CANDIDATE) / EC_{50}(REFERENCE)	READING
Heat · H	4,200	0.238	~4.2× as potent as reference
Zero-Bitter · B	2,200	0.455	~2.2× as potent
Kokumi · K	1,900	0.526	~1.9× as potent
Umami · U	1,800	0.556	~1.8× as potent
Astringency · A	1,400	0.714	~1.4× as potent
Sweetness · S	**1,100**	**0.909**	**~1.1× as potent**
Cool-Burn · C	800	1.250	0.8× — *less* potent than its reference
Pungency · P	252	3.97	0.25× — *less* potent than its reference

Three of those rows teach the whole discipline. Heat, at 4,200, means Phoenix's capsaicin reaches half-activation at under a quarter of the reference concentration — a genuine roar, our model predicts approximately 2.95 million SHU. Cool-Burn, at 800, is the quiet one: a score *below* 1,000 means the cooling compound is **less** potent than its own reference, needing 1.25× the concentration to reach half-loudness. Pungency, at 252, is quieter still — a deliberate background note, needing roughly four times the reference concentration to reach half-loudness, so the sweetness and the heat stand fully forward. In a full-spectrum design those scores below

1,000 are not failures; they are chosen restraints, notes kept deliberately in the background so the dominant axes can speak. The index does not flatter. It reports.

The honesty of the number: CV, error, and what the index will not do

A score is only as trustworthy as the measurement under it, and here the mathematics is unforgiving in a way we find clarifying. The IFI is a **ratio** of two measured EC_{50} values, and the relative uncertainties of a ratio combine in **quadrature**:

```
CV(IFI) ≈ √( CV(reference)² + CV(candidate)² )
```

Our platform holds batch-to-batch variation to a **coefficient of variation of ≤7%** — against the **42%** typical of a wild-type plant left to its own genetics. Feed both figures through the propagation and the stakes become plain. With a well-characterized reference measured so many times that its own CV is negligible, the candidate's 7% carries straight through: the index lands in a band of about **±7%**. Even in the conservative case where reference and candidate each carry 7%, the combined uncertainty is only

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

Run the wild-type's 42% through the same two-source formula and you get $\sqrt{(0.42^2 + 0.42^2)} \approx$ **0.59 — about ±59%.** A sweetness nominally scored 1,100 could truthfully be anywhere from roughly 450 to 1,750. That is not a measurement; it is a rumor. The ≤7% consistency is not a vanity metric — it is the precondition that lets a sweetness score of 1,100 *mean* 1,100 and not "somewhere sweetish." Precision is what turns a reading into a number.

We are just as careful about what the index does **not** capture, because an honest instrument declares its blind spots:

- **It measures intensity, not pleasure.** The score is the firing rate of the labeled line; liking is assigned later in the cortex and is nobody's EC_{50}. A perfectly scored sweetness can still be wrong for a dish.
- **It is a steady-state snapshot, not a shape.** The hyperbola describes the plateau the signal climbs toward; it says nothing about *onset* and *lingering* — whether a sweetness arrives fast and clean or blooms slowly and overstays. Two compounds can share an EC_{50} and feel entirely different in time.
- **It does not see adaptation.** The cascade desensitizes as it runs, so the second bite of a dessert is measurably less sweet than the first. A single number taken at one instant cannot carry that fade.
- **Cross-modal effects live outside it.** Warmth raises TRPM5's gain, acidity and saltiness shift perceived sweetness, and aroma colors all of it through the nose. The index pins one receptor honestly and leaves the orchestra to the cook.

These are not apologies. They are the terms on which a scalar is allowed to stand for a sensation — and the labeled-line biology is exactly what keeps the scalar honest within those terms.



The sequencer bench where EC_{50} becomes a score — every sweetness our family designs is read against the same reference, on the same machine, to the same ≤7%.

Sweetness across the plants — one receptor, many hosts

Because the receptor is a hinge that many hands can close, the sweetness axis is not bound to any one species, and our family reads every sweetness-bearing host against this identical cascade and this identical reference. The constructs differ; the labeled line does not.

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

The tomato is the one that moves us most, and it is the clearest case of the receptor's reach. A tomato's sweetness and its savory depth are read by two separate Type II cells — sweet through TAS1R2·TAS1R3, umami through the glutamate machinery — and both can be grown back into a fruit that industrial handling had stripped nearly bare. The sweet protein brazzein, a 54-residue chain that grips the cysteine-rich stalk from outside, reaches the very same line as sugarcane's sucrose, which fills the TAS1R2 flytrap from within. Different doors, one wire, one number. That the same cascade can be honestly read in a tomato, a stalk of cane, a monk-fruit glycoside, and a designed full-spectrum pepper is the whole reason a single index is worth building — and the whole reason we built it.

Adaptation and the shared wiring — which is the lever

Two features of this cascade are design-critical, and we close on them because they are where the craft actually lives.

The first is that the machinery **adapts**. Hold a sweetness steady and the signal fades — receptors desensitize, the calcium store needs refilling, the channels relax — which is the physiology behind the second bite tasting less sweet than the first. A designed sweetness has to account for its own decay, front-loading just enough that the fade lands where you want it and not before.

The second is **shared wiring**, and it is the most useful fact in this book. Gustducin, PLCβ2, and TRPM5 are not private to sweetness. The bitter TAS2R receptors and the umami machinery run their

signals through the **very same enzymes and the very same channel**. Sweet and bitter and savory are therefore not processed in sealed tubes; they compete and bleed into one another at the level of the single cell, borrowing from a common pool of second messengers. This is why a trace of bitterness can flatten a beautiful sweetness — the two are drawing on the same amplifier — and it is not a complication to wish away. It is the **lever.** Quiet the bitterness feeding into that shared pool and the identical sweetener reads cleaner and louder, having stopped sharing its machinery with a rival. Designing the bitterness out of a plant, as a craft in its own right, works for exactly this reason — the room it frees up is room the sweetness was already built to fill. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The loudness, the fade, the cross-talk, and the lever that redeems it all live right here, in the wiring behind the lock, and that is where we set the number and where we spend our care.

Chapter 3 · The mathematics of loud — *why sweetness needs a number.*

For almost the whole history of cooking, sweetness was measured the only way a flavor could be: in agreement. A fruit was sweet if the people around the table said so, and the seed of that fruit was the one you kept. The first real number came from the sugar trade — degrees Brix, the percentage of dissolved solids in a liquid, read off the bent light in a refractometer — and Brix is still the workhorse of the field, honest and fast and cheap. But Brix carries a quiet, enormous flaw that this whole book exists to exploit: *it measures sugar, not sweetness.* It counts molecules of dissolved solid. It cannot see a sweetness that arrives without them. Drop a steviol glycoside or a sweet protein into water and the needle barely moves while the tongue sings. To measure what the tongue actually does — to put a number on the singing — you have to measure the receptor. That is what IFI-S does, and the rest of this chapter is the arithmetic of how.

The number the sugar trade gave us — and the one it couldn't

Brix is a mass measurement wearing an optical disguise. Dissolved sucrose bends light in proportion to how much of it is present, so a refractometer reading of ten degrees Brix means, near enough, ten grams of sugar in every hundred of solution. It is a beautiful instrument for exactly what it does, and our family grew up trusting it the way you trust a good knife. But a refractometer is counting grams of a particular solid, and sweetness is not grams of anything — it is an event at a receptor. The two track each other faithfully for table sugar and come apart completely everywhere else, and that gap is the whole doorway modern sweeteners walk through.

Fix the reference first, because every number in the chapter is read against it. We define the anchor as a ten-percent sucrose solution — one hundred grams of cane sugar made up to a liter of water. Sucrose has a molar mass of 342.3 grams per mole, so that solution works out to a tidy concentration:

```
100 g/L ÷ 342.3 g/mol = 0.292 mol/L ≈ 0.29 M
```

About a third of a molar — and, not by accident, a concentration sitting right near the half-point of sucrose's own response. That ten-percent reference is *defined* to read **IFI-S = 1,000.** Everything sweet in the world is then measured as a ratio to it.

Now watch Brix fail. A steviol glycoside potent enough to match that reference is present at roughly a thousandth of the molar concentration and a far smaller mass still; a sweet protein like brazzein or thaumatin is a few hundred amino acids folded into a shape the receptor reads as intensely sweet, dosed in milligrams. On a refractometer all of them are nearly invisible — a Brix of a fraction

of a degree against the reference's ten — while on the tongue they are every bit as loud. The sugar is gone; the sweetness is not. Brix sees the first fact and is blind to the second. The index was built precisely to see the second.



Sugarcane in the research block — the plant that gave the world its first number for sweetness, and the one that number could never fully describe.

Beidler's law: the shape every sweetener obeys

In 1954 the physiologist Lloyd Beidler, recording from taste nerves, found that taste intensity follows a simple, beautiful law — the same law that governs any molecule reversibly binding any receptor:

```
R / R_max = c / (K + c)
```

Here *c* is the concentration of the sweetener, *R* the response it produces, *R_max* the largest response the cell can give, and *K* a constant with a precise, physical meaning: the concentration at which the response reaches exactly half of maximum. Set *c = K* and the algebra is immediate — R/R_max = K/(2K) = 1/2. Half, every time. That half-maximal constant is what pharmacology calls **EC_{50}**, and it is the single most honest number you can hang on a sweetener.

Read the curve the way the tongue reads it. At low concentration — *c* much smaller than *K* — the *K* in the denominator dominates and response climbs almost linearly with how much you add; every extra molecule is heard. As *c* approaches *K* the curve bends over. Past *K* it flattens, crawling toward *R_max* no matter how much more you pile in: the saturating plateau. We can put numbers on the entire shape just by reading the equation at multiples of *K* — at *c* = 0.1*K* the receptor sits at 9% of full response, at 0.5*K* it is at 33%, at *K* exactly 50%, at 2*K* it is 67%, at 10*K* it is 91%. From barely detectable to all but saturated, concentration has to travel about a hundredfold. And the middle of that journey — right at the half-maximum — is the steepest, most informative stretch, where a small change in concentration buys the largest and most repeatable change in response.

A low EC_{50} means the receptor reaches half-response at a tiny concentration — the molecule is *loud.* A high EC_{50} means you need a great deal of it to be heard — it is *quiet.* Sucrose, our reference, is quiet; you need a lot of it, which is exactly why humanity spent centuries moving it around the planet in sacks.

Inside the receptor: one lock, many doors

The sweet receptor is a single machine assembled from two proteins, TAS1R2 and TAS1R3 — two class-C G-protein-coupled receptors that clasp together into one working heterodimer. Neither half tastes sweet alone; the pair does. Each subunit carries, on the outside of the cell, a large bilobed structure shaped and named for a Venus flytrap: a clamshell that closes around a molecule that fits it. A cysteine-rich stalk connects that flytrap to the seven-helix bundle buried in the membrane, and it is that bundle which, when the clamshell shuts, finally turns and speaks to the inside of the cell.

What makes sweetness strange and designable is that this one receptor has *many doors.* Sugars and small sweeteners settle into the flytrap of the TAS1R2 lobe. Other sweeteners — and lactisole, a sweetness *inhibitor* — work instead on the transmembrane bundle of TAS1R3. The sweet proteins — brazzein, thaumatin, monellin — are far too large to sit inside the clamshell at all, so they dock against its outer face like a wedge pressing the shutters closed. One receptor, several chemically unrelated ways in. Two facts we build on fall straight out of that architecture, and the math of this chapter leans on both.

The first is **synergy.** Because two different sweeteners can occupy two different sites on the same receptor at once, both nudging it toward the closed, active state, their combined signal can exceed the plain sum of what each would give alone. Well-chosen pairs are supra-additive. This is not a kitchen superstition; it is a direct reading of a receptor with more than one keyhole.

The second is **partial agonism.** A ligand that closes the flytrap only partway, or stabilizes a lazier conformation of the transmembrane bundle, drives the receptor to a *lower* ceiling than sucrose does — its own personal *R_max* sits below sugar's. Hold that thought; it is the reason a sweetener's "×sucrose" multiple is not a fixed number but a sliding one.

Then comes the cascade — the amplifier that turns one closed clamshell into a nerve impulse loud enough to matter, and the reason the response plateaus at all. The active receptor couples to gustducin, its dedicated G-protein; gustducin frees a subunit that switches on phospholipase C-β2; PLCβ2 cleaves a membrane lipid into IP_3; IP_3 opens intracellular stores that spill calcium into the cell; that calcium throws open TRPM5, a channel whose opening depolarizes the membrane; and the depolarization drives ATP out through CALHM1 channels to the waiting gustatory nerve. One binding event, a flood of ions, a single clean spike of *sweet.* Every stage of that chain has a finite supply — so many receptors, so many channels, so much ATP — so when the upstream signal is already saturating them, adding yet more sweetener upstream changes nothing downstream. The molecular ceiling you feel as a plateau is simply the cascade running out of room. Beidler's hyperbola is that ceiling, written as an equation.

Anchoring the scale: 1,000 × EC_{50}

IFI-S takes the half-maximal concentration, inverts it so that louder reads as larger, and ties the whole scale to the sugar everyone already knows:

```
IFI-S(candidate) = 1,000 · EC50(reference) / EC50(candidate)
```

The reference — the ten-percent, 0.29 M sucrose solution — has, by construction, $EC_{50} \approx 0.29$ M, and therefore IFI-S = 1,000 × 0.29 / 0.29 = 1,000, exactly as defined. A molecule with half the reference's EC_{50} is twice as loud and scores 2,000. A molecule a hundred times louder scores 100,000. Nothing in the formula is cleverer than a ratio; the cleverness is in having an honest number to take the ratio of.

Put three real sweeteners through it. Read each half-maximal concentration, convert into the reference's units, and divide.

SWEETENER	EC_{50}	1,000 × 0.29 M ÷ EC_{50}	MOLECULAR IFI-S
Sucrose (reference)	0.29 M	1,000 × 0.29 / 0.29	1,000
Rebaudioside M	1.16 mM	1,000 × 0.29 / 0.00116	≈ 250,000
Sucralose	0.48 mM	1,000 × 0.29 / 0.00048	≈ 600,000

Run the formula backward once to feel its reach: a sweetener whose receptor reaches half-response at 0.29 mM — a thousandth of the reference's concentration — scores 1,000 × 0.29 M / 0.29 mM = 1,000,000, a flat million, a thousand times the reference from a thousandth of the sugar. And here sweetness parts ways with every other axis. The high-potency sweeteners are so far above the reference that their raw scores run into the hundreds of thousands and the millions. A pure steviol glycoside, about three hundred times louder than sucrose, carries a molecular IFI-S near 300,000. Neotame, roughly eight thousand times louder, approaches eight million. These are true numbers, and on a scorecard they are useless — a radar that must also plot a heat of 4,200 cannot also plot a sweetness of 8,000,000. The loudest molecules in the pantry span four orders of magnitude, and no honest picture can draw them all at once. So the index does two sensible things, and keeping them straight is the key to reading every profile in this book.

Molecular loudness vs. delivered sweetness

First, for scorecards the index **clamps the per-axis display at 5,000** — the shared normalization maximum of all eight axes — so the picture stays legible and the axes stay comparable to one another. Second, and far more important: a scorecard does not report a pure compound's molecular loudness at all. **It reports the delivered sweetness of a food as eaten** — how loud the sweet axis reads in *this* tomato, *this* beverage, *this* finished plate, at the concentration actually arriving on the tongue. A cultivar's IFI-S of 2,400 does not claim its sugars are 2.4 times louder per molecule. It says the whole experience of eating it reads about two and a half times as sweet as the reference. That is the number that matters to a person holding a spoon, and it is the number we design toward.

The distinction dissolves a paradox that trips up newcomers. A high-potency sweetener has an astronomical *molecular* score and is nonetheless dosed into a drink to land a *delivered* IFI-S of, say, 1,100 — a touch sweeter than a flat sucrose cola — because you use a vanishing amount of it. The molecule is a shout; you simply whisper it. And the arithmetic of the whisper is clean. To make a food read as sweet as the reference (delivered IFI-S 1,000), you dose each sweetener to its own EC_{50}, where it too produces exactly a half-

maximal receptor response — the same half-max the reference produces. So the concentration you need is just the sweetener's EC_{50}, and the factor by which you under-dose it against sugar is exactly its molecular loudness.

SWEETENER	MOLECULAR IFI-S	DOSE FOR DELIVERED IFI-S 1,000	FEWER MOLECULES THAN SUCROSE
Sucrose	1,000	0.29 M	1×
Rebaudioside M	≈ 250,000	1.16 mM	250×
Sucralose	≈ 600,000	0.48 mM	600×

Read the last column against the second and the whole strategy is visible in a single line: *the dilution factor is the loudness.* Reb M is 250 times louder than sugar, so you spend 250 times fewer molecules to say precisely the same sweet thing; sucralose, 600 times louder, 600 times fewer. The reference tastes its 1,000 with a third of a mole of sugar in every liter — a hundred solid grams; Reb M says the identical 1,000 with a thousandth of that, a dusting you could lose on a fingertip. All of reformulation, in a sentence, is the art of composing a target *delivered* sweetness out of ingredients whose *molecular* loudness ranges across four orders of magnitude — and doing it so the body is given the sweetness without being handed the sugar to carry.

A tomato, restored: reading a delivered profile

Nowhere is the gap between sugar and sweetness plainer than in a tomato. Take a modern supermarket fruit — picked green, gassed red in a truck, its sweetness a ghost — and set beside it the same variety restored on our platform:

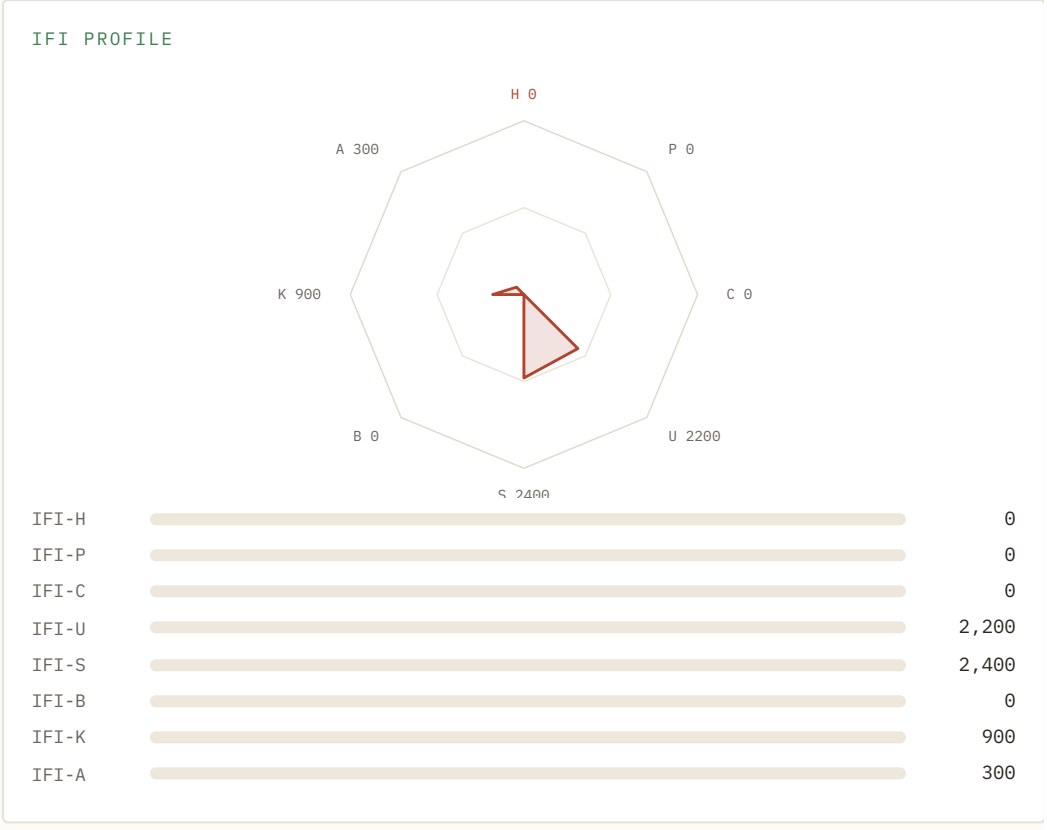

Read it the way you now can. The S of 2,400 says the restored fruit *tastes* about two and a half times as sweet as the reference — a vivid, ripe, almost startling sweetness, the kind that stops a child at the edge of the garden. Measured on a refractometer, its Brix is only modestly higher than the pale supermarket fruit's. Most of that

2,400 is not extra sugar at all. It is sweetness recovered through other routes: the volatile aroma compounds that make a tomato *read* sweet without being sugar, carried up the back of the throat to the nose, where the brain folds them back into taste as added sweetness; a floor of kokumi (K 900) that gives the sweetness body and continuity; a whisper of umami (U 2,200) that makes it land fuller and linger longer. One number, and the strategy of the whole book is legible inside it: *the sweetness went up; the sugar barely did.* We did not pour sugar into the fruit. We gave the plant back the ability to taste like itself.



A restored tomato, cut — the sweetness you read as S 2,400 is mostly aroma and body, not added sugar.

The delivered number is a real number on *every* food, not only the sweet ones. Our full-spectrum flagship — a pepper we designed gene by gene to register on all eight axes at once — carries this profile:

IFI PROFILE

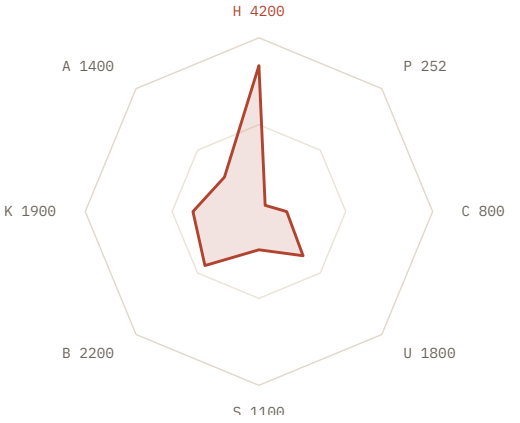

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

Set the fire aside for a moment and read only its S: 1,100, a predicted sweetness a hair above the reference, held there on purpose so the heat leads and the sweetness only rounds it off. Same index, same

anchor, same arithmetic — a tomato at 2,400 and a planet-hot pepper at 1,100, each sweetness weighed on the one identical scale. That universality is the entire point of a number.

And the sweetness work is not one plant's story. The same platform carries a roster of organisms whose primary job is the sweet axis, each with its own construct library:

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

The tomato is the easy win and the deep library; stevia and monk fruit are the natural high-potency routes, their glycosides and mogrosides the loud molecules of the arithmetic above; sugarcane, strangely, is the frontier — because improving the plant that already *defines* sweetness is a subtler problem than teaching a quiet plant to sing.

Cooperativity, thresholds, and the steepness of the knob

Two refinements turn Beidler's sketch into a working tool. The first is **cooperativity.** The plain equation assumes one molecule, one site; a real sweetener binding a receptor with several sites often climbs more steeply than a simple hyperbola, and we capture that with a **Hill coefficient** *n*:

```
R / R_max = cⁿ / (Kⁿ + cⁿ)
```

For most sweeteners *n* sits between about 1 and 1.5 — mild positive cooperativity — and that modest number has a very practical consequence, because *n* sets the *steepness* of the knob. Ask how far you must move the concentration to climb from a tenth of full response to nine-tenths of it. The Hill equation answers c_{90}/c_{10} = 81^(1/n). At *n* = 1 that span is a broad 81-fold — a gentle, forgiving knob you can turn with your eyes closed. At *n* = 1.5 it collapses to about 19-fold; at *n* = 2 it would be just 9-fold. A steeper curve means intensity changes faster with every gram you add, which makes a sweetener easier to overshoot and harder to dial to a target. Two molecules with the *same* EC_{50} — the same loudness, the same IFI-S — can be utterly different to formulate with, one lazy and biddable, the other hair-triggered, and the only thing that tells them apart is *n*. The index reports the loudness; the Hill coefficient warns you how carefully to turn it.

The second refinement is the ladder of **thresholds.** Every sweetener has a *detection* threshold — the least concentration distinguishable from plain water — then a higher *recognition* threshold, where it is identifiably sweet rather than merely present, and, well above both, the EC_{50}. We anchor the index at the half-maximum and not at either threshold for the reason the curve already told us: a threshold lives at the shallow, noisy foot of the response, where tiny changes in concentration barely move the signal and the measurement scatters badly. The half-maximum sits in the steep middle, where the signal is

most sensitive to concentration and the number comes back the same every time you measure it. Anchoring where the receptor is most informative is not an aesthetic choice; it is how you earn a number you can trust to three figures.

Potency is a moving target

Now the subtlety that bedevils reformulation and that no single EC_{50} can hold: **potency is not constant with concentration.** The familiar "×sucrose" multiple of a high-intensity sweetener *falls* as you climb toward higher sweetness. A molecule rated two hundred times sucrose at a gentle, iced-tea level of sweetness may be only a hundred times at a bold, dessert level — not because its chemistry changed between the two, but because its partial-agonist curve is already flattening toward that lower ceiling while sucrose's keeps climbing. The two curves are not parallel, so the ratio between them slides as you ascend. This is why you cannot linearly scale a recipe's sweetener to make the finished thing sweeter, and why any sweetness-equivalence value is meaningful only when it is quoted *at a stated intensity.* The map of equivalence is not flat country; it tilts, and the index measures where on the slope you are actually standing.

Perceived sweetness bends the arithmetic a second time, because it does not climb linearly with the receptor signal either. Across the usable range it follows a power law — Stevens' law — with an exponent near 1.3 for sucrose:

```
ψ = k · c^1.3
```

Double the sugar and perceived sweetness does not merely double; it rises by 2^1.3 ≈ 2.46-fold, a little faster than the sugar itself, before saturation finally sets in at the very top. Turn the statement around and it is just as useful: to make something taste *twice* as sweet with sucrose, you raise the concentration by only 2^(1/1.3) ≈ 1.7-fold, not 2. The tongue is doing its own nonlinear arithmetic on top of the receptor's, and a formulator who forgets the exponent will miss the target in both directions.

And when two sweeteners share a glass, their perceived sweetnesses frequently refuse to simply add. Well-chosen pairs show genuine **synergy** — a combined sweetness greater than the sum of the parts — traceable, as the receptor itself told us, to their binding at different sites on the same dimer and both driving it toward the active state. The sucrose-equivalent value of a blend is therefore something you *measure,* not always something you can calculate from its parts. Which is the whole reason the index has to exist: you cannot design what you cannot measure, and sweetness, until you measure the receptor, is a quantity that hides from every instrument that only counts sugar.



Sweet-heat strawberry on the platform — where a designed sweetness and a designed warmth are composed against one shared scale.

What the index measures — and what it does not

A number is a form of respect for the reader only if it is honest about its own error, so we end at the limits. The index is a ratio of two quantities, EC_{50}(reference) over EC_{50}(candidate), and error flows through a ratio by a fixed rule: the squared relative errors add. If each half-maximal concentration is pinned to a coefficient of variation of about 5% — a good, achievable dose-response fit — then the index inherits

```
CV(IFI-S) = √(5%² + 5%²) ≈ 7%
```

Two careful measurements, combined, land the index at roughly seven percent — which is exactly the batch-to-batch discipline we hold across the whole platform, a CV at or under 7% where unselected plants scatter by forty. We read EC_{50} from a nonlinear least-squares fit of response against the logarithm of concentration, six to eight doses placed deliberately to bracket the midpoint, because — one last time — the steep half-maximum is where the fit is tightest and the error smallest.

Now the honest boundaries of what the number can say. The index measures *loudness,* not liking: a sweetness of 2,400 tells you how intensely sweet, never whether you will love it. It does not capture the *shape in time* — the quick, clean arrival of sugar versus the slow bloom and long tail of some glycosides read the very same S yet eat entirely differently, and the lingering aftertaste that makes or breaks a diet soda is a temporal story the single number does not tell. It does not score the *off-notes* — the faint licorice edge of a stevia, a metallic whisper, a cooling cross-talk — that ride alongside the sweetness without being it. Delivered sweetness is *matrix-bound:* the same molecule reads differently in acid, in fat, in the cold, so a delivered profile belongs to a food as eaten and does not transfer, molecule for molecule, to the next recipe. And because perception varies from one mouth to the next, the index is calibrated on a trained reference panel — a population's honest average, not any single tongue's private truth.

None of this is apology; it is the fine print that lets the number be trusted. We built IFI-S not to replace the pleasure — the warm fruit, the hush at the table, the first sweetness arriving on the tongue — but because for ten thousand years that pleasure could be grown and tasted and never once *designed,* and you cannot design what you cannot measure. Now sweetness has a number, and a number is a thing you can aim. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.



PART II OF V

The Sugar Problem

the five jobs nobody tastes · the potency ladder · the off-notes · body without sugar

Chapter 4 · What sugar actually does — *the five jobs nobody tastes.*

The naïve idea of sugar reduction is subtraction: take out the sugar, put back the sweetness with something stronger, done. Every formulator who has tried it knows the result — a thin, hollow, oddly cold drink, a cookie that bakes pale and goes stale by Tuesday, an ice cream that freezes to a brick. The failure is instructive, and it is the reason this book is not a list of sweeteners but a design discipline. **Sweetness is only one of sugar's jobs, and it is the only one the tongue consciously tastes.** The other five are invisible, and taking the sugar out takes them with it.

We write about sucrose with real affection, because to design around something honestly you have to first understand what it was quietly doing for you. A spoonful of table sugar is a small machine running six jobs at once: it tastes sweet, it fills space, it thickens the liquid, it keeps the food from spoiling, it browns in the heat, and it sets the physical state of the finished thing — whether a candy is glassy, a dessert is scoopable, a dough can rise. Five of those six never reach the conscious tongue at all. They are physics and microbiology wearing a flavor's clothes. The whole of this chapter is spent pulling those five apart from the one, because the single most expensive mistake in reformulation is to replace the taste and assume the rest came along for free. It did not. When you lift the sucrose out of a food you open five vacancies, and each one has to be filled on purpose, by a different ingredient, obeying a different law.

The one taste the tongue admits — *and the receptor that reads it.*

Begin with the job you can feel. Sweetness is a receptor event, and the receptor is a single molecular machine: **TAS1R2 and TAS1R3**, two class-C G-protein-coupled receptors that clasp together into one working heterodimer on the sweet-sensing taste cells of the tongue. Each half carries, at its tip, a large hinged extracellular structure that structural biologists call the **Venus flytrap domain** — two lobes on a hinge, open and waiting. A sugar molecule drifts into the cleft of that flytrap, hydrogen-bonds across both lobes, and the trap closes. That closure is the entire beginning of sweetness: a mechanical shutting, a billion times a second, transmitted down through a cysteine-rich neck into the seven-helix bundle buried in the membrane.

What happens next is a cascade we can name step by step, because naming it is how we learned to build on it. The closed flytrap changes the shape of the transmembrane bundle; the bundle grips its intracellular partner, the G-protein **α-gustducin**; gustducin's release frees a βγ pair that switches on **phospholipase C-β2**; PLC-β2 cleaves a membrane lipid into the second messenger **IP_3**; IP_3 pries open a calcium store inside the cell; the flood of calcium throws open **TRPM5**, a cation channel, and the cell depolarizes; the depolarization opens a **CALHM1/3** channel that puffs ATP out into the synapse, where the gustatory nerve is waiting. Sweetness, the

sensation, is the far end of that chain. Nothing about it requires that the molecule in the flytrap be sucrose. It requires only that something close the trap. That single fact is the door the whole field of sugar-free sweetness walks through — and it is why the receptor, not the sugar, is the thing we measure.

Measuring it is the purpose of the Index. The quantity that matters is **EC_{50}**, the concentration of a stimulus at which the receptor runs at half of its full response — the half-shut flytrap, the midpoint of the dose-response curve. A loud stimulus reaches half-maximal activation at a low concentration; a quiet one needs to pile up before the receptor answers. We fix sweetness on the Index by anchoring it to a reference and inverting:

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

Sucrose is the natural reference for the sweet axis, so sucrose scores itself: substitute EC_{50}(sucrose) into the top and bottom and **IFI-S(sucrose) = 1,000** exactly. Everything louder than table sugar climbs above a thousand; everything quieter falls below it. The formula is a ratio of concentrations, so the units cancel — express both EC_{50} values in the same terms, grams per litre or moles per litre, and the number that comes out is pure loudness, a score a sugarcane, a stevia leaf, and a designed tomato can be read against on one page.

Work a real one. Our full-spectrum flagship, **Phoenix (SS-0096)**, registers **IFI-S = 1,100** on its scorecard. Put that back through the formula and read what it means at the receptor:

> $1{,}100 = 1{,}000 \times EC_{50}(\text{sucrose}) \div EC_{50}(\text{Phoenix})$, so $EC_{50}(\text{Phoenix}) \div EC_{50}(\text{sucrose}) = 1{,}000 \div 1{,}100 = 0.909$

Phoenix's sweet stimulus reaches half-activation at about **91% of the concentration sucrose needs** — roughly ten percent more potent, or, in the language of the axis, ten percent louder than table sugar. That is the whole claim of the number, stated with no poetry and no exaggeration: a little louder than sucrose, and carried by the body more lightly. Here is where that sweetness sits among the eight axes of a plant we designed to speak all of them at once:

IFI PROFILE

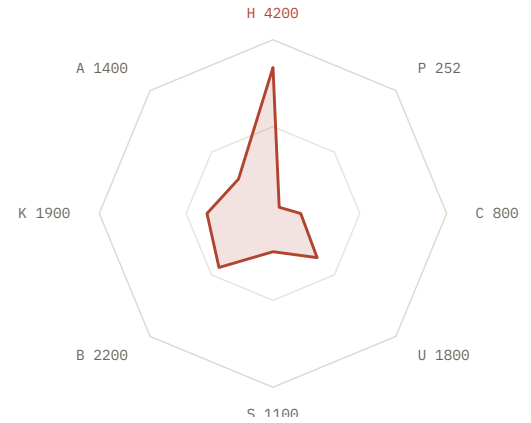

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

The axis has an enormous dynamic range, and the formula shows exactly why. A high-potency sweetener in the familiar ~600× class — sweet, by weight, six hundred times over sucrose — has an EC_{50} roughly six hundred times smaller, so its raw loudness is **1,000 × 600 ≈ 600,000**. A number that large is useless on a radar plot next to a tomato's umami, so scorecards compress every axis onto a shared ceiling: the **per-axis normalization maximum is 5,000**. The raw 600,000 clamps to 5,000 for display, while a gentle bulker and an engineered fruit sit legibly in the hundreds and low thousands.

The ladder below carries the sweet axis from a quiet polyol to the screaming high-potency ceiling, with the implied receptor concentrations shown so the arithmetic is never hidden.

SWEET STIMULUS	SWEETNESS VS SUCROSE	IMPLIED EC_{50} RATIO (CAND ÷ REF)	IFI-S
Allulose	0.70×	1.43	≈ 700
Sucrose *(reference anchor)*	1.00×	1.00	1,000
Phoenix · SS-0096, sweet axis	1.10×	0.91	1,100
High-potency class *(≈600×, illustrative)*	≈ 600×	≈ 0.0017	≈ 600,000 → 5,000 cap



Sugarcane in cross-section. Everything the plant packed into that sweetness — the bulk, the body, the browning, the keeping — comes out of the cane when you take the sugar away. Designing sweetness without sugar means designing those back in, one law at a time.

Bulk — *the volume you cannot dose in parts per million.*

Now the five silent jobs, beginning with the most brutally physical. In a cookie, sugar is often a fifth of the total mass — it is not a seasoning, it is a structural member, physical volume holding the crumb apart. The arithmetic of its replacement is where the trouble starts. Take a cola-strength sweetness of about **100 grams of sucrose per litre** and try to supply it with a 600× sweetener. You need 100 ÷ 600 ≈ **0.17 grams per litre** — about 170 parts per million. The taste is covered. But you have removed 100 g/L of dissolved solid and put back 0.17 g/L: a **99.8% hole** in the dissolved mass of the product, a vacancy the tongue and the structure both feel immediately. This is the single most common failure in reduced-sugar work, and it is not a flavor failure at all. It is a missing solid.

The hole gets filled with **bulking agents** — ingredients dosed in gram quantities, like sugar, that happen to carry little or no metabolizable energy. The workhorses are the sugar alcohols **erythritol**, the fibers **inulin** and **resistant dextrin**, and above all **allulose**, a rare sugar the body absorbs but largely does not metabolize. Allulose is worth singling out because it is quietly the backbone of serious reformulation: it bulks like sugar, it browns like sugar, it holds moisture like sugar, and it tastes — by the source's own measure and ours — about **seventy percent as sweet** as sucrose. Put that 0.70 through the sweet-axis formula and you can read allulose's loudness directly. A stimulus seventy percent as sweet needs about 1 ÷ 0.70 ≈ **1.43 times** the concentration to reach the same half-activation, so its EC_{50} is 1.43× sucrose's, and

$$IFI\text{-}S(\text{allulose}) = 1{,}000 \times EC_{50}(\text{sucrose}) \div [1.43 \times EC_{50}(\text{sucrose})] = 1{,}000 \div 1.43 \approx 700$$

An **IFI-S near 700** is exactly the profile a bulker should have: present enough to help, quiet enough that it fills the hole without hijacking the sweetness you are tuning on purpose with the high-potency component. A bulking agent is, by design, a sweetener that whispers.



Sugar beet, whole. Cane and beet are the two great reservoirs of sucrose — and the whole argument of reduction is that a reservoir is not one thing but six jobs bound into a single molecule. Pull the sugar and every one of them needs a new home.

Body — *the texture wearing a flavor costume.*

Dissolved sugar does something the tongue reads as richness but a chemist reads as rheology: it raises the liquid's viscosity and its density. Those dissolved solids give a beverage weight, a slow coat

across the palate, a *presence* the mind files under "flavor" because it arrives with the flavor. Strip them out and the drink tastes **watery** — and the complaint is almost never about sweetness. A diet soda is usually sweet enough; it is thin, because 100 g/L of dissolved solid became 0.17 g/L and the viscosity went with it. This is a texture problem in a flavor costume, and diagnosing it correctly is half the craft, because the fix is not more sweetener.

The fix is to rebuild the body separately. Bulkers restore some of it simply by being dissolved mass again. The rest comes from **hydrocolloids** — pectin, gellan, cellulose gums, soluble fibers — dosed to return the viscosity curve to something near the original. The rebuild has a ceiling and a failure mode worth naming: push a hydrocolloid past its window and "body" tips into **slime**, a gummy, mucilaginous drag that reads as cheap and wrong. The target is not maximum viscosity but the *original* viscosity, matched across the whole shear range from the slow tilt of a glass to the fast pull across the tongue. Body is a curve, not a number, and it has to be fitted, not maximized.

Keeping — *sweetness as a colligative preservative.*

Sugar's third silent job is the one that keeps food safe, and it runs on a quantity called **water activity (a_w)** — the fraction of a food's water that is chemically free, available for a microbe to drink. Pure water sits at a_w = 1.00; dissolve a solute and you lower it, because the solute ties up water molecules. To a good first approximation the law is Raoult's: a_w equals the mole fraction of water in the solution,

```
a_w ≈ n(water) ÷ [ n(water) + n(solute) ]
```

Work it for a jam. Take 100 grams of a classic preserve at about 65% sugar: roughly 65 g of sugar and 35 g of water. If the sugar has partly inverted to glucose and fructose (molar mass ≈ 180 g/mol, against sucrose's 342), the moles are 65 ÷ 180 ≈ 0.36 mol of sugar and 35 ÷ 18 ≈ 1.94 mol of water, so

```
a_w ≈ 1.94 ÷ (1.94 + 0.36) = 1.94 ÷ 2.30 ≈ 0.84
```

Real jams measure a little lower still, near **0.80**, because the solution is non-ideal and pectin and acids tie up more water than Raoult's clean line predicts — an honest gap between the model and the jar, and a reminder that the equation sets the direction, not the last decimal. Either way the number matters because microbes have thresholds: **most spoilage bacteria cannot grow below a_w ≈ 0.90, and most molds stop below ≈ 0.70.** A jam at 0.80 has pulled its water below the bacterial line, which is why an opened jar survives on a shelf — and it sits *above* the mold line, which is precisely why the failure you eventually see on old jam is a bloom of surface mold, not a bacterial rot. The sugar was never a flavor there. It was the preservative, hidden in plain sight.

Pull the sugar and a_w climbs back toward 1.00, and the shelf-stable jar becomes a petri dish. A 170-ppm high-potency sweetener contributes essentially **zero** mole fraction — it cannot lower a_w at all. So the preservation has to be rebuilt with molar quantities of

something: a **humectant** like glycerol, a bulking polyol, added acid to drop the pH, or simply the cold of a refrigerator. The molar burden is real. To hold a_w near 0.80 in that same 35 g of water you need roughly 0.49 mol of dissolved solute; supplied as glycerol (92 g/mol) that is about **45 grams** of humectant — gram-scale work, not parts per million, and a vivid measure of how much the old sugar was quietly doing to keep the food safe.

State — *the freezing point, the glass, and the yeast's supper.*

The fourth silent job is physical state — whether the finished food is scoopable, glassy, or risen — and it is pure physics and biology, done by the molecule that happened also to taste sweet.

Start with ice cream, where sugar's trick is **freezing-point depression**. A dissolved solute lowers the temperature at which water freezes by a colligative law that counts molecules, not sweetness: ΔT_f = K_f × m, where K_f for water is about **1.86 °C per mole of solute per kilogram of water** and m is the molality of the unfrozen serum. Say the serum phase of a mix carries sucrose at **3.0 mol/kg** — that is 3.0 × 342 ≈ 1,027 g of sugar per kilogram of water, a rich but ordinary dessert syrup. Its depression is

```
ΔT_f = 1.86 × 3.0 = 5.58 °C
```

so the mix begins to freeze near **−5.6 °C** and stays soft and scoopable at freezer temperatures, part liquid serum and part ice. Now swap the sugar for a 0.5 g/kg high-potency sweetener of molar mass ~300: its molality is 0.5 ÷ 300 ≈ 0.0017 mol/kg, and its depression is 1.86 × 0.0017 ≈ **0.003 °C** — nothing. The water freezes at essentially 0 °C and the dessert sets to a block of ice you cannot put a spoon through. To rebuild the 5.58 °C you need molecules back in the serum, which is why a sugar-free frozen dessert leans on a bulking polyol dosed to molality: matching 3.0 mol/kg with allulose (180 g/mol) takes about **540 g/kg** of water.

And here the molar law hides a sharp failure mode. Because depression counts *moles*, a smaller molecule depresses harder per gram. Erythritol (122 g/mol) packs 342 ÷ 122 ≈ **2.8 times** as many moles into a gram as sucrose does, so a lazy one-for-one swap by weight **over-depresses** the freezing point almost threefold — the dessert never sets, or sets slushy and weeping. Erythritol has to be dosed to a molar target, blended, and balanced, never poured in pound for pound. The arithmetic that makes it a good bulker is the same arithmetic that ruins it when ignored.

The same molecule governs two more states. In hard candy and the crisp edge of a cookie, sugar sets by its **glass transition** — amorphous sucrose turns from a flowing melt to a rigid glass near **60 °C**, and that glassiness is the snap of a lollipop and the shatter of a brittle. Nothing dosed at trace levels builds a glass; the bulker has to form one, and water plasticizes it, which is why a reduced-sugar candy picks up ambient moisture and goes sticky faster than the original. And in bread, sugar is not tasted at all — it is **eaten**, by the yeast, as the fermentable substrate that raises the loaf. This is the one job that resists replacement entirely: a high-potency sweetener gives the yeast nothing to metabolize, so a leavened dough still

needs a real fermentable sugar even when the finished bread is meant to taste barely sweet. The organism must have its supper, whatever the label says.

Color — *the two roads to brown.*

The fifth silent job is color and roasted flavor, and sugar reaches it by two different roads at two different temperatures. The first is **caramelization**: sugar alone, heated past about **160 °C**, breaks and recombines into a cascade of brown, bitter-sweet, nutty compounds — no other ingredient required. The second is the **Maillard reaction**, which begins lower, from around **140 °C**, when a reducing sugar's carbonyl meets the amino group of a protein. That meeting runs a long branching chemistry — a Schiff base, the Amadori rearrangement, fragmentation and Strecker degradation of amino acids, and finally the polymerization of brown nitrogenous **melanoidins** — and along the way it throws off hundreds of aroma molecules: the crust of bread, the top of a cookie, the toffee edge, the roast. That bouquet is not a single flavor you can bottle and add back; it is a *population* of compounds generated in place by the sugar reacting with the food.

Which is why a sugar-free bake comes out pale and tastes flat on two counts at once — short the sweetness, and short the hundreds of Maillard volatiles the sugar was quietly making. Here allulose earns its place a second time: it is a **reducing sugar**, so it enters the Maillard reaction readily and browns beautifully, often faster than sucrose, while contributing almost none of sucrose's metabolizable energy. A bulker that also browns closes two vacancies with one ingredient — the practical reason it has become a backbone of the work rather than one option among many.

The honest unit is a system — *a committee, not a molecule.*

Read the six jobs on one page and the discipline of the whole enterprise is visible at a glance.

SUGAR'S JOB	RECEPTOR OR PHYSICAL LAW	CANONICAL ANCHOR	WHAT MUST REPLACE IT
Sweetness	TAS1R2·TAS1R3 heterodimer	IFI-S = 1,000 (sucrose anchor)	high-potency sweetener (usually a blend)
Bulk / volume	physical mass	~⅙ of a cookie's mass	allulose, erythritol, inulin, resistant dextrin
Body / viscosity	dissolved-solids rheology	matched viscosity curve	bulkers + hydrocolloids
Keeping	water activity (Raoult)	bacteria < 0.90, molds < 0.70; jam ≈ 0.80	humectants, acid, refrigeration
Freezing & glass	colligative; glass transition	K_f = 1.86 °C·kg/mol; T_g ≈ 60 °C	polyols dosed to molality
Color / aroma	Maillard + caramelization	Maillard ≥ 140 °C; caramel ≥ 160 °C	allulose (a reducing sugar; it browns)
Fermentable substrate	yeast metabolism	—	a real fermentable sugar (unreplaceable)

Seven rows, and the high-potency sweetener addresses only the first. The other six are where sugar reduction is actually won or lost, which is why a serious reformulation is never a swap but a **system** — a small committee of ingredients, each covering one job the single molecule of sucrose used to do alone. A high-potency sweetener, often a blend, for the taste; a bulker such as allulose for the volume, the browning, and some of the body; a hydrocolloid to finish the viscosity; a humectant or acid to hold the water activity; a polyol dosed by molality for the freeze; a **kokumi** system — the calcium-sensing-receptor richness on its own axis — for the continuity and mouthfeel the sugar used to carry; and, if the dough must rise, a fermentable sugar the yeast can eat. Each member has a personality and a ceiling. Erythritol is tolerated in the largest amounts of the polyols because it is absorbed and excreted largely intact, but it carries a faint **cooling bite** from its negative heat of solution, a coolness you feel on the tongue as the crystals dissolve. The soluble fibers lend a gentle, sweetness-adjacent roundness and double as texture, but each has an upper use level above which it reads heavy or turns a product gassy. The art is the balance of the committee, not the strength of any one member.

Holding that balance steady, batch after batch, is its own quiet achievement, and it is where the mathematics turns back into something you can taste. We hold the engineered sweet output to a batch-to-batch coefficient of variation of **≤ 7%**, against the **42%** that conventional selection drifts across. That gap has a sensory meaning, because the tongue's own just-noticeable difference for sweetness sits near a tenth — roughly a 10% change before a person reliably says "sweeter." A 7% CV lands *under* that threshold: two jars off the same specification taste the same, every time, which is the whole promise of a designed ingredient. A 42% swing lands well past it: the product tastes different from one lot to the next, and no amount of elegant formulation survives that. Consistency is not a manufacturing nicety here; it is the difference between a flavor you designed and a flavor that merely happened.

The sweetness itself we would rather grow than synthesize, and the platform speaks the sweet axis through several living hosts. Each is scored on the same scale and written toward the fullest flavor — and the fullest nourishment — it can carry.

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



Stevia, the quiet high-potency leaf. Its glycosides close the Venus flytrap at a fraction of sucrose's concentration — the taste, sourced from a plant the body carries more lightly, while the bulk and the keeping and the browning are rebuilt around it.

Why we linger on what sugar does — *reverence, made quantitative.*

We spend a whole chapter on sugar's hidden work because it reframes the entire enterprise, and, honestly, because it is where our respect for the old ingredient lives. Sucrose was never a villain at our table. It was a brilliantly multifunctional molecule doing six jobs for the price of one, and the reason it is hard to replace is simply that it was *good* — good in a way you only measure once you try to live without it and watch the cookie go pale, the jam go moldy, the dessert go to stone. The failures are a form of praise.

What the Index finally lets a family of cooks and chemists do is see those six jobs **separately**: name each one, attach it to its law — a receptor, a mole fraction, a colligative constant, a browning temperature — keep the ones we want, and source the one that matters most, the sweetness, from a plant the body carries more lightly. The computing that found those edits was one pair of hands; the table it serves, and the care that sends us looking for the cleaner, gentler version of a thing we love, belongs to all of us. We are not fighting sugar. We never were. We are unbundling it — taking a molecule that did six beautiful things at once, understanding each one exactly enough to rebuild it, and giving the body back the sweetness it loves with a little less to carry. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 5 · The potency ladder — *every sweetener, ranked.*

At the foot of everything sweet sits one quiet molecule. Sucrose — the sugar pressed from cane and beet, the sweetness our family cooked by for three generations before we ever put it on a bench — is the reference the whole ladder leans on, and by the convention we built the Index around, its loudness is exactly **1,000**. Everything else is read against it: how many times louder, or softer, a molecule speaks to the same pair of receptor proteins, by the same rule, on the same page. The ladder is nothing more than that one comparison, performed on every sweetener we could find or grow, and then set in order.

Here is the whole cast, rung by rung, from that quiet sugar at the bottom to molecules twenty thousand times louder at the top. The ordering is by **potency** — how many times sweeter than sucrose a molecule reads — and beside it we give the half-maximal concentration (EC_{50}; lower is louder) and the raw **molecular lFI-S** that potency implies. Hold one caution the whole way up: these molecular scores are the loudness of the pure compound in a clean buffer, not the delivered sweetness of a food. The giants at the top run clean off any scorecard. They are dosed to a whisper, never to their shout.



Our sugarcane host — two hundred constructs on the sweet axis, and the botanical origin of the quiet 1.0 at the foot of the ladder. Half-maximal sweetness arrives near 0.29 moles per litre, roughly a ten-percent solution: about the sweetness of a cola, and the one reading the entire Index is anchored to.

The ladder, rung by rung

SWEETENER	CLASS	POTENCY (× SUCROSE)	APPROX. EC_{50}	MOLECULAR IFI-S	TIME-INTENSITY
Sorbitol	polyol	0.6	0.48 M	600	clean, slow, cooling
Allulose	rare sugar	0.7	0.41 M	700	sugar-like, browns
Erythritol	polyol	0.7	0.41 M	700	cooling, quick off
Glucose	sugar	0.75	0.39 M	750	clean, quick
Sucrose	**sugar (reference)**	**1.0**	**0.29 M**	**1,000**	**clean, quick on/off**
Xylitol	polyol	1.0	0.29 M	1,000	cooling, clean
Fructose	sugar	1.5	0.19 M	1,500	fast, fruity
Cyclamate	synthetic (7TM)	40	7.3 mM	40,000	slow, clean, round
Glycyrrhizin	triterpene	50	5.8 mM	50,000	slow, licorice, lingering
Aspartame	synthetic (VFT)	200	1.45 mM	200,000	slight delay, clean
Acesulfame-K	synthetic	200	1.45 mM	200,000	fast, metallic tail
Stevioside	steviol glycoside	200	1.45 mM	200,000	slow, licorice + bitter
Mogroside V	cucurbitane	200	1.45 mM	200,000	rounded, slight cooling
Rebaudioside M	steviol glycoside	250	1.16 mM	250,000	cleanest glycoside
Rebaudioside A	steviol glycoside	300	0.97 mM	300,000	clean, mild bitter at dose
Saccharin	synthetic	400	0.73 mM	400,000	fast, metallic/bitter tail
Sucralose	synthetic	600	0.48 mM	600,000	slight delay, clean, lingers
Brazzein	sweet protein	1,000	0.29 mM	1,000,000	clean-ish, lingers, heat-stable
Monellin	sweet protein	1,800	0.16 mM	1,800,000	lingering, heat-labile
Thaumatin	sweet protein	2,500	0.12 mM	2,500,000	very slow, licorice linger
Neotame	synthetic (VFT)	8,000	36 µM	8,000,000	slow onset, long linger
Advantame	synthetic	20,000	14.5 µM	20,000,000	slow, very long linger

(The EC_{50} values are illustrative, computed on a common sucrose-equivalent basis from published potency; the true figure drifts with assay, temperature, and pH. For the proteins, potency is given by weight, as the field conventionally reports it — per molecule, a protein like thaumatin at roughly twenty-two kilodaltons is louder still, and we do that arithmetic before the chapter is out.)

Two facts about this ladder are worth settling before we climb it, because both save a great deal of confusion later.

The first is bookkeeping. The eight-axis scorecard we profile a *food* on is normalized to a ceiling of **5,000** per axis — ample room for a designed tomato or a flagship pepper. The **molecular** IFI-S in the right-hand column obeys no such ceiling; advantame's twenty million is not a scorecard value and never will be. The two numbers answer two different questions. Molecular IFI-S asks *how loud is this pure compound at its receptor* — an unbounded property of the molecule. The scorecard axis asks *how sweet does this food eat* — a bounded property of a portion. A pinch of advantame is twenty-million-loud and vanishingly dilute; it lands on a food scorecard as a modest, pleasant sweetness because you dose it to a whisper. The ladder reports the shout so that, knowing it, you can choose the whisper.

The second fact is the quiet shape of a pure stimulus. Read as a clean compound in water, the reference sweetener lights exactly one of the eight axes and nothing else — a single note with no chord behind it:

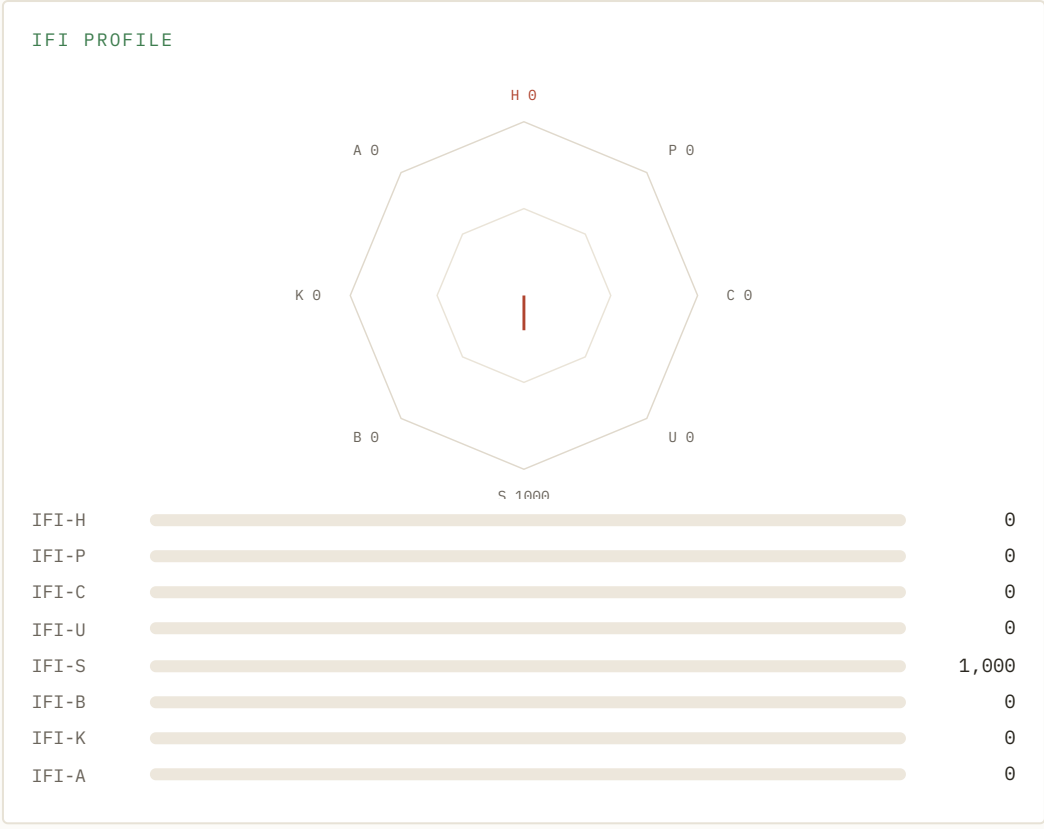

That is the fingerprint of purity: one axis bright, the other seven dark. It is also the exception, not the rule. Almost nothing you actually eat is a pure stimulus, and almost nothing on the upper rungs of this ladder stays on one axis when you taste it — saccharin drags a bitterness, the glycosides a licorice, the big proteins a slow tail. A single number tells you the loudness. It cannot, by itself, tell you the rest of the chord, and most of the craft of sweetness lives in that gap.

Reading loudness in molar terms

To put a molecule on the ladder we have to agree, exactly, on what *loud* means. We measure it the way pharmacology has measured agonists for a century: a dose–response curve. Take the sweet receptor in a cell assay, raise the concentration of a sweetener step by step, and record how hard the receptor fires. The answer is never a straight line. It is a sigmoid — silence, then a steep climb as binding sites fill, then a plateau at saturation where every receptor is already working and more sweetener changes nothing. The single most honest summary of that whole curve is one point on it: **EC_{50}**, the concentration at which the response reaches half of its own maximum. Low EC_{50}, the receptor is half-roused at a trace, and the molecule is loud. High EC_{50}, you must flood the tongue before it answers, and the molecule is quiet. Sugar, for all we love it, is quiet.

The IFI-S is a stated convention: sucrose is defined as 1,000 on the sweetness axis, by exactly the same logic as the Scoville scale defining sweetness-to-threshold by dilution. The Index does exactly one thing to EC_{50}: it inverts it against that fixed reference and scales it, so that *louder* reads as a *higher* score.

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

The reference is sucrose, anchored by definition at 1,000. Its EC_{50} sits near **0.29 moles per litre** — and because sucrose weighs 342.3 grams per mole, that is 0.29 × 342.3 ≈ **99 grams per litre**, a touch under ten percent weight-per-volume: the sweetness of a soft drink, a concentration everyone's tongue already knows. Keep that trio in a back pocket, because it is the same reading written three ways: **0.29 M = 290 mM = 290,000 µM ≈ 10% w/v.** The reason the formula is a *ratio* is that these absolute units then cancel. We never have to defend one particular molar reading of sucrose; we only have to measure every candidate in the same assay and divide. What survives the division is pure relative loudness.



EC_{50} is not a matter of opinion. We read it off the receptor assay on the bench — the loudness of a molecule in molar terms — and every rung on the ladder is anchored to a number that came from here, not from a tasting panel's adjectives.

The derivation that links the two columns of the ladder is short and worth doing once by hand. Potency, as the field reports it, is the factor by which a molecule out-sweetens sucrose: if a sweetener is *P* times as sweet, you need *1/P* of the concentration to reach the same half-maximal firing, so $EC_{50}(\mathrm{candidate}) = EC_{50}(\mathrm{sucrose}) \div P$. Substitute that into the Index formula and the EC_{50} of sucrose cancels top and bottom:

> $\text{IFI-S} = 1{,}000 \times EC_{50}(\text{sucrose}) \div (EC_{50}(\text{sucrose}) \div P) = 1{,}000 \times P$

So molecular IFI-S is simply a thousand times potency — which is why Reb A at 300× reads 300,000, and advantame at 20,000× reads 20,000,000. Run the formula forward from EC_{50} instead, on a handful of rungs, and the chain from the bench to the score is laid bare:

SWEETENER	EC_{50} (REPORTED)	EC_{50} (MOLAR)	$1{,}000 \times 0.29 \div EC_{50}$(M)	MOLECULAR IFI-S
Sucrose *(ref.)*	0.29 M	0.29	1,000 × 0.29 ÷ 0.29 = 1,000	1,000
Fructose	0.19 M	0.19	1,000 × 0.29 ÷ 0.19 = 1,526	1,500
Aspartame	1.45 mM	0.00145	1,000 × 0.29 ÷ 0.00145 = 200,000	200,000
Mogroside V	1.45 mM	0.00145	1,000 × 0.29 ÷ 0.00145 = 200,000	200,000
Rebaudioside A	0.97 mM	0.00097	1,000 × 0.29 ÷ 0.00097 = 298,969	300,000
Sucralose	0.48 mM	0.00048	1,000 × 0.29 ÷ 0.00048 = 604,167	600,000
Thaumatin	0.12 mM	0.00012	1,000 × 0.29 ÷ 0.00012 = 2,416,667	2,500,000
Advantame	14.5 µM	0.0000145	1,000 × 0.29 ÷ 0.0000145 = 20,000,000	20,000,000

Notice that Reb A computes to 298,969 and rounds to 300,000, and thaumatin to 2,416,667 and rounds to 2,500,000. That is not sloppiness; it is honesty about where the inputs came from. The EC_{50} values here are themselves rounded back out of published potencies, so the recomputed IFI lands within a percent or two of the listed figure and we report the round number. A score with more than two significant digits, on a measurement this soft, would be a lie told in the grammar of precision.

What efficacy means at the receptor

A molecule's EC_{50} tells you *where* its climb sits on the concentration axis. It says nothing about *how high* the climb goes — and that second number decides more about a sweetener's fate than its potency ever will. To see why, we have to look at the instrument the whole ladder is read against.

Sweetness is a single receptor: **TAS1R2 and TAS1R3**, two proteins that must find each other and pair before either can do a thing. Alone, each is deaf; bound together as an obligate heterodimer, they are the one lock the human body uses to hear sweetness. They belong to the class C G-protein-coupled receptors — the deliberative, nutrient-reading family — and each subunit carries an enormous **Venus flytrap domain**, a bilobed clamshell that hinges shut around a ligand, joined by a cysteine-rich collar to a seven-pass transmembrane bundle threaded through the membrane. There is not one keyhole here but several, and which one a sweetener uses is printed right in the ladder's class column: small sugars and the steviol glycosides and mogrosides settle into the cleft of the **TAS1R2 flytrap** and pull it shut; aspartame and neotame engage that same flytrap, which is why the ladder tags them *(VFT)*; cyclamate instead binds the **transmembrane bundle of TAS1R3**, tagged *(7TM)*,

nowhere near the sugar site; and the sweet proteins — brazzein, monellin, thaumatin — are far too large to fit inside any pocket, so they dock against the *outside* of the folded flytrap, a wedge pressed to a wall.

When the flytrap shuts, the motion travels down the collar, twists the transmembrane bundle, and releases the bound G-protein **gustducin** inside the cell; gustducin wakes phospholipase C-β2, which frees the messenger IP_3, which opens calcium stores, which throw open the channel TRPM5, which depolarizes the cell and vents ATP onto the waiting gustatory nerve. A few hundred milliseconds after the molecule lands, you *know* it is sweet.

Here is the part that matters for the ladder. The depth of that firing — how fully the receptor is driven toward its ceiling — is a property of the agonist, not just its dose. A molecule that locks the flytrap all the way shut and holds it there is a **full agonist**: push its concentration high enough and it drives the receptor to its true maximum. A molecule that can only coax the clamshell part-way closed, or couples to gustducin less efficiently, is a **partial agonist**: its dose–response curve flattens at a plateau *below* the full maximum, and no concentration on earth lifts it past that ceiling. Two sweeteners can share an identical EC_{50} — read equally loud at their half-max — and still be utterly different molecules, because one tops out at full sweetness and the other cannot reach it. EC_{50} is a map of *position*. Efficacy is a map of *height*. The ladder sorts by the first and stays silent on the second, which is the single most expensive thing a formulator can forget.

Potency is not quality — the shape in time

Even height is not the whole story, because sweetness is not a quantity so much as an event, and events have shape. The right-hand column of the ladder only gestures at it — *clean, quick on/off* against *slow onset, long linger* — but in the lab that shape is a measured curve. A **time–intensity** trace records perceived sweetness several times a second across an entire tasting, and from it fall the numbers that separate a sweetness you love from one you merely tolerate: the **time to peak** (sucrose peaks almost at once; thaumatin climbs for seconds), the **maximum intensity**, and the **extinction time** — how long until the sweetness is simply gone.

Sucrose draws a clean, near-symmetrical hill: a fast rise, a brief peak, a prompt fall to nothing. It is the square pulse the tongue reads as *honest* — it arrives when the food arrives and leaves when the food leaves. A high-potency sweetener's trace is often a different animal entirely: a slow rise to a long, flat, unwelcome plateau that outstays its welcome by seconds, and in that lingering tail the molecule's *second* signal — saccharin's metal, stevioside's licorice-bitter — has all the time it needs to surface and be noticed. This is the mechanical reason a single potent sweetener so rarely satisfies on its own. It is not that it is not sweet enough. It is that its sweetness is the wrong *shape*, and shape is most of what the mouth is actually judging.

Matching a reformulation's whole time–intensity curve to the original's — not merely its peak height — is what makes a sugar swap imperceptible, which is why we profile sweeteners dynamically

rather than at a single point. And there is one more moving part: sweeteners **adapt and cross-adapt** at different rates, the receptor quietly turning down its own gain across a long exposure, so a sweetness that was correct on the first sip can drift by the bottom of the glass. A molecule read once, at one instant, at one concentration, is a photograph of a thing that is actually a piece of music.

The ceiling the EC50 hides

Let us make the efficacy point quantitative, because it governs every real blend. The dose–response curve obeys the Hill relation:

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

where [L] is concentration, R_max is the receptor's own maximum firing, and *n* — the Hill coefficient — sets how steeply the curve climbs (for the sweet receptor it sits a little above 1, the mild cooperativity of the multi-site proteins). Set sucrose's R_max at 100, its EC_{50} at 0.29 M, and *n* at 1, and the whole climb falls out of arithmetic:

- at **0.145 M** (half the EC_{50}): 100 × 0.145 ÷ (0.29 + 0.145) = **33**
- at **0.29 M** (the EC_{50} itself): **50**, by definition
- at **0.58 M** (twice): 100 × 0.58 ÷ 0.87 = **67**
- at **1.16 M** (four times): 100 × 1.16 ÷ 1.45 = **80**
- at **2.9 M** (ten times): 100 × 2.9 ÷ 3.19 = **91**

Sucrose, a full agonist, keeps climbing toward 100 as long as you keep pouring. Now a partial agonist. The honest observation here is one we can anchor exactly: pure stevia and pure acesulfame-K, pushed all the way to saturation, top out near the sweetness of an **eleven-to-fourteen-percent** sucrose solution and will not climb past it. Take the midpoint, 12.5% w/v — that is 125 ÷ 342.3 ≈ 0.365 M of sucrose, which on the curve above gives 100 × 0.365 ÷ 0.655 ≈ **56**. So the partial agonist's entire ceiling is about **R_max ≈ 56** on a scale where sucrose reaches toward 100. No dose lifts it higher; past its plateau, adding more only piles on the off-notes its tail already carries.

Sort the ladder's upper rungs by efficacy rather than by potency and they split cleanly down this line. Sucrose, sucralose, and aspartame are **full agonists** — reach for the dose and they reach full sweetness. The steviol glycosides and acesulfame-K are **partial agonists**, each flattening near that eleven-to-fourteen-percent ceiling and climbing no further however hard you push. Two of those molecules — a full one and a partial one — can read the same EC_{50} on the ladder and share a rung, yet one can carry a dessert and the other cannot. This is the mathematical reason a pure-stevia cola tastes *both* not-quite-sweet-enough *and* faintly bitter at the same time: you are pressing a partial agonist up against its ceiling, where it can neither reach full sweetness nor stop broadcasting its second signal. For a tabletop sweetener the ceiling is irrelevant — nobody needs cola-strength sweetness in their morning tea. For a dessert or a concentrated

syrup, where the target sits *above* 14% sucrose-equivalent, the ceiling is decisive, and it is the reason the answer is never "more stevia." The answer is a blend.

Why the answer is a blend

Blends are not a compromise; they are the geometry of the receptor made practical. Because the lock has several keyholes — the TAS1R2 flytrap, the TAS1R3 transmembrane bundle, the protein wedge on the outer wall — two sweeteners can occupy *different* sites on the *same* receptor at the same time. When they do, the response is frequently **supra-additive**: the pair stabilizes the active conformation more completely than either could alone, and the mouth reads more sweetness than the doses, summed, should buy.

Put numbers on it. Suppose sweetener A, at its chosen dose, delivers a perceived sweetness of 6 in sucrose-equivalent units, and B delivers 6 at its dose. A purely additive model predicts 12. If a careful panel reads the blend at 15, the synergy factor is

> $\gamma = \text{observed} \div \text{additive} = 15 \div 12 = \mathbf{1.25}$

— a quarter more sweetness for no additional molecule, because A sits in the flytrap while B works the transmembrane site and each deepens the other's hold. (That arithmetic is illustrative of the model, not a measured datum; the point is the *form* of the gain, which is real and which we exploit deliberately.) The same geometry runs the other, subtler way too: two sweeteners each dosed *below* their own detection threshold can, together, cross it — summation out of two silences. Formulators draw this as an **isobologram**, a plot of "how much A alone" against "how much B alone," joined by the straight line of simple additivity; a synergistic pair lands *below* that line, meaning less total sweetener reaches the target, and with less sweetener comes less of every tail and off-note the sweeteners carry.

Synergy is why the working answer to the partial-agonist ceiling is to pair a ceilinged partial agonist with a full-agonist partner that carries the top, and to spread the load across several binding sites rather than flogging one. It is also a hedge against adaptation: a receptor turning down its gain on one site keeps answering the others, so a blend's sweetness drifts less across a long glass than a single molecule's does. The ladder tells you how loud each voice is. It takes a blend to make several voices into a chord that holds.

The price of a number — error, and honest limits

Every number in this chapter has a cost, and we would rather quote it than hide it. Begin with measurement error. An IFI-S is a ratio of two EC_{50} values, each with its own spread, and for a ratio the relative variances add in quadrature:

> $(\text{CV of IFI})^2 = (\text{CV of reference } EC_{50})^2 + (\text{CV of candidate } EC_{50})^2$

If each EC_{50} is measured to a coefficient of variation of about 5%, the IFI inherits CV = $\sqrt{(0.05^2 + 0.05^2)} = \sqrt{0.005} \approx$ **7.1%**. That is not a coincidence we are fond of: roughly seven percent is the same

batch-to-batch reproducibility our platform holds its cultivars to — **≤7% CV**, against the 42% a wild crop swings across sun and soil and season. A sucralose reading of 600,000 at 7% is really (6.0 ± 0.4) × 10^5; we write the round number and carry the band. A score without its band is half a measurement.

Then the limits of the number itself, which are more interesting than its error bars:

- **The vacuum problem.** Every EC_{50} here is a clean-buffer reading, and the footnote says so — the true figure drifts with assay, with temperature, and with pH. A food is not a buffer. Acids, viscosity, competing tastants, even the fat that coats the receptor, all shift the perceived midpoint. Molecular IFI-S is loudness *in vacuo*; delivered sweetness is what survives the matrix.
- **Species specificity.** The sweet proteins are sweet to us and to our primate kin, and essentially silent to rodents, because the wedge they press against lives on a stretch of the human flytrap that the mouse receptor does not share. An assay run on the wrong receptor would score brazzein near zero. The ladder is written for a human tongue, and only a human — or humanized — receptor earns a place reading it.
- **Thermal fate.** Potency measured cold can mislead a bake. Thaumatin and brazzein are disulfide-stapled knots that hold their sweetness through heat; monellin is heat-labile and falls silent when it denatures. The ladder's time–intensity column whispers this; a formulator baking at 180 °C had better read it aloud.
- **Weight versus molecule.** Potency by weight and potency by molecule diverge exactly in proportion to molar mass, and for the proteins the gap is the whole story. Molar potency equals weight potency times the ratio of molar masses, so against sucrose's 342.3 g/mol:

PROTEIN	MASS	POTENCY BY WEIGHT	× MOLAR-MASS RATIO	PER-MOLECULE POTENCY
Brazzein	~6.5 kDa	1,000×	×19	~19,000×
Monellin	~11 kDa	1,800×	×32	~58,000×
Thaumatin	~22 kDa	2,500×	×64	~160,000×

A single molecule of brazzein, gripping the outside of the flytrap, does the work of nineteen thousand molecules of sugar; thaumatin, heavier, more than a hundred thousand. That is what "the loudest sweetness in nature" means, made exact — and why the Index keeps its books in grams for the formulator reading an invoice while we keep the molar figure for the biology, because the receptor counts molecules and the kitchen counts grams, and an honest standard serves both readers on one page.

Where the ladder points

A ladder of pure compounds is, in the end, a list of targets — and the reason we built it is that nature already makes every rung with enzymes, and enzymes are what our crops can be asked to write. Four of our forty-two established host organisms speak the sweetness axis as a primary voice, and their construct counts measure how deeply each has been worked:

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

Tomato sits at the top with 2,040 constructs not because a tomato is the sweetest thing we grow but because it is the most instructive — a fruit people believe they already know, where a restored sweetness registers as the fruit simply tasting the way memory promised. Sugarcane sits near the bottom, a frontier host with 200 constructs, for the opposite reason: it already makes the reference molecule to perfection, and there is little to teach a plant that invented the benchmark. Stevia, between them, is where the glycosyltransferases earn their keep, dressing steviol in sugars at the positions that keep it loud on the sweet receptor and quiet on the bitter one.

And the axis does not only live in the obvious sweet crops. Read the full profile of Phoenix — SS-0096, a *Capsicum chinense* our model predicts at approximately 2.9 million SHU, a computed IFI-H of 4,200, the loudest plant our family has ever designed, lit on all eight axes at once — and the sweetness line sits, calm, at **1,100**:

IFI PROFILE

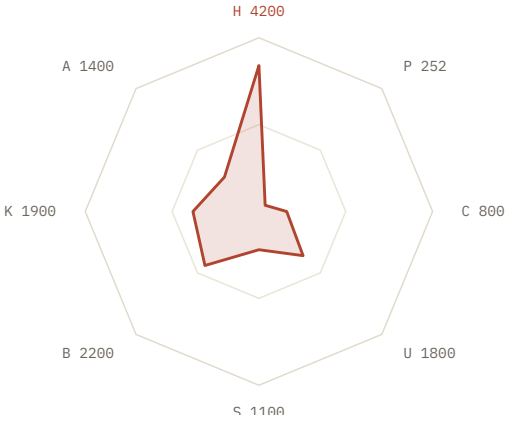

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400



Tomato — 2,040 constructs, and the clearest lesson on this axis: sweetness lifted not by adding a grain of sugar but by tuning what the fruit already builds. The receptor reads the fruit as sweeter; the molecular loudness on the ladder becomes a delivered taste on the tongue.

That one figure — `S: 1100`, sitting unbothered inside a wall of nearly three million Scoville units in prediction — is the whole argument of the ladder in a single number. A sweetness read from a receptor can be addressed independently of every other signal a plant sends; it is not borrowed from sugar, or it could not survive in a fruit this fierce. Decouple the taste from the carrier and the combinations open without limit, and every one of them is the same quiet act performed on the one pair of cupped hands the body uses to hear sweet. We catalogued the whole cast of sweetness and set it on a ladder for the people who have to solve sweetness for a living — the formulator facing a sugar-reduction target, the clean-label founder, the team that has tasted licorice in eleven stevia blends and wants to know, exactly, why. The math is shown rather than waved, because to us a number is a form of respect for the reader and for the plant. And the argument underneath it is the one our family makes about everything we grow: that food should give the body pleasure and nourishment in the same bite, with nothing to apologize for. Sweetness was never the thing to fear. It was only waiting to be understood.

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

Chapter 6 · The off-notes — *bitterness rides with sweetness.*

Nature almost never offers a clean, loud sweetness for free. The molecules potent enough to matter — the ones that can carry a dessert or a morning cup on a few parts per million rather than a heaping spoonful — tend to arrive with a passenger: a bitterness, a licorice note, a metallic edge, a slow tail that will not leave the tongue. We have tasted every one of them at our own table, steeping the leaf, stirring the glass, waiting for the second where the sweetness turns on itself. That passenger is not a flaw in the manufacturing and it is not an impurity to be filtered away. It is written into the same molecule, and frequently into the very same receptors, as the sweetness. Understanding where it comes from is the whole difference between masking it clumsily and designing it out — and designing it out is the move the platform exists to make. So we begin, as always, by saying exactly where the off-note lives, because you cannot turn down a thing you have not first learned to measure.

Two cells, one wire — why a bitter note subtracts from sweet

Sweetness begins at a single protein assembled from two: **TAS1R2 and TAS1R3**, a heterodimer of the class C G-protein-coupled receptors, standing side by side in the membrane of a taste cell. Each half carries a large clamshell-shaped **Venus flytrap domain** at its tip, a cysteine-rich hinge below it, and a seven-pass transmembrane bundle anchoring it to the cell. Sugars and most small sweeteners settle into the flytrap of the TAS1R2 lobe and pull it closed; the big sweet proteins wedge against the outer face of the same domain; a few compounds instead slip into the transmembrane pocket of the TAS1R3 half. However they bind, the geometry does one thing — it snaps the clamshell shut, and that mechanical closing is the signal.

What happens next is a cascade the cell runs the same way every time. The closed receptor activates **α-gustducin**, its G protein; the freed βγ subunits switch on **phospholipase C-β2**, which splits a membrane lipid into the second messenger **IP_3**; IP_3 opens a calcium store inside the cell; the calcium gates **TRPM5**, a cation channel that lets sodium flood in and depolarizes the membrane; and that voltage swing opens **CALHM1**, the channel that releases ATP onto the nerve fiber waiting beneath. Lock, clamshell, gustducin, PLC-β2, calcium, TRPM5, ATP — a six-step relay repeated a billion times a second, and the loudness you perceive is set by how far along that relay the stimulus can drive the cell.

Here is the fact that makes an off-note so costly. Bitterness runs on a *different* set of receptors — the **TAS2R** family, two dozen of them — carried on *different* taste cells. But those cells run the **identical downstream wire**: the same gustducin, the same PLC-β2, the same TRPM5 and CALHM1. Sweet and bitter are two locks on two cells reporting up one shared kind of wire, and the brain reads their

signals in the same region, at the same instant, where they interfere. This is why a bitterness that rides alongside a sweetness does not merely *add* a bad taste on top of a good one. Through the central phenomenon of **mixture suppression**, the bitter signal actively *subtracts* from the perceived sweetness, flattening it, so you get less sweet *and* a bad note for the same molecule. If we write the perceived sweetness as the intrinsic loudness S_0 discounted by the bitter drive B,

$$S_{\text{perceived}} = S_0\,(1 - \beta B),$$

then even a modest passenger is expensive. Take an illustrative suppression weight $\beta = 0.4$ and a bitter drive sitting at half of full, $B = 0.5$: the sweetness is cut by $0.4 \times 0.5 = 0.20$, a fifth of the loudness gone, not to any lack of sweetener but to the company it keeps. This is the licorice-and-lingering-bitter people mean when they say they can *taste the stevia*. It is a receptor fact, and it is a subtraction — which is also the good news, because a subtraction can be removed.

The passenger is written into the molecule

Start with the case that has defeated a thousand clean-label projects, because it is the one we have spent the most of our own evenings on: **steviol glycosides and bitterness.** Stevioside, the dominant sweet compound in the raw *Stevia rebaudiana* leaf, and to a lesser degree rebaudioside A, do not only close the sweet receptor. They also open two bitter ones — **TAS2R4 and TAS2R14** — and that bitter signal arrives on the tongue at the same moment as the sweet, through the shared wire, where it subtracts. The licorice tail of a stevia-sweetened drink is not a purity problem to be solved with a finer filter. It is the molecule doing two jobs at once.



The sweetness and the licorice live in the same leaf — and, at the receptor, nearly in the same molecule. Shifting a plant's output from stevioside to rebaudioside M moves the sweetness off the bitter receptors it was riding.

Now the structural detail that is the whole opening, the hinge on which this chapter turns: **the different steviol glycosides open those bitter receptors to very different degrees.** They share one rigid core — the diterpene **steviol** — and differ only in how many glucose units are hung on it, and where. Stevioside carries three. Rebaudioside A carries four. Rebaudioside D carries five. **Rebaudioside M carries six.** And as the sugars accumulate, the sweetness is largely preserved while the fit at TAS2R4 and TAS2R14 is progressively spoiled. Reb M is sweet and nearly clean; stevioside is sweet and

distinctly bitter; the two sit on the same scaffold, a glucose or three apart. The bitterness and the sweetness, in other words, can be *pulled apart by molecular structure* — the extra glucose units that distinguish Reb M happen to clash with the bitter pocket while still closing the sweet flytrap. This single fact is why the stevia story is really a biosynthesis story. The goal is never to grow *more* steviol glycoside. It is to grow the *right* one.

A glucose at a time — the fix is an enzyme, not a filter

The leaf builds its glycosides on an assembly line of **UDP-glucosyltransferases**, each enzyme hanging one glucose at one position. Steviol is glucosylated at its C-13 hydroxyl to make steviolmonoside; a second glucose is added to make steviolbioside; the carboxyl at C-19 is capped to make stevioside; and then **UGT76G1** reaches in and adds the branch glucose that converts stevioside into rebaudioside A — and, with its partners, carries A onward to D and to M. The bitter-tinged early product and the clean late product are separated by a handful of enzymatic steps. Shift the balance of those enzymes — give the plant more of the activity that finishes the molecule — and the leaf's output slides down the ladder, away from bitter stevioside and toward clean Reb M, with no change to the sweetness the plant was already making.

GLYCOSIDE	GLUCOSE UNITS	MW (G/MOL)	TAS2R4 / TAS2R14 ENGAGEMENT	WHAT THE TONGUE READS
Steviolbioside	2	642.7	High	Sharp, herbal-bitter
Stevioside	3	804.9	High	Sweet with a bitter-licorice tail
Rebaudioside A	4	967.0	Moderate	Cleaner, a soft lingering edge
Rebaudioside D	5	1129.2	Low	Clean, slow, rounded
Rebaudioside M	**6**	**1291.3**	**Minimal**	**Sugar-like, near-silent passenger**

(Molecular weights are the compounds' own chemistry; the receptor-engagement column is the qualitative ordering our panels and the published pharmacology agree on — the heavier the glycoside, the quieter the bitter.)

The engineering target follows directly, and it is a number we can state: **910 constructs** stand behind our stevia work, and not one of them is aimed at making the leaf sweeter in the crude sense. They are aimed at the ratio — at moving output from the three-glucose molecule to the six-glucose one. This is also where the mass-versus-molecule distinction earns its keep for a formulator. Dose a beverage with Reb M at 300 parts per million — that is 300 mg/L — and in molar terms you are delivering

$$\frac{300\ \text{mg/L}}{1291.3\ \text{g/mol}} = 0.232\ \text{mmol/L}.$$

Dose the same 300 mg/L of stevioside and, because its molecule is lighter, you deliver

$$\frac{300\ \text{mg/L}}{804.9\ \text{g/mol}} = 0.373\ \text{mmol/L},$$

which is $0.373/0.232 = 1.6\times$ as many molecules — *more* of them, and every one the kind that also finds the bitter receptor. At equal mass the cheaper leaf hands you a crowd of the wrong molecule. The clean line hands you fewer molecules of the right one. The sweetness can be made to match; what you have quietly removed is the passenger.

Reading loudness as a number — the index on the sweet axis

None of this is governed by poetry. It is governed by **EC₅₀** — the concentration at which a stimulus drives its receptor to half of its full response — and the index is built on it:

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

A lower EC_{50} means a louder stimulus: it reaches half-activation at a lower concentration, so it speaks more loudly for the same presence. The formula inverts that half-activation against a fixed native reference and scales by a thousand, so the **reference scores exactly 1,000** — put the reference in for the candidate and the ratio is one. The relationship is a simple hyperbola. Halve a stimulus's EC_{50} and you double its score; quarter it and you quadruple it; the axis is normalized so scores saturate toward a per-axis ceiling of **5,000**.

Behind that single number lies the shape of a whole curve. The receptor's response climbs with concentration along a sigmoid,

$$r = R_{\max}\,\frac{[S]^{\,n}}{\mathrm{EC}_{50}^{\,n} + [S]^{\,n}},$$

and EC_{50} is just the concentration where that climb reaches half of $R_{\max}$ — set $[S] = \mathrm{EC}_{50}$ and the fraction is exactly one-half. The exponent n, the **Hill coefficient**, sets how *steep* the climb is. Because the sweet receptor's flytrap can close cooperatively, many sweeteners run steep, with $n > 1$: a narrow band of concentration separates *not sweet enough* from *cloying*, which is the knife-edge every beverage reformulator feels under the thumb. That same steepness governs the onset and offset slopes — how fast the sweetness rises as a sip floods the mouth, how fast it falls as the mouth clears — so two stimuli can share an EC_{50}, and a score, and still feel unlike each other because their curves have different shapes. The index pins the half-activation; the curve around it is where the drinking happens.

EC₅₀(C) ÷ EC₅₀(REF)	IFI-S	READING
1.000	1,000	the reference, by definition
0.909	1,100	1.10× the reference's loudness
0.500	2,000	twice as loud
0.250	4,000	four times as loud
0.200	5,000	the normalized ceiling



Cane sugar — the clean, passenger-free sweetness every engineered sweetener is measured against: loud at the peak, and almost silent in the tail.

One guard against a common confusion: IFI-S is *not* "times sweeter than sugar." A high-potency sweetener can be hundreds of times as potent as sucrose by weight, and if the index reported raw molar potency the scale would run to the hundreds of thousands and tell you nothing a tongue can feel. The axis instead anchors to a native reference and reports *loudness at the concentration actually presented*, compressed onto a human 0–5,000 range. It answers the only question that matters at the table — how loudly does this speak, right now, in this glass — and it lets a leaf, a fruit, and a cane sugar be read against one another on a single line.

Watch the arithmetic run the other way, from a known score back to the chemistry. Our full-spectrum flagship, **Phoenix**, carries a predicted sweetness reading of **1,100**. Invert the formula:

$$\frac{\mathrm{EC}_{50}(\mathrm{ref})}{\mathrm{EC}_{50}(\mathrm{Phoenix})} = \frac{1100}{1000} = 1.10,$$

so its sweet stimulus reaches half-activation at a concentration about **9% lower** than the reference — it is 1.10 times as loud. That is a small, honest number, and it sits on the same plant as a predicted heat reading of 4,200. Phoenix is the clearest illustration for this chapter, because its design profile predicts a loud sweetness and a strongly attenuated bitterness held **on one organism, at once**:

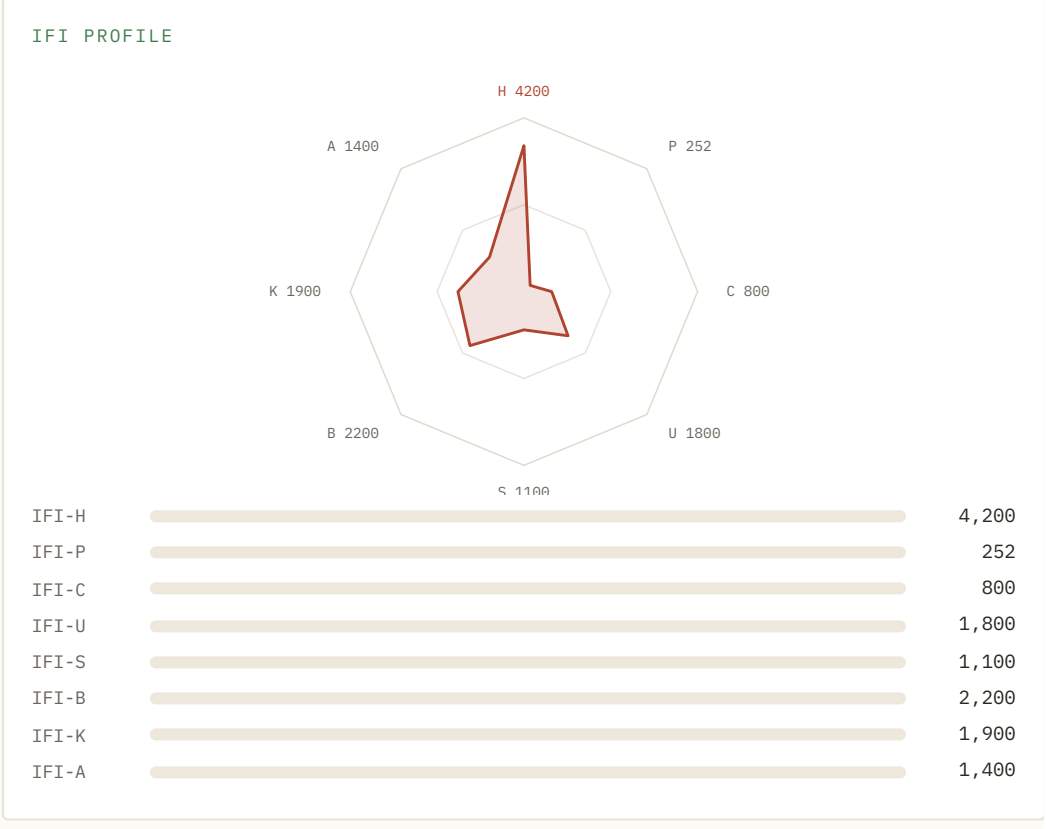

Read the two axes this chapter cares about side by side: **S: 1100** and **B: 2200** — a real sweetness, and a Zero-Bitter reading twice as high, meaning the bitter channel is designed to be significantly attenuated. Loud sweet, quiet bitter, printed as two numbers on one card. That pairing — not either number alone — is the honest fingerprint of a clean sweetener, and it is the thing the whole platform is built to deliver.

The other addresses — saccharin, the polyols, the sweet proteins, the tail

The passenger has more than one address, and each one is a specific receptor with a specific remedy.

Saccharin and, for some people, acesulfame-K read metallic and bitter because they open a *different* pair of bitter receptors — **TAS2R43 and TAS2R44** — rather than the TAS2R4/TAS2R14 pair the steviol glycosides use. Whether you taste that bitterness is partly written in your genes, which is exactly why saccharin divides a room: half the table tastes a clean sweetness and half tastes a tin spoon.

The polyols — erythritol, xylitol — carry something subtler and genuinely physical: a **cooling** sensation. It is not a flavor and not a flaw; it is a *thermal event*. These crystals dissolve endothermically, pulling heat out of the saliva as they go into solution, and that small real drop in temperature brushes **TRPM8**, the cold-and-menthol channel, so a mouthful reads faintly of mint. Lovely in a lozenge; a problem in a warm, round flavor it was never meant to chill. You cannot taste your way past this one — you have to know it is a heat of solution and formulate around the thermodynamics.

Glycyrrhizin and thaumatin bring a licorice quality and a slow, creeping onset. Thaumatin is one of the sweet *proteins*, a molecule too large to tuck into the flytrap the way a sugar does; it wedges against the outer face of the receptor's clamshell, which is part of why its sweetness arrives late and leaves slowly. That profile makes both of them more at home as background enhancers — a quiet floor under another sweetness — than as the whole sweetness of a drink.

And nearly every high-potency sweetener, pushed to a high dose, grows a **lingering tail** — a sweetness that overstays. The cause is kinetic. Once the flytrap closes, how long it stays closed is set by the stimulus's **off-rate**, k_{off}; the mean time the receptor is held shut is its reciprocal, the residence time

$$\tau = \frac{1}{k_{\text{off}}}.$$

A clean-cutting sweetener with $k_{\text{off}} \approx 1\ \text{s}^{-1}$ releases in about a second and the sweetness ends with the sip. A slow one at $k_{\text{off}} \approx 0.1\ \text{s}^{-1}$ holds for roughly ten seconds, re-closing the flytrap again and again after the glass is down — the tail you remember when the drink is gone.

OFF-NOTE	WHERE IT RIDES	RECEPTOR(S)	HOW WE ADDRESS IT
Licorice / bitter tail	Stevioside, Reb A	TAS2R4, TAS2R14	Shift biosynthesis toward Reb M
Metallic / bitter	Saccharin, ace-K	TAS2R43, TAS2R44	Choose a sweetener off these receptors
Cooling	Erythritol, xylitol	TRPM8 (thermal)	Formulate around the heat of solution
Slow licorice onset	Glycyrrhizin, thaumatin	TAS1R2·R3 (protein wedge)	Use as a background floor, not the lead
Lingering tail	Most high-potency sweeteners at dose	TAS1R2·R3 (slow k_{off})	Dose down; blend faster-clearing partners

The passenger on two clocks — time-intensity and the aftertaste

An aftertaste is not a single note fading. It is **two notes falling out of step.** A lingering bitterness and a lingering sweetness decay on different schedules — different off-rates, different clearance from the taste pore — so a sweetener can read clean at its peak and then sour the memory of itself a few seconds later, as the sweetness drops away faster than the bitter tail it was masking. We record this directly: a **time-intensity** trace, where a trained panelist rides a slider through the whole life of a sip, turns the aftertaste from an argument into a pair of curves. And matching the *bitter* extinction to the *sweet* extinction — making the two curves end together — is as much a part of a clean reformulation as matching their peaks. A sweetener that is right for one second and wrong for the next has not been finished; it has only been started.

There is a mirror-image curiosity on the same clock, and it is worth knowing because it is the receptor's dynamics made audible. After the sweet receptor has been strongly adapted, or just after a sweetness *inhibitor* has washed away, plain water can briefly taste **sweet** — the receptor rebounding past its own baseline before it settles. These after-effects are not noise in the measurement. They are the system's kinetics, and a formulator who ignores them ships a sweetness that is honest at the peak and dishonest in the tail.

The passenger is not tasted equally — genetics and the panel

Here is the fact that reshapes every sensory panel we run: **the passenger is not tasted equally by everyone.** Bitter perception is strongly genetic, and the cleanest illustration is **TAS2R38**, the receptor for the probe compounds PROP and PTC. It comes in two common forms that differ at three amino-acid positions — 49, 262, and 296 — giving a "taster" haplotype and a "non-taster" haplotype. Whether a person tastes the sharp bitterness of PROP loosely predicts how keenly they register the bitter edge of saccharin, of a stevia drink, of a raw cruciferous leaf from the garden. The consequence for our work is direct and humbling: a reformulation that reads perfectly clean to a *non-taster* on the bench can read distinctly bitter to a *taster* across the counter. "Clean" asserted from one tongue is not clean. It has to be measured across the distribution.



We genotype the panel, not just the sweetener — because whether the passenger is heard at all is written in the taster.

So a serious sweetness panel is **screened for taster status**, and the IFI-B axis exists precisely so that bitterness becomes a number agreed across panelists rather than a dispute between them. But agreement across people is only half of consistency; the other half is agreement across *batches*, and this is where the platform's discipline shows as arithmetic. The coefficient of variation — the standard deviation as a fraction of the mean — is **≤ 7%** on our lines, against **42%** in the wild-type leaf. Write it out. For a target loudness μ at 7% CV, the batch-to-batch standard deviation is $0.07\,\mu$, and 95% of batches land inside

$$\mu \pm 1.96 \times 0.07\,\mu = \mu\,(1 \pm 0.137),$$

a window of about **±14%**. At the wild-type's 42% the same window is $\mu\,(1 \pm 1.96 \times 0.42) = \mu\,(1 \pm 0.82)$ — **±82%**, a sweetness that can nearly double or very nearly vanish from one harvest to the next. A sweetener you cannot repeat is not a sweetener a formulator can build on, and a bitterness you cannot repeat is a passenger that boards at random. Holding the loudness inside a seventh is not a vanity metric. It is what lets a clean note *stay* clean.

Mask it, or design it out

When masking is the chosen route, its tools are specific and worth knowing precisely. **Sodium** salts suppress bitterness broadly — a few hundredths of a percent will quiet a bitter edge with no meaningful saltiness, which is one honest reason a pinch of salt makes a sweet thing taste cleaner as well as sweeter. Certain **nucleotides** and flavor compounds damp particular bitter receptors. **Neohesperidin dihydrochalcone** is the unusual case of a molecule that is *both* a sweetener and a bitter-masker at once. And at the far extreme sits **gymnemic acid**, the compound from *Gymnema sylvestre* that binds the sweet receptor and switches sweetness *off* for many minutes — sugar turning to sand in the mouth. We keep it on the bench not to use but as a probe: if a candidate sweetness vanishes under gymnemic acid, we have confirmed it genuinely runs through TAS1R2·TAS1R3 and not some other route. Between the sweet switch and the bitter switches, the whole axis can be interrogated from outside the cell.

There are, then, two ways to deal with a passenger, and we hold a strong and stated preference for the second.

The first is **masking**: add something — a salt, a nucleotide, a blocker — that damps the off-note's receptor without killing the sweet one. It works, it is widely used, and it is a patch. You are adding an ingredient to hide a problem rather than removing the problem, which means one more line on the label and one more thing that can drift from batch to batch and lot to lot.

The second is **designing the trigger out**, and this is where IFI-S meets **IFI-B, the Zero-Bitter axis** — TAS2R38 and the wider bitter-receptor family — whose remit is bitterness attenuation in the **20–65%** range. Because we can engineer a crop's bitter chemistry at the source, we can deliver a sweetness that never raises the bitter signal in the first place. Run the mixture-suppression arithmetic forward to see what that buys. Return to $S_{\text{perceived}} = S_0(1 - \beta B)$ with the illustrative weight $\beta = 0.4$ and a passenger at $B = 0.5$, which cost us a fifth of the sweetness. Attenuate that bitter drive by the platform's upper reach, 65%, so the residual is $B = 0.5 \times (1 - 0.65) = 0.175$; the suppression falls to $0.4 \times 0.175 = 0.07$. The loss drops from 20% to 7%, and most of the sweetness that the passenger was quietly stealing simply **returns** — recovered not by adding sweetener but by removing the thing that was subtracting it.

HOST	CONSTRUCTS	PRIMARY AXES	THE NATIVE PASSENGER	THE DESIGN MOVE
Stevia	910	S / F / M	Stevioside bitterness (TAS2R4/14)	Shift output to Reb M at the UGT step
Monk Fruit	80	S	Slow, faintly medicinal tail	Favor the cleaner mogrosides
Sugarcane	200	S	The reference sweetness — near-silent	Hold the clean baseline, raise yield of loudness

The cleanest stevia is not stevia with a masker bolted on top. It is a stevia line bred and edited to make rebaudioside M and not stevioside, so there is **no bitterness to mask** — nothing on the label, nothing to drift, nothing to manage across a thousand batches. On the index it reads as a cultivar carrying a high IFI-S beside a correspondingly high IFI-B: loud sweetness, low bitterness, the two numbers printed side by side as the honest fingerprint of the thing. A trigger removed is worth more than a trigger masked, because an absent bitterness never has to be managed at all. That is the whole conviction of this work, said plainly: an off-note is just another axis. We measure it. And then we turn it down. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 7 · Body without sugar — *building the floor back.*

Taste a high-potency sweetener dissolved in plain water and you will learn, in one sip, the deepest problem in sugar reduction — deeper than any off-note, deeper than any metallic tail. The sweetness is *there*, correct and loud, and yet the whole thing tastes wrong: thin, hollow, fleeting, a note played with no instrument behind it. The mouth was promised richness and handed a signal. It is the whole problem of sugar reduction arriving, undisguised, on the tongue. Sugar gave a food a *floor* — body, viscosity, a dissolved-solids presence the mouth reads as richness and continuity — and the sweetener, at parts per million, gives sweetness standing on nothing.

Rebuilding that floor is the quiet, decisive craft of sugar reduction, and the secret we keep returning to at our own table is that it is mostly not about sweetness at all. The sweetness is the easy half: the Index can set it to a number and hold it there, batch after batch. The hard half is everything the sugar was *also* doing while it happened to taste sweet — and almost all of that lives on a different axis, with a different receptor, a different family of molecules, and nearly no taste of its own. Its name is kokumi. So this chapter about sweetness is, honestly, a chapter about a sweetness's whole neighborhood: the floor beneath it, the acid against it, the warmth it is served at, the savory depth that lets it last. We do not design a sweetener. We design a *sweetness* — and a sweetness, to feel like one, needs a body to live in.

The hollow note — a sweetness standing on nothing

Pour a glass of a sugar-sweetened beverage and you are drinking, by mass, mostly sugar. Something near a tenth of its weight — on the order of a hundred grams in every liter, which for sucrose at 342.3 g/mol is roughly 290 millimoles — a genuine crowd of molecules jostling in the water. That crowd is the floor. It thickens the liquid a little, slows it across the tongue, raises its density and the way it bends light, and gives the mouth a steady physical presence to read as *body* in the moments before and after the sweet taste itself ever fires. Now match that same glass for sweetness with a steviol glycoside or a monk-fruit mogroside — compounds on the order of two hundred times sweeter than sucrose by weight — and you have swapped a hundred grams of dissolved solids for a few tenths of a gram. The sweetness survives the trade intact. The floor does not. Two-hundred-odd times less mass is now in the water, and the mouth, which was never only tasting sweetness, feels the subtraction as thinness.

The reason the two cannot quietly stand in for each other is written into the receptor. Sweetness is read by **TAS1R2·TAS1R3**, a heterodimer — two different proteins clasped together — in the class C family of G-protein-coupled receptors, the same structural family as the kokumi receptor we are about to meet. Each half carries a large hinged **Venus flytrap domain**, and that architecture is why a

single receptor can be sweet to such wildly different things: sugars and small sweeteners settle into the flytrap cleft of TAS1R2; other ligands bind the transmembrane bundle; sweet *proteins* like brazzein dock against the outside of the domain entirely. Different keys, different keyholes, one lock — and the lock does not care how much mass delivered the key. Sugars reach half-activation at tens of millimolar; the high-potency sweeteners get there at micromolar or below, a potency gap of three to four orders of magnitude. That gap is exactly why one can be swapped for the other at a thousandth of the dose, and exactly why the floor vanishes when you do.

Loudness, though, is not quality, and the receptor keeps a second secret the mouth hears plainly: *timing*. Sugar binds and releases on gentle kinetics — it ramps up, rounds over, trails away. Many high-potency sweeteners snap on faster, cling longer to their site, and leave a tail the tongue reads as lingering or, worse, as a second, wrong taste arriving late. A sweetness is not a height; it is a shape in time. Restoring the floor is partly about restoring that shape — giving the sweetness a slow attack and a clean, carried finish instead of a spike over emptiness.

The sweet receptor, exactly — and the number we hang on it

When a sugar seats in the flytrap of TAS1R2·TAS1R3, the domain closes like a book around a finger and the motion is transmitted down through the cysteine-rich stalk to the seven-helix transmembrane bundle, which pivots against the taste cell's inner membrane and lets go of its G-protein. That G-protein is **α-gustducin**; freed, it drives **phospholipase C-β2**, which splits a membrane lipid into **IP_3**; the IP_3 opens stores that flood the cell with calcium; the calcium gates **TRPM5**, a cation channel, and the cell depolarizes and releases ATP onto the nerve. Sweetness, as a sensation, is that whole cascade firing — a lock and key repeated a billion times a second and summed into a single bright report. We can feel the result exactly. For ten thousand years we could describe it only in adjectives.

The Index ends the adjectives by hanging one number on the first step of that cascade: not how much sweetener a food holds, but how *loudly* it speaks to the receptor. The anchor is the half-activation point — the **EC_{50}**, the concentration at which a stimulus drives its receptor to half of full response — of a fixed reference, scored at exactly 1,000 by convention. Everything else is read against it.

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

Read the formula as a sentence and it is almost obvious. A candidate that reaches half-activation at a *lower* concentration than the reference is more potent — louder — and its EC_{50} sits in the denominator, so its score climbs above 1,000. A candidate that needs *more* concentration to be heard scores below 1,000. Potency and loudness are the same quantity seen from two sides, and the inversion is the whole trick: it turns a concentration, which falls as things get stronger, into a score that rises.

The reference is the point where the two EC_{50} values are equal, so every axis of it reads 1,000 by convention — the plain wild-type tongue's-eye view of the world, before any design:

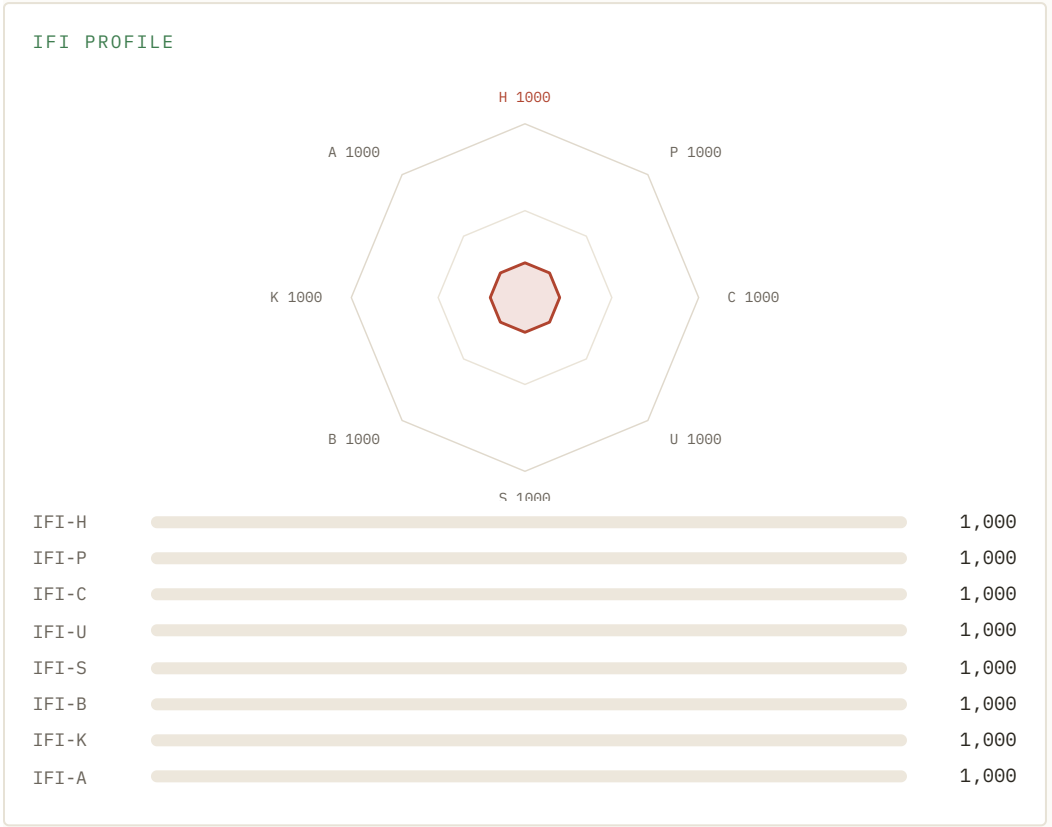

There is a ceiling, and it matters for sweetness more than for almost any other axis, because sweetness is where the most potent ligands in nature live. We normalize each axis to a maximum of **5,000**. By the formula, a score of 5,000 means $EC_{50}(c) = EC_{50}(ref) / 5$ — a ligand five times as potent as the anchor. Past that the number clips, and the sweet proteins make the clipping vivid. Brazzein (UniProt P56552), which binds the outside of the flytrap rather than its cleft, runs on the order of a thousand times sweeter than sucrose by weight; a ligand that potent would carry an EC_{50} perhaps a thousandth of the reference's, and the raw formula would return $1000 \cdot EC_{50}(ref) / (EC_{50}(ref)/1000) = 1{,}000{,}000$. It reads 5,000. Beyond roughly five-fold loudness the tongue itself saturates, more potency buys no more perceived sweetness, and reporting a million would be reporting a difference the mouth cannot feel. The ceiling is not a limitation of the scale; it is an honesty about the receptor, which has a loudest it can go.

Working the index on a sweetness

The arithmetic is worth doing by hand once, because it is the arithmetic of every reformulation. Take a candidate exactly twice as potent as the reference — half the EC_{50}. Then $IFI = 1000 \cdot EC_{50}(ref) / (EC_{50}(ref)/2) = 2{,}000$. Make it five times as potent and you reach the 5,000 ceiling. Run it the other way, from score back to concentration, and the Index tells you the physical state of the molecule in the food:

READING ON IFI-S	IFI	EC_{50} ÷ EC_{50}(REF)	WHAT IT MEANS AT THE RECEPTOR
Reference anchor	1,000	1.000	the convention half-activation point
Phoenix sweetness	1,100	0.909	~10% more potent than the anchor
A depleted tomato's umami floor	300	3.333	a whisper — more than 3× the reference concentration to be heard
A restored tomato's umami floor	2,600	0.385	better than twice as loud as the anchor
Phoenix kokumi floor	1,900	0.526	nearly twice the anchor — the floor louder than the sweetness it carries
Normalization ceiling	5,000	0.200	five times the anchor; the tongue saturates

That tomato line is the method in miniature, and it is pure canonical arithmetic: a shipping-depleted tomato sat at IFI-U 300 — an effective EC_{50} more than three times the reference, a savoriness you had to strain to hear — and our model predicts the restored plant at 2,600. The ratio 2,600 / 300 is 8.7, meaning the designed route drops the effective EC_{50} by nearly ninefold. Nothing was invented between those two numbers; a weak signal was made loud, and the formula simply records how much.

Dose-response is the curve the EC_{50} names a single point on. A receptor's fractional response follows the Hill form, $\theta = [L]^n / (EC_{50}^n + [L]^n)$, and at $[L] = EC_{50}$ the two terms are equal and $\theta = ½$ — half-max, by construction, which is the whole reason EC_{50} is the honest anchor rather than threshold or saturation, both of which wander with conditions. Step off that midpoint and the arithmetic is worth seeing: for a simple slope of $n = 1$, doubling the concentration to $[L] = 2 \cdot EC_{50}$ gives $\theta = 2 / (1 + 2) = 0.67$, and halving it to $EC_{50}/2$ gives $\theta = 0.5 / 1.5 = 0.33$ — a fourfold swing in dose moving the response only from a third to two-thirds of full, the gentle ramp of sugar. Steepen the exponent to $n = 2$ and that same doubling lifts θ to $4 / (1 + 4) = 0.80$: the sweetness that snaps on and is part of why some high-potency sweeteners feel abrupt. The slope, in other words, is the part the mouth feels as *shape*. Set the sweetness to its number and you have fixed the midpoint; you have not yet fixed the curve around it, and the floor is largely how you fix the curve.

The last piece of the math is the one a formulator lives or dies by: **error**. Our batches are designed to a coefficient of variation of **≤7%**, against **42%** in the wild-type plants our designs replaced. Because the reference EC_{50} is fixed and the Index is linear in $1/EC_{50}(c)$, the score inherits the measurement's relative error almost one-for-one: a sweetness specified at IFI-S 1,100 carries about ±77 at a 7% CV — a true window of roughly 1,023 to 1,177, tight enough to pour the same drink twice. At 42% that same target swings ±462, from near 640 to past 1,560 — not one sweetness but a dozen, batch to batch. The floor cannot be built on a number that will not hold still, which is why consistency is not a virtue we add at the end but a precondition for designing a sweetness at all.

Kokumi — the floor is a different receptor

Here is the move that makes the floor designable rather than mystical: it has its own receptor, its own molecules, and its own axis. Kokumi, **IFI-K**, is read by the **calcium-sensing receptor, CaSR** — another class C GPCR, structural kin to the sweet receptor, carrying the same hinged Venus flytrap architecture. Its agonists are small **γ-glutamyl peptides**: **glutathione** (γ-Glu-Cys-Gly), the familiar one, and, far more potent, **γ-glutamyl-valyl-glycine** (γ-Glu-Val-Gly), both built around the unusual γ-linkage off glutamate's side chain that ordinary digestion is slow to cleave. They occur in aged, fermented, and long-cooked foods — the gentle depth of a matured cheese, a long broth, a well-made soy sauce — which is to say the floor was always there in traditional cooking; we are only naming it and learning to dose it.

The remarkable part, and the reason kokumi sits at the center of a book on sweetness, is *how* these peptides act. They are **positive allosteric modulators**: they bind CaSR not at the orthosteric site where its true agonists (calcium, certain amino acids) land, but at a separate site, and from there they left-shift the receptor's sensitivity to whatever is already present. They do not pull the trigger; they lighten it. On their own, in plain water, they are nearly silent — almost no taste of their own. Set them beside a sweetness, a saltiness, a savoriness, and they *amplify and prolong* it. This is exactly the right tool for a sugar-free sweetness, because the sweetness is not weak — it is thin and brief — and a modulator that adds body and duration without adding a flavor of its own is precisely a floor.



Soybean — one of our kokumi hosts, and a natural reservoir of the γ-glutamyl peptides that act on CaSR. The richness that makes a sugar-free sweetness feel whole is engineered here, on a separate axis from the sweetness itself.

The sensory literature names what kokumi contributes with real precision, and a sugar-free sweetness usually fails on the third of the three: **thickness** (the sense of body), **mouthfulness** (how fully the flavor occupies the mouth), and **continuity** (how seamlessly it carries from first taste to swallow). A sweetener spikes and vanishes; a sugar carried. Continuity is the floor's signature, and it is what a kokumi base restores. We build that base by engineering the peptides' biosynthesis — the γ-glutamyl route through glutamate-cysteine ligase and the γ-glutamyltransferases — directly into hosts that already lean toward richness, and then we profile the result on IFI-K so that "fuller" stops being an adjective and becomes a dose. These are the plants we reach for:

HOST	CONSTRUCTS	PRIMARY AXES	ROLE IN A SWEETNESS	DIFFICULTY
Tomato	2,040	S / U / F / M	the sweet-and-savory floor in one fruit	Easy
Stevia	910	S / F / M	the sweetness itself	Moderate
Garlic	560	K / P	a pungent-edged kokumi floor	Moderate
Soybean	360	U / K	the broad savory-rich floor	Frontier
Yeast	240	U / K	a clean, neutral kokumi base	Hard
Sugarcane	200	S	the sweetness itself	Frontier
Monk Fruit	80	S	the sweetness itself	Hard
Onion	40	K	a sweet-savory floor	Moderate
Vanilla	40	K	kokumi as mouthfeel and roundness	Hard

A note on honesty before we build on it: because kokumi is a *modulator*, IFI-K is a measure of modulating power, not of a taste you could point to alone. A high IFI-K peptide over plain water still does almost nothing — there is nothing present for it to lengthen. The floor is only a floor under something. That is a mechanism, not a flaw, but it is the first thing to understand before specifying a formula, because it tells you the order of operations.

The three physical layers — bulk, hydrocolloid, kokumi

A serious reconstruction of sugar's floor uses three layers, and they are not interchangeable, because they fix three different absences.

Bulking agents replace the dissolved solids themselves — the physical hundred grams. Allulose (a rare sugar), erythritol (a polyol), and soluble fibers such as inulin and resistant dextrin put mass back in solution, restoring density and some viscosity; allulose, being a reducing sugar, also browns through the Maillard reaction, so a baked or glazed food gets its color back too. This is the structural subfloor — it fixes *volume*. **Hydrocolloids** fix *flow*: pectins and specific gums tune viscosity directly where bulk is not enough, and they carry a trade-off a formulator must respect, because viscosity also slows flavor release. A thicker liquid holds its aroma and its sweetener back from the receptors — diffusion to the taste cells scales roughly as the inverse of viscosity — so past a point more gum makes a drink feel fuller while tasting *less* sweet. The art is a window: enough hydrocolloid for body, not so much that it mutes the sweetness you are building. **Xanthan** is prized here for its shear-thinning habit — thick at rest, thin as it crosses the tongue — which gives body without the muting, because the viscosity that builds the floor falls away at the moment of tasting.

LAYER	AGENTS	RESTORES	ACTS THROUGH
Bulk	allulose · erythritol · inulin · resistant dextrin	dissolved solids, density, browning	mass in solution
Hydrocolloid	pectin · xanthan · specific gums	viscosity, flow, the tactile "round"	shear-dependent rheology
Kokumi	γ-Glu-Val-Gly · glutathione	perceived fullness, continuity, duration	CaSR (positive modulation)



Sugarcane — the sweetness itself, and a host on the Index. The grass that fed the world's sweet tooth, now read against the same half-activation anchor as every other axis, so a sweetness can be specified rather than guessed.

Kokumi does the perceptual work on top of the other two, via CaSR, making the whole reach of the sip feel continuous. And because the three layers interact, the *order* in which you build them is itself a method: set the sweetness to its number first; build continuity with kokumi; adjust thickness with bulk and hydrocolloid; balance the acid last. Each later step changes how the earlier ones read, so the sequence runs from the hardest-to-move to the easiest-to-trim. Build it out of order and you spend the afternoon chasing a sweetness that keeps changing underneath you, because you were adjusting the floor while standing on it.

The cross-axis arithmetic — a sweetness is built from its neighbors

There is one more set of interactions that governs how sweet a finished food actually reads, and once you know them they are a gift, because they let you spend less sweetener for the same sweetness. Sweetness is strongly shaped by the axes it sits among. **Acidity suppresses it** — a tart lemonade needs more sweetness to balance, which means trimming a formula's sourness can *replace* a share of its sweetener outright. **Salt, in small amounts, enhances it. Umami and kokumi lengthen and round it**, which is why the deepest clean sweetnesses we make sit over a savory floor rather than beside a sour one — a glaze, a drinking chocolate, a slow-cooked tomato. **Warmth lifts it**, because TRPM5, the channel at the end of the sweet cascade, is itself temperature-sensitive, so the identical formula reads sweeter hot than cold — a genuine lever, and, served cold, a genuine trap. And **aroma can read as sweetness with no sweetener present at all**: the warm, ripe smell of a tomato arrives at the brain pre-labeled "sweet" from a lifetime of ripe fruit, and the mind adds sweetness the tongue never tasted. None of these is sugar. All of them are sweetness.

The cleanest proof that these layers can live in one organism is a plant we designed to register on every axis at once — and reading only its relevant lines is a lesson in composition:

IFI PROFILE

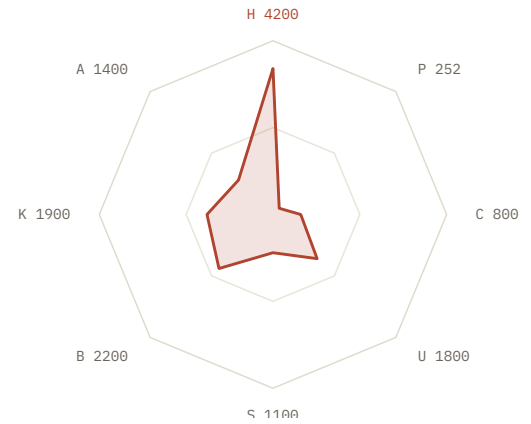

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

Ignore the heat for a moment and look at the chord underneath it. The sweetness sits at **S 1,100** — a touch louder than the anchor. But it does not sit alone: beneath it is **K 1,900**, a kokumi floor nearly twice the anchor's loudness, and **U 1,800** of umami, a savory depth that lengthens continuity; and holding the whole thing clean is **B 2,200** on the zero-bitter axis — bitterness attenuated by the designed route that trims the off-notes a sugar-free system would otherwise expose. A modest sweetness carried on a loud floor, with the bitterness pulled down so nothing muddies it: that is not a sweetener, it is a *sweetness*, and the Index shows it as a vector rather than a single number. The floor, designed at K 1,900 against a sweetness of S 1,100, is literally louder than the thing it supports — which is the whole lesson of this chapter rendered as arithmetic.

Phoenix and the honest limits — what a number can and cannot promise



Vanilla, in cross-section — the world's most complex and most faked flavor, and on our bench a kokumi host. Its contribution to a sweetness is not more sweetness but mouthfeel: roundness, a slow finish, the sense that the sweetness has somewhere to rest.

We are careful, at our table, about what a number is allowed to promise, because the discipline of the Index is only worth as much as its honesty. Several limits are worth stating plainly.

The viscosity window is real and unforgiving: build too much body with gum and the sweetness you worked to place quietly mutes, because the same thickness that reads as rich also holds the sweetener back from the receptor. Temperature is a second edge — a formula tuned warm, where TRPM5 runs hot and sweetness lifts, can read flat served over ice, so a cold product must be sweetened to its serving temperature, not to the bench. Kokumi has the limit already named: a modulator with nothing to modulate does nothing, so a kokumi base cannot rescue a sweetness that is otherwise empty; it can only enrich one that is present. And the error budget, generous as ≤7% is against a wild-type's 42%, is not zero — a sweetness specified to the last point is specified past what the measurement can honestly hold, and we round to the precision the CV supports.

There are limits the Index does not even try to cross, and naming them is part of respecting it. It scores loudness at the receptor, not *liking* — hedonics are a separate, personal thing, and a perfectly placed sweetness can still be a sweetness someone does not care for. Bitter perception varies person to person with the genetics of the bitter receptors, so the same attenuated coffee lands differently on different tongues, and a single number cannot carry that spread. The temporal shape — the attack and the tail that the mouth feels so acutely — is only partly captured by a midpoint EC_{50}, and remains a place where the craft still outruns the measurement. The Index is a floor for design, not a ceiling on it.

What it does promise, it keeps. Across roughly nineteen thousand design packages spanning forty-two host organisms — with the method set down across twenty-one peer-reviewed papers and each design held to the same ≤7% CV specification — a sweetness can be written as a specification and reproduced to spec. The sweet axis is never ordered alone on our platform. Every sweetness ships inside a small composition: a floor of kokumi over CaSR, a managed acid, a whisper of salt, the right aromatics, the warmth it will be served at. We do not sell a sweetener. We design a *sweetness* — with everything a sweetness needs to feel like one, and with the particular care of people who believe food is something you give the body on purpose, grown clean, built to nourish, and made to taste the way the living plant always meant it to.

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



PART III OF V

The Plants We Teach to Be Sweet

stevia · monk fruit · sugarcane · the sweet proteins · the tomato

Chapter 8 · Stevia — *the leaf that out-sweetens sugar three hundred times.*

A stevia leaf is a small, unassuming thing — toothed, soft green, a shrub from the highlands of Paraguay and Brazil where the Guaraní called it *ka'a he'ê*, sweet herb, and steeped it for centuries. Chew one and the sweetness arrives slowly and enormously, hundreds of times louder than the sugar in any crop, and then, if the leaf is an ordinary one, it turns — a licorice edge, a lingering bitterness, the flaw that has kept stevia a compromise through a hundred years of food science. Our family came to this leaf the way we come to every plant in the library: wanting to give the body a pleasure cleanly, a sweetness with nothing hidden in it. This chapter is about where that sweetness is made inside the leaf, where the bitterness rides along with it, and how — molecule by molecule, glucose by glucose — the two can be separated. Of all the plants we work on, stevia is the one where the biosynthesis *is* the product: the pathway and the flavor are the same object, read from two ends.

How the tongue reads sweet — TAS1R2·TAS1R3

Before we touch the plant, we have to be exact about the lock the leaf is built to fit. Sweetness is read by a single receptor: a heterodimer of two proteins, **TAS1R2 and TAS1R3**, sitting together in the membrane of the sweet-sensing taste cells. It belongs to the class C G-protein-coupled receptors, the same family as the glutamate and calcium sensors, and like its cousins each half carries a large clamshell on its outside — a **Venus flytrap domain**, two lobes hinged around a cleft. Sugars settle into the cleft of the TAS1R2 lobe and pull it closed; the hinge motion is the signal. This is the elegance and the vulnerability of the system at once: the receptor does not count molecules, it reads a *shape* closing a *hinge*, and anything of roughly the right geometry — a sugar, a sugar alcohol, a diterpene dressed in glucose, even certain small proteins wedged against the TAS1R3 lobe — can throw the latch.

What follows the latch is a cascade we design around as carefully as the binding itself. The closed receptor couples to **α-gustducin**, which frees a βγ pair that drives **phospholipase C-β2**; PLCβ2 cleaves a membrane lipid into **IP_3**; IP_3 opens stores that flood the cell with calcium; the calcium gates **TRPM5**, a monovalent cation channel, which depolarizes the taste cell; and the depolarization opens **CALHM1** to release ATP onto the afferent nerve. Eight molecular steps between a leaf on the tongue and the word *sweet* in the mind. We do not touch that cascade — it is the body's, and we revere it — but knowing its shape tells us where our only lever is: the very first step, the fit of the ligand in the flytrap. Everything the platform does to a sweetener, it does to make that first closing cleaner and louder.

Loudness is the word we mean precisely. The Index score on any axis is not how much of a compound a food carries but how *loudly* that compound speaks to its receptor, and loudness is the inverse of a single measured number — the **EC_{50}**, the concentration at which a stimulus drives its receptor to half of full response. A low EC_{50} means the receptor is half-lit by a whisper of the compound; that compound is loud. The Iyengar Flavor Index turns that whisper into a score with one formula:

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

We fix the reference sweet stimulus as a sucrose solution — table sugar, the thing every tongue already calibrates against — and read every other sweetener against it. This is a stated convention, like the Scoville scale's dilution-to-threshold definition: each reference constant sets IFI = 1,000 on that axis by convention. Invert the EC_{50} ratio, scale by a thousand, and a sugar, a stevia leaf, and a monk-fruit extract can be laid side by side on one page.

Sweetness, though, makes the Index work harder than any other axis, and it is worth saying why. Heat spans a narrow band — even the loudest engineered pepper is only a few times louder at the receptor than a hot wild-type reference — so its scores stay small and nearly linear. Sweetness spans three orders of magnitude: an intense glycoside is not three times louder than sugar but a *thousand* times louder. A raw inverse-EC_{50} score would run past a million and off the edge of any page. So the scorecard value you will see in these pages is the raw loudness compressed into a fixed **0–5,000 band**, by the same logarithmic law the tongue itself obeys — Weber and Fechner's old observation, later sharpened by Stevens, that perceived intensity rises with the logarithm of the stimulus, not with the stimulus itself. We write the compression plainly:

$$\text{scorecard IFI-S} = 1{,}000 + 1{,}000 \times \log_{10}\left(EC_{50}(reference) \div EC_{50}(candidate) \right)$$

so that sucrose, the reference, lands at exactly 1,000, a thousandfold-louder glycoside lands near 4,000, and the band saturates at 5,000. The raw ratio is the physics; the scorecard is the physics read the way a person tastes it.

And here is the whole trouble with stevia, stated in receptor terms. The sweet molecules of the leaf do not speak only to the sweet receptor. The same diterpene core, at the same doses, also throws the latch on two **bitter** receptors — **TAS2R4** and **TAS2R14** — and the tongue reads both messages at once, sweetness and a licorice bitterness arriving together. One ligand, two receptors, two words. The leaf's famous flaw is not a contaminant to be filtered out. It is the sweet molecule talking to the wrong lock. Which means the fix cannot be a filter. It has to be chemistry.

Where the sweetness is made — GGPP to steviol

The sweet molecules are the **steviol glycosides**: a steviol core, which is a diterpene, hung with glucose units. The core is not built from sugar at all. It comes from the plant's isoprenoid machinery — the plastidial **MEP pathway**, the same assembly line that makes the

backbones of carotenoids and chlorophyll — which converges on a twenty-carbon building block, **geranylgeranyl diphosphate (GGPP)**. From GGPP, a defined cascade folds and oxidizes the carbon skeleton into steviol, and then a second cascade dresses it in glucose.



Stevia specimen, isolated. The sweetness is a diterpene dressed in glucose; the bitterness is a matter of exactly how many glucoses, and in which linkage. We breed and edit for the latter.

STEP	ENZYME	REACTION	PRODUCT
1	CPS	cyclization	*ent*-copalyl diphosphate
2	KS	cyclization	*ent*-kaurene
3	KO (CYP701A)	three oxidations	*ent*-kaurenoic acid
4	KAH (CYP716D)	C-13 hydroxylation	**steviol**
5	UGT85C2	glucosylate C-13	steviolmonoside
6	UGT91D2	β(1→2) glucose	steviolbioside
7	UGT74G1	glucosylate C-19	**stevioside**
8	UGT76G1	β(1→3) glucose	**rebaudioside A**
9	UGT91D2	β(1→2) glucose	rebaudioside D
10	UGT76G1	β(1→3) glucose	**rebaudioside M**

Two facts in this table do the work of the whole chapter. The first: steps 1 through 4 are not stevia's at all — they are shared, residue for residue, with **gibberellin** biosynthesis, the pathway that makes the plant's growth hormones. Stevia branches off at step 4, where the enzyme KAH hydroxylates carbon 13 of *ent*-kaurenoic acid and sends it toward steviol instead of toward the hormones. The leaf's sweetness is, quite literally, a diverted hormone precursor. That shared stem is why the sweetness can never be made free: every steviol molecule is drawn from a pool the plant also needs to grow, and we will have to pay for it in the arithmetic of flux.

The second fact: the last six steps, the ones that decide everything about taste, are all performed by a small family of **UDP-glycosyltransferases** — **UGTs** — adding glucose units one at a time, in a specific order, with specific chemical linkages. β(1→2) here,

β(1→3) there. The identity of the final glycoside, and therefore both its sweetness *and* its bitterness, is decided entirely by which glucoses get added where. The leaf is not making one sweetener. It is making a ladder of them, and where it stops on the ladder is the whole game.

The off-note is an earlier molecule

Walk the ladder with the two receptors in mind. **Stevioside** (step 7), the dominant glycoside of an ordinary leaf, is intensely sweet — and it is also the loudest talker to TAS2R4 and TAS2R14. That is the licorice tail. **Rebaudioside A** (step 8) carries one more β(1→3) glucose and is cleaner, which is why it became the industry's first workhorse. Keep climbing — add the β(1→2) glucose of step 9 to reach **rebaudioside D**, then the final β(1→3) of step 10 to reach **rebaudioside M** — and the molecule grows two more sugars, loses almost all of its grip on the bitter receptors, and keeps very nearly all of its sweetness.

So the off-note is not an impurity sitting *beside* the good molecule. It is an *earlier molecule on the same pathway.* Stevioside is bitter because stevioside is unfinished. And that rewrites the entire design goal, cleanly and almost gratefully: do not try to grow more steviol glycoside and then scrub the bitterness out. Grow the pathway **forward** — past stevioside, past Reb A, all the way to Reb D and Reb M — so the leaf's output is dominated by the finished, quiet glycosides and starved of the unfinished, loud ones.

GLYCOSIDE	TYPICAL WILD-TYPE LEAF	PLATFORM TARGET	BITTERNESS
Stevioside	~50%	~5%	high (TAS2R4/14)
Rebaudioside A	~25%	~10%	moderate
Rebaudioside C	~10%	~2%	moderate
Rebaudioside D	~2%	~25%	very low
Rebaudioside M	<1%	~55%	**very low**
others	~12%	~3%	—

Read the wild-type leaf on the Index and the problem is visible at a glance — a loud sweet axis dragged down by a loud bitter signal, a sweetness you cannot give a child without an apology:

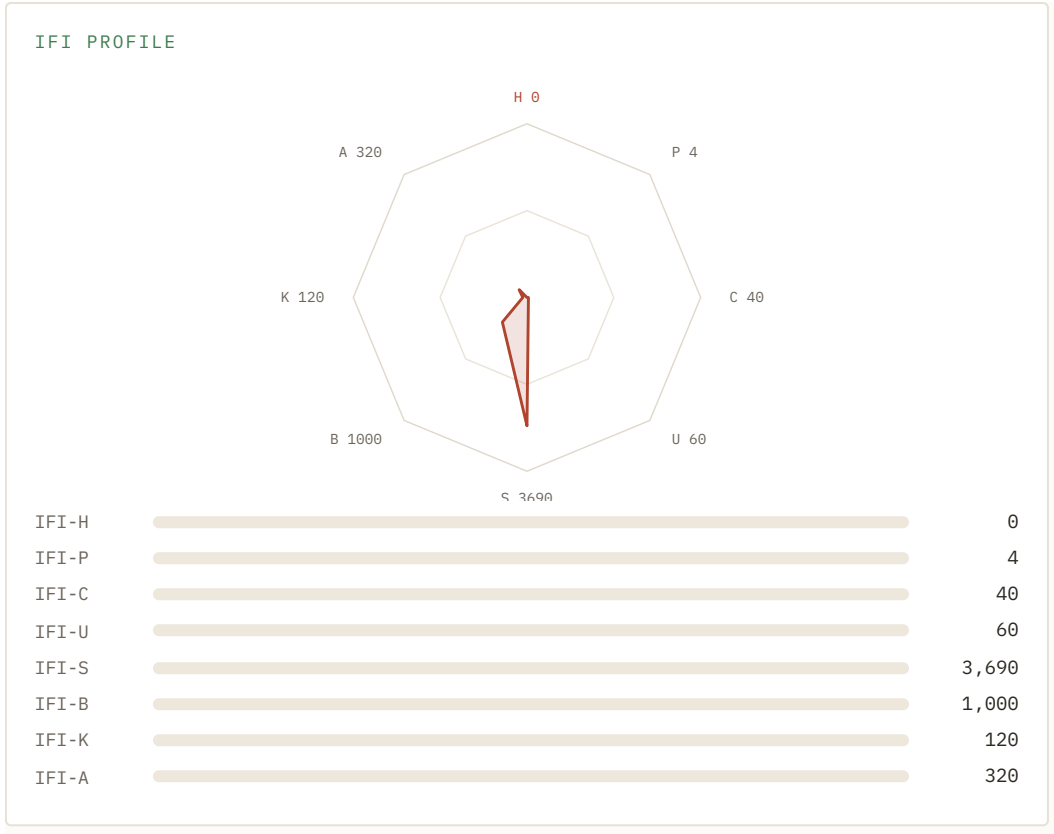

The sweetness (S) is already magnificent. The axis we have to move is **B** — freedom from bitterness — which the unfinished leaf scores at the floor. The rest of this chapter is the arithmetic and the engineering of lifting that one number without losing any of the first.

The mathematics of a clean sweet

Start with the sweet receptor and the numbers our calcium-flux panel reads. We express every sweetener against a sucrose reference whose half-maximal concentration is **EC_{50} = 20 mM** in our assay, and we measure each glycoside's own sweet EC_{50} the same way, in human TAS1R2·TAS1R3 expressed in cultured cells, reading calcium as the output.

The conversion that surprises people is the unit one. The old food-science shorthand — "three hundred times sweeter than sugar" — is a *weight* ratio, and the receptor does not weigh molecules, it counts them. A steviol glycoside is a heavy thing next to sucrose (MW 342.30), so its per-molecule loudness is larger still. Take Reb M, roughly 300× sweeter than sucrose by weight, molecular weight 1,291.29:

> `molar potency = 300 × (1,291.29 ÷ 342.30) = 300 × 3.772 =` **`1,132×`**

On a per-molecule basis Reb M is not 300 times louder than sugar but about eleven hundred. Its sweet EC_{50} follows directly: 20 mM ÷ 1,132 = **17.7 µM**. Feed that into the Index formula and the raw loudness is

> `IFI-S(Reb M) = 1,000 × (20,000 µM ÷ 17.7 µM) = 1,000 × 1,132 ≈` **`1,131,900`**

— a true number, and an unreadable one, which is exactly why the scorecard compresses it: $1{,}000 + 1{,}000 \cdot \log_{10}(1{,}132) = 1{,}000 + 3{,}054 \approx$ **4,050**. Do the same arithmetic down the ladder and a quiet pattern appears.

GLYCOSIDE	MW (G/MOL)	SWEET EC_{50}	MOLAR POTENCY VS SUCROSE	RAW IFI-S	SCORECARD IFI-S	BITTER EC_{50} (TAS2R4/14)	IFI-B	BITTER OCCUPANCY AT EQUI-SWEET DOSE
Sucrose (ref)	342.30	20 mM	1×	1,000	1,000	—	—	–
Stevioside	804.87	40.5 µM	494×	493,700	3,690	70 µM	1,000	36.6%
Rebaudioside A	967.01	28.3 µM	706×	706,300	3,850	180 µM	2,570	13.6%
Rebaudioside D	1,129.15	21.7 µM	924×	923,700	3,970	650 µM	5,000 (cap)	3.2%
Rebaudioside M	1,291.29	17.7 µM	1,132×	1,131,900	**4,050**	2,500 µM	5,000 (cap)	**0.7%**

Look at the sweet column first. Stevioside to Reb M, the scorecard sweetness climbs only from 3,690 to 4,050 — a span of under ten percent across the whole ladder. At the receptor, the finished glycoside is barely louder than the unfinished one. This is the honest and slightly humbling heart of the stevia problem: **you cannot taste your way up this ladder by sweetness.** The leaf that fixes the flaw is not meaningfully sweeter. Everything that matters is happening on the other receptor.

So read the bitter columns the same way, and notice that the arithmetic has to invert. On the sweet, umami, or heat axes, loudness is the thing you want, and a *lower* EC_{50} scores *higher*. Bitterness is the thing you want *gone*; the IFI-B axis measures freedom from it. So we anchor IFI-B to the loudest bitter glycoside — stevioside — and write the ratio the other way up:

$$\text{IFI-B(candidate)} = 1{,}000 \times EC_{50}\text{(candidate at TAS2R4/14)} \div EC_{50}\text{(stevioside at TAS2R4/14)}$$

Now a molecule that needs a *high* concentration to half-light the bitter receptor — one that barely talks to it — scores high. Stevioside, the reference, lands at 1,000. Reb A, with a bitter EC_{50} of 180 µM, lands at 1,000 × 180 ÷ 70 = **2,570**. Reb D and Reb M need 650 µM and 2,500 µM respectively to reach the bitter receptor at all — 1,000 × 650 ÷ 70 = 9,286 and 1,000 × 2,500 ÷ 70 = 35,700 — both far past the ceiling, both reported at the capped **5,000**. The one molecule you can grow forward is the one that goes silent on the lock you want silent.

The most useful number in the table is the last column, because it is the one your mouth actually reports. Sweet and bitter are not read at separate doses — the *same* molecule, at the *same* concentration, speaks to both receptors, and the fraction of bitter receptor it occupies follows the simplest dose–response law, the Hill equation with coefficient one:

$$\theta\text{(bitter)} = C \div (C + EC_{50}\text{(bitter)})$$

Set the dose C to each glycoside's own sweet EC_{50} — the concentration that gives you half of full sweetness — and ask how much bitter receptor comes along for free. Stevioside at 40.5 µM: θ = 40.5 ÷ (40.5 + 70) = **36.6%** of the bitter receptor lit, a third of the

bitter channel firing while you taste the sweet. Reb M at 17.7 µM: θ = 17.7 ÷ (17.7 + 2,500) = **0.7%**. The bitterness has not been covered. It is simply no longer spoken.

Push the worked example all the way to a glass, because that is where the number has to survive. Sweeten a beverage to the equivalent of a 10% sugar solution — 100 g/L of sucrose-equivalent sweetness, ordinary soft-drink territory, well above half-max and close to full sweetness. With Reb M at 300× by weight you need 100 ÷ 300 = 0.333 g/L, which is 333 ppm, which in molar terms is 333 mg/L ÷ 1,291.29 g/mol = **258 µM**. At that real dose, Reb M's bitter occupancy is 258 ÷ (258 + 2,500) = **9.4%**. Sweeten the *same* glass with stevioside — 100 ÷ 210 = 0.476 g/L = 592 µM — and its bitter occupancy is 592 ÷ (592 + 70) = **89.4%**. The unfinished leaf saturates the bitter channel nine times in ten; the finished leaf grazes it. The attenuation between them, at the dose a person actually drinks, is (89.4 − 9.4) ÷ 89.4 ≈ **90%**.

That 90% is worth setting against the established bitterness-attenuation band of **20–65%**, which is what masking-based approaches reach across coffee, cacao, and the brassicas — adding a blocker, balancing with another taste, hiding the signal. Stevia sits past that band not because our chemistry is better at hiding, but because it is not hiding at all. You cannot mask your way past the ceiling of masking; you can only get there by removing the ligand. Forward-pathway engineering removes it.



The sequencer bench where candidate glycoside profiles are read back — the leaf's chemistry returned as numbers, batch after batch, against a ≤7% tolerance.

One last piece of arithmetic keeps us honest about every number above: the error on them. The platform holds batch-to-batch variation to a coefficient of variation of **≤7%**, and an EC_{50} carries that same CV. Because the IFI is a ratio of two independently measured concentrations, the raw score's relative error is the quadrature sum, $\sqrt{(0.07^2 + 0.07^2)}$ = **9.9%** — call it ten percent on the raw ratio. The scorecard, though, is the logarithm of that ratio, and the logarithm compresses the error exactly as it compresses the range: a ±9.9% swing in the ratio becomes 1,000 × (1 ÷ ln 10) × 0.099 ≈ **±43 points** on the scorecard. Reb M's sweetness is 4,050 ± 43 — a little over one percent. The same compression that makes a thousandfold sweetener readable also makes it reproducible, which is why we report the band, not the point.

Our model predicts the designed leaf reads like this — the sweetness essentially unchanged from the wild leaf, the bitterness axis lifted off the floor and against the ceiling:

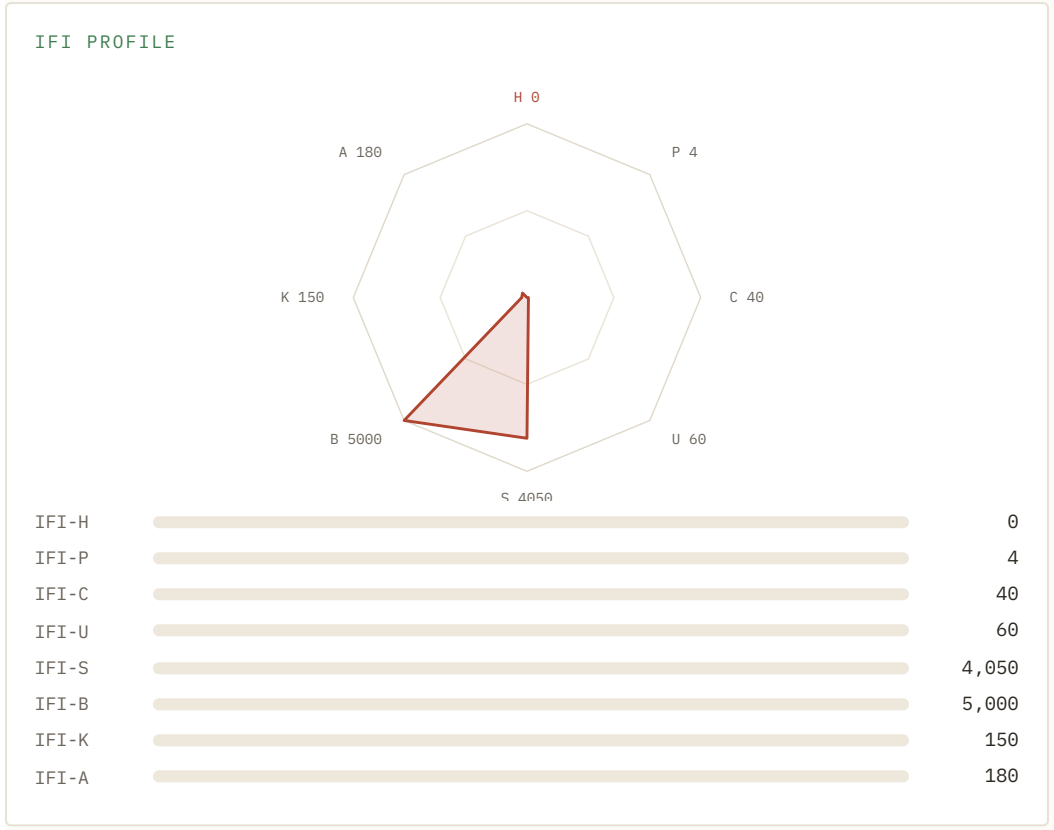

Two numbers moved. S from 3,690 to 4,050, inside the error bars of "as sweet as ever." B from 1,000 to 5,000, the entire height of the axis. That is the clean sweet, written as coordinates.

Two levers — selectivity and flux

Our stevia program runs **910 constructs**, and the work divides along exactly the two facts the biosynthesis table handed us: the UGTs decide the glycoside, and the shared stem decides the budget.

The first lever is **selectivity engineering** of the glycosyltransferases, and the workhorse is **UGT76G1** — the enzyme that performs both β(1→3) glucosylations, step 8 (making Reb A) and step 10 (making Reb M). We want it faster at that linkage and more specific for the Reb M product. Rather than simply overexpress the wild-type enzyme — which only piles up intermediates and can throw off a dozen minor glycosides — we redesign its acceptor pocket, the cradle that holds the growing glycoside while the new glucose is positioned. An **ESM2 650M-parameter protein-language model** scores every single-residue substitution in that pocket by its predicted effect on the β(1→3) rate and on Reb M specificity; the top-ranked edits are written in and tested. The β(1→2) step that builds Reb D — step 9, run by **UGT91D2** — gets the same treatment, because a leaf that reaches Reb D cleanly is a leaf poised to finish at Reb M.

The second lever is **flux balance**, and it carries a constraint the sweeter-is-louder axes never face. GGPP is not free. It is drawn from the one pool that also feeds carotenoids, chlorophyll, and the gibberellins the plant needs to grow — the shared stem from steps 1 through 4. Draw too hard toward steviol and you get a sweeter leaf on a sicker plant: pale, stunted, short of the pigments that capture light and the hormones that let it stand up. So the draw rate is not a dial we are free to turn to maximum. We hold the entire leaf's

metabolism in a genome-scale **flux-balance model** and solve for the partition that maximizes clean-glycoside yield *subject to* the plant keeping enough isoprenoid for its carotenoids, its chlorophyll, and its growth. The model finds the edge — the most sweetness the leaf can spend without starving itself — and stops there. A sweeter leaf that is still, unmistakably, a healthy plant.

PROGRAM MODULE	CONSTRUCTS	WHAT IT TUNES
UGT76G1 acceptor-pocket redesign — β(1→3)	420	Reb A → Reb M flux and specificity
UGT91D2 engineering — β(1→2)	180	Reb D formation
Flux-balance / GGPP partitioning	160	draw rate toward steviol without starving isoprenoids
UGT74G1 · UGT85C2 upstream priming	90	C-19 and C-13 first glucosylations
Controls, enzyme stacks, agronomic lines	60	validation and field trial
Total	**910**	

The two levers are not independent, which is the point of running them against one shared model rather than one at a time. A faster UGT76G1 is useless on a leaf that cannot afford the steviol to feed it; a generous GGPP draw is wasted on a UGT that leaks its flux into minor glycosides. The 910 constructs are how we find the pairs that agree.

Bioconversion — the same enzymes in a tank

There is a second way to make a clean steviol glycoside, and we keep it running beside the plant work because it teaches us faster than the plant can. The same UGT enzymes will run **outside** the leaf — in engineered yeast fed a simple sugar, or in a cascade of purified enzymes in a tank — carrying cheap, abundant stevioside or Reb A the last few steps into Reb D and Reb M. Most of the clean Reb M on the market today is made exactly this way.

We take two things from it. The first is a proof: bioconversion demonstrates, unanswerably, that the glycoside profile is **enzymatic, not varietal.** Give the right UGTs the right substrate and the bitter molecule becomes the clean one in a steel tank, with no plant involved at all — which means the leaf's flaw was never a fact about the leaf, only about which enzymes were busy. The second is speed. A candidate UGT76G1 variant can be expressed, purified, and run against a substrate in a tank in days, its product read on the same bench that reads the glycoside profiles, long before we would know how the edit behaves in a growing plant. So we test in the tank and build in the leaf. Our in-planta program takes the lesson of bioconversion and moves it upstream — rather than fix the profile after harvest, grow a leaf that made the right glycoside to begin with — but the fastest road to each new variant runs through the tank first.

Harvest, rules, and honest limits

A stevia leaf carries its glycosides at roughly a **tenth of its dry weight**, and the fraction is highest just before the plant flowers, after which the leaf begins spending its chemistry on reproduction. So

harvest timing is itself a lever on both yield and profile — pick early and you trade total mass for a cleaner, richer leaf; let it flower and the ratio slips. The agronomy and the chemistry are the same decision, made twice.

The regulatory ground is well-trodden, which is part of why we chose this leaf to prove the sweetness axis on. High-purity **Reb A** has been generally recognized as safe since **2008**, and **Reb M** since the later 2010s, each with a defined acceptable daily intake. A leaf whose output is dominated by Reb M and Reb D lands inside an established category rather than a novel one — a real advantage when the goal is to put a clean sweetness on an ordinary table, not inside a decade of filings.

But we would be poor stewards of our own numbers if we did not mark their edges, so here are the honest limits of everything above.

- **Receptor math predicts loudness, not liking.** An EC_{50} tells you how fully a molecule lights a receptor; it does not tell you the shape of the sweetness in time — the onset, the lingering, the way real tasters describe a sweetness as "slow" or "flat." Those are temporal and cognitive, and our cell assays are blind to them. Every construct that passes the EC_{50} panel still faces a human sensory panel, and the two do not always agree.
- **The residual is real.** Reb M at full beverage sweetness still grazes the bitter receptor at about 9% occupancy, and the most sensitive palates can read it. "Essentially clean" is not "undetectable," and we do not pretend otherwise; for the most demanding applications the lever is to raise per-molecule sweet loudness so the needed dose — and with it the residual — falls further.
- **A forward pathway can overshoot.** An over-active UGT does not only make more Reb M; it can make unintended minor glycosides the wild leaf never produced, and a molecule the food supply has not seen is a molecule that owes a toxicology review before it owes anything else. Specificity, not just speed, is why we engineer the pocket rather than flood the enzyme.
- **Flux has a floor.** The genome-scale model finds the edge, but the edge moves with light, temperature, and the plant's stage of life; a partition that is healthy in a controlled chamber can starve a field plant in a cold week. The model is a guide to the trade, not a guarantee against it.

It is worth placing stevia among its neighbors on this axis, because the lesson is the same wherever a plant makes an intense sweetness. Monk fruit (an **80-construct** program) builds its sweetness from a different chemistry entirely — cucurbitane triterpene glycosides, the mogrosides — yet carries the identical structure of problem: a glycoside ladder where the finished molecule is clean and the unfinished ones carry the off-note, and the same forward-pathway answer applies. Sugarcane (**200 constructs**) comes at sweetness from the opposite side, as the body's own sugar made more generously rather than more loudly, a frontier problem of yield rather than of receptor selectivity. And the sweet proteins, brazzein among them, throw the sweet latch by a different

mechanism altogether — wedging against the receptor rather than settling into its cleft — a reminder that the lock accepts more than one kind of key. Stevia is where the method is clearest because stevia is where the pathway and the flavor are most nearly the same object.



Inside the leaf — the mesophyll where the diterpene is built from GGPP and dressed, glucose by glucose, into the glycoside we grow forward.

The result, on the Index, is a stevia our model predicts with sweetness as loud as it ever was — a scorecard S near the top of its band — and freedom from bitterness high rather than low, an IFI-B lifted from the floor to the ceiling. A Reb-M-dominant leaf that reads clean in our computational design, with the licorice tail designed out at its source rather than masked at the mixing tank. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. That is the whole stevia problem solved the way our family likes to solve things, the way a vegetarian kitchen taught us to think about everything we grow: not by hiding the flaw, but by walking the pathway forward, one glucose and one linkage at a time, until the flaw is simply no longer made. The sweetness was never the problem. It was only ever unfinished.

Chapter 9 · Monk fruit — *the rounder sweetness.*

If stevia is a sharp, bright, slightly difficult sweetness, monk fruit is its rounder cousin — slower to arrive, fuller through the middle, with a clean exit and far less of the licorice argument. The fruit is *Siraitia grosvenorii*, a small green gourd from the mountains of Guangxi in southern China, where it has been trained on trellises in cool mist and dried slowly over low fires for centuries — *luo han guo*, the Buddha fruit, the monk's fruit. Its sweetness belongs to a different chemical family than stevia's, and that difference in chemistry is exactly the difference the tongue reports. We came to monk fruit the way our family comes to every plant in this book: first at the table, tasting what a thousand Guangxi harvests had already made of it, and then at the bench, asking the only question we ever really ask of a flavor — *how much, and why.* This chapter is the long answer, written out: the molecule, the enzymes that build it, the receptor it speaks to, the number we give it, and the honest edges of what that number means.



Monk fruit — *luo han guo*, the Buddha fruit. Its sweetness is a triterpene wearing five glucoses, and its gentleness is why it so often rounds the sharper edge of stevia in a blend.

A triterpene wearing five glucoses

Where stevia's sweeteners are diterpene glycosides built on a twenty-carbon skeleton, monk fruit's are **mogrosides** — cucurbitane-type *triterpene* glycosides built on a thirty-carbon skeleton. That ten-carbon difference is not bookkeeping. It is the reason the two sweetnesses feel different in the mouth and behave differently in a hot pan, because a larger, more rigid carbon cage presents a different face to the receptor and holds together under conditions that would unmake a smaller molecule.

The sweet core of every mogroside is **mogrol**, an aglycone — a sugar-free skeleton whose sweetness is only latent until it is dressed. What makes a mogroside sweet, and how sweet, is the number and arrangement of glucose units hung on that core. Each glucose added by the plant climbs a ladder: more glucoses, and the right branched geometry, mean a tighter fit at the sweet receptor and a louder, cleaner signal. The prize molecule is **mogroside V** — five glucoses on the mogrol core, roughly two hundred times as sweet as sucrose by weight, heat- and pH-stable, with a rounded profile and essentially no glycemic load. The minor glycoside **siamenoside I** is sweeter still.

This is the quiet theme of the whole plant section, and monk fruit states it as plainly as any crop we grow: **sweetness lives at the end of a glycosylation pathway, and bitterness lives in its unfinished intermediates.** The *lower* mogrosides — the ones carrying only two or three glucoses — are the problem children: less sweet, and some of them frankly bitter. Good design therefore means one thing above all, driving the pathway to completion so the fruit accumulates the five-glucose sweeteners and starves the half-finished ones.

STEP	ENZYME (CLASS)	PRODUCT	GLUCOSES ON MOGROL	CHARACTER
—	squalene synthase → squalene epoxidase	2,3-oxidosqualene	0	the C_{30} precursor, not yet folded
1	cucurbitadienol synthase (OSC)	cucurbitadienol	0	the triterpene fold
2	P450s (incl. CYP87D18) + epoxide hydrolase	mogrol	0	the aglycone — sweetness latent
3	UGTs (glucosylation)	mogroside IIE / III	2–3	low-sweet, some bitter
4	UGTs	mogroside IV	4	sweet
5	branching UGT (UGT94-type)	**mogroside V**	5	~200× sucrose, clean
5′	UGT	siamenoside I	branched	sweeter than V, minor



Monk fruit in cross-section. The sweet mogrosides concentrate in the dried flesh of the ripe gourd; the pale network of the cut fruit is where a thousand years of slow curing taught the sweetness to build.

From acetyl-CoA to cucurbitadienol, step by step

Follow the carbon and the whole thing becomes designable. The mogroside skeleton is assembled in the cytosol by the **mevalonate pathway**, the same route that in other tissues makes sterols and the backbone of rubber. Two molecules of acetyl-CoA condense to acetoacetyl-CoA; a third is added to make 3-hydroxy-3-methylglutaryl-CoA (HMG-CoA); HMG-CoA reductase — the committed, rate-setting step — reduces it to mevalonate. Mevalonate is phosphorylated twice and decarboxylated to isopentenyl diphosphate (IPP), the five-carbon building block, which equilibrates with its isomer dimethylallyl diphosphate (DMAPP). Head-to-tail additions build the ten-carbon geranyl diphosphate, then the fifteen-carbon **farnesyl diphosphate** (FPP). Squalene synthase joins two FPP molecules tail-to-tail into **squalene**, a thirty-carbon chain, and squalene epoxidase installs a single oxygen to give **2,3-oxidosqualene** — the folded-up precursor of nearly every triterpene a plant makes.

Here the monk-fruit-specific chemistry begins. An oxidosqualene cyclase called **cucurbitadienol synthase** folds that long epoxide, in one committed cyclization, into the four-ring cucurbitane cage of **cucurbitadienol**. A series of cytochrome-P450 oxidations — **CYP87D18** prominent among them — hang hydroxyl groups onto the cage at specific carbons, and an **epoxide-hydrolase** step opens a terminal epoxide to a diol. Those installed oxygens matter twice over: they finish the sweet aglycone **mogrol**, and they become the exact handles onto which the plant will later attach glucose. Then the decorating begins. A cascade of **UDP-glucosyltransferases (UGTs)** adds glucose units one at a time, each raising the molecule's sweetness, until a **branching UGT of the UGT94 family** does the decisive thing: it adds a glucose onto a glucose already in place, building the branched glucosyl chain that distinguishes mogroside V — and, with a small shift in that final linkage, the even-sweeter siamenoside I.

Every one of those waypoints — the cyclase, the P450s, the epoxide hydrolase, the finishing UGTs and the branching UGT — is now a named enzyme with a known sequence, because the *Siraitia* genome has been read in detail. That inventory is what separates a prized folk sweetness from a designable one. It lets us apply the

same selectivity logic we use everywhere: balance the P450 oxidations that make mogrol against the UGT glucosylations that finish it, push flux past the lower, sometimes-bitter mogrosides toward mogroside V and siamenoside I, and leave the intermediates with nowhere to accumulate.

How a mogroside speaks to TAS1R2·TAS1R3

A sweetness is a number only because it is first a binding event. The sweetness axis of our index, **IFI-S**, reads a single receptor: the heterodimer **TAS1R2·TAS1R3**, a class C G-protein-coupled receptor on the sweet-sensing taste cells of the tongue. Each half of the dimer is a large protein with three parts — a bilobed extracellular **Venus flytrap domain (VFT)**, a small **cysteine-rich domain (CRD)** that acts as a hinge, and a **seven-transmembrane domain** that talks to the cell's interior. Small sugars like sucrose nestle into the cleft of the TAS1R2 flytrap and pull its two lobes shut; that closing is relayed down through the cysteine-rich hinge to the transmembrane bundle, and the receptor switches on.

Mogroside V cannot work quite that way, because it is enormous — a rigid thirty-carbon cage carrying a branched spray of five glucoses, far too large to tuck into a single flytrap cleft the way a sugar does. The evidence favors a **multipoint model**: the mogroside drapes across more than one surface of the dimer, its glucose arms making several simultaneous contacts, so that even a very low concentration of molecules holds the receptor in its active, closed conformation. This is the beautiful place where the enzymology and the pharmacology meet. The **branched** glucose chain built by that UGT94-family enzyme is not an incidental flourish; it is the precise three-dimensional array of hydroxyls that makes the best multipoint fit. The branching enzyme is, in effect, carving the key to fit the receptor's lock — and the unfinished mogrosides, with fewer glucoses and no branch, present a worse key, bind more weakly, and read fainter. Some of them present a face that the **bitter-receptor family (the TAS2Rs)** reads instead, which is the molecular root of the bitterness that lives in the half-finished intermediates.

Once the switch is thrown, every sweetener converges on the same downstream wire. The activated receptor couples to the taste G-protein **gustducin**; its released subunits turn on phospholipase C-β2; that enzyme makes the second messenger IP_3; IP_3 opens a calcium store inside the cell; the rising calcium gates the cation channel **TRPM5**; the cell depolarizes and releases ATP through **CALHM1** channels onto the afferent nerve fibers that carry the word *sweet* to the brain. Because that cascade is shared, the brain cannot tell a sucrose sweetness from a mogroside sweetness by its wiring — only by its loudness and its timing. We are honest about the limit here: the exact footprint of a mogroside on the dimer is still being resolved in the structural literature, and we treat the multipoint picture as the best current model, not a settled map. What is settled is the consequence — a handful of mogroside molecules can drive the same channel that a mouthful of sugar drives, and that ratio is what the index measures.

Putting a number on it — the index, worked

The Iyengar Flavor Index turns that loudness into a single figure with one formula:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. It is defined by the dose–response curve itself. If θ is the fraction of maximal receptor response at concentration [A], then

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

where n is the Hill coefficient. Set [A] = EC_{50} and the expression collapses to θ = ½ for any n — that is what half-activation means. A more potent sweetener reaches half-activation at a *lower* concentration, so it has a smaller EC_{50}, and by the formula a larger IFI. The axis is anchored so that the reference sweetness, sucrose, scores exactly **1,000** by convention — a stated definition, the same way the Scoville scale is anchored to dilution-to-threshold. Substitute EC_{50}(candidate) = EC_{50}(reference) and IFI = 1,000 × 1 = 1,000.

Now the worked example. Mogroside V is about **200× as sweet as sucrose by weight** — meaning one gram of it matches the sweetness of roughly two hundred grams of sucrose. On a weight basis, then, its EC_{50} is 1/200 of sucrose's, and

$$\text{IFI-S}(\text{mogroside V, by weight}) = 1{,}000 \times 200 = \mathbf{200{,}000.}$$

But a receptor does not count grams; it counts molecules, and the two sweeteners are wildly different in size. Sucrose ($C_{12}H_{22}O_{11}$) has a molar mass of 342.30 g/mol; mogroside V ($C_{60}H_{102}O_{29}$) weighs 1,287.45 g/mol — nearly 3.8 times heavier. Converting the weight potency to a **molar** potency means multiplying by that mass ratio:

$$200 \times (1{,}287.45 \div 342.30) = 200 \times 3.761 = \mathbf{\sim 752\times\ by\ moles,}$$

and therefore IFI-S(mogroside V, by moles) = 1,000 × 752 ≈ **752,000.** The two numbers are both true and answer different questions: the formulator, dosing a blend by weight, lives in the 200,000 world; the receptor, counting molecules settling into its cleft, lives in the 752,000 world. We carry both because conflating them is how a spec sheet lies.

QUANTITY	SUCROSE (REFERENCE)	MOGROSIDE V	HOW WE GET IT
Molecular formula	$C_{12}H_{22}O_{11}$	$C_{60}H_{102}O_{29}$	composition
Molar mass (g/mol)	342.30	1,287.45	sum of atomic masses
Relative sweetness, by weight	1×	~200×	isosweet matching
Relative potency, by moles	1×	~752×	200 × (1,287.45 ÷ 342.30)
EC_{50} vs. sucrose	1×	~1/752 (molar)	inverse of potency
Raw IFI-S, weight basis	1,000	200,000	1,000 × weight potency
Raw IFI-S, molar basis	1,000	~752,000	1,000 × molar potency
Scorecard IFI-S (cap 5,000)	1,000	~4,800 (pinned)	compressed near the ceiling

The radar scorecard normalizes every axis to a maximum of 5,000, which is why a raw value in the hundreds of thousands cannot be drawn to scale. The scale reads linearly through the ordinary range — so sucrose sits honestly at its anchor of 1,000 — and compresses the extremes toward the ceiling, so a two-hundred-fold sweetener pins near the top. The reference and the fruit, side by side, make the point without a word:

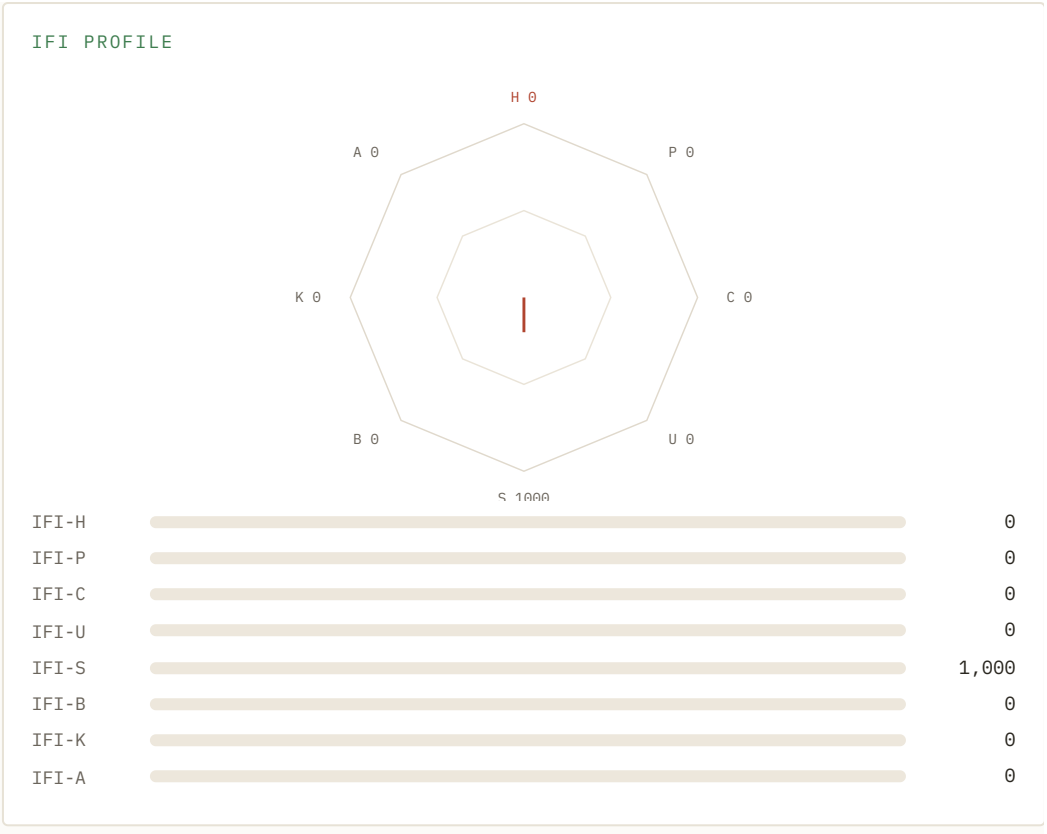

That is sucrose: pure sweetness, parked at the anchor, silent on every other axis. And this is an engineered, mogroside-V-rich monk fruit:

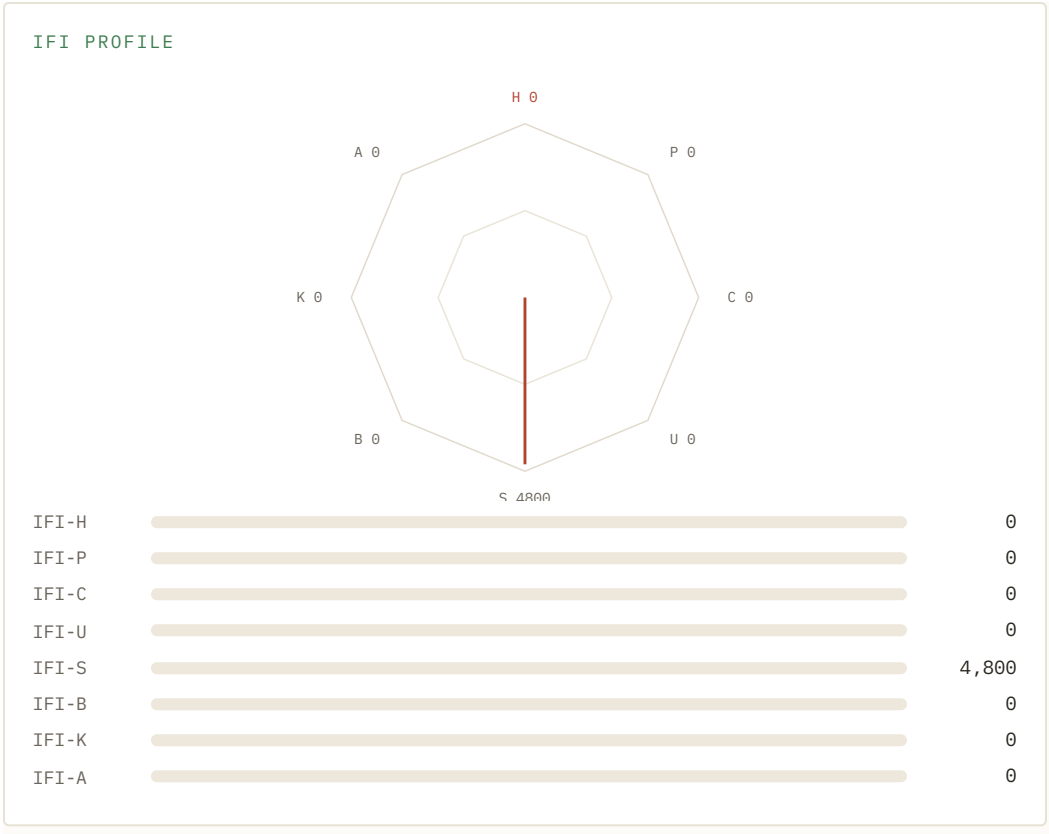

One axis, and only one — pinned near the ceiling, flat everywhere else. This is what a true specialist reads like. Monk fruit does one thing, and the index shows it doing that one thing louder than almost anything a tongue has met.

Stability, and the honest end of what we say

Mogroside V's robustness is no small virtue, and it has a structural reason. The sweetness of many high-potency molecules is carried on fragile architecture that heat or acid will rearrange; the mogrosides carry theirs on a compact, saturated triterpene cage whose carbon–carbon framework simply does not care about the inside of an oven, and whose glycosidic bonds hold through the low pH of a soda far better than one might expect. The practical result is that monk fruit is one of the few high-potency sweetnesses that survives a hot process essentially unchanged — it holds its sweetness through a bake and through a long simmer, which is why it shows up in blends meant for cooking and not only for cold use. A sweetener you can put in the pan at the start is a different kind of tool from one you must stir in at the end.

Like several of the molecules in this book, the mogrosides pass through the body largely without being metabolized for energy, which is a compositional fact about the molecule and the honest end of what we will say about it. The receptor is thrown; the ATP is released; the brain hears *sweet* — and the body draws almost no fuel from the few milligrams that did it. We state that as chemistry and stop there. We make no claim about what monk fruit does for anyone's health, because that is not ours to claim and not what the number measures. These are computational designs; Scoville Splice's monk fruit constructs have not been experimentally validated in a laboratory. The pleasure is the point, and the pleasure is at the tongue. Our family's reverence for feeding the body well is exactly why we hold that line: a sweetness offered honestly, with its real properties named and no halo painted around them, is a better gift than a promise we cannot keep.

Why "round" is a time, not an eighth axis

People reach for *round* and *full* to describe monk fruit, and those words tempt a designer to invent an axis for them — a little kokumi here, a touch of cooling there. We resist that, because it would be false at the receptor. Monk fruit's roundness is not a second taste; it is the **shape of one taste over time.** Every sweetness has a time–intensity curve: how fast it rises as the molecule reaches the receptor, how high and how broad its plateau, how cleanly it falls as the tongue clears it. Sucrose has a characteristic arc — a prompt rise, a broad middle, a clean fall. Stevia's steviol glycosides tend to spike early and then linger, and that lingering tail is where the licorice note lives. Monk fruit's mogrosides do the opposite: they rise more slowly, hold a full, broad plateau through the middle of the bite, and leave without a tail.

That slow, broad plateau *is* the roundness. The faint cooling impression at the finish is the same phenomenon read from the other end — a potent sweetener dosed in milligrams clears the tongue quickly and leaves a clean, almost cool absence where a heavy dose of sugar would leave a syrupy film. None of this is a TRPM8 signal or a CaSR signal, so none of it belongs on the cooling or kokumi axes; put it there and the profile would lie about which receptors are actually lit. The discipline matters. An index is only worth citing if every axis means exactly one thing, and we would rather describe monk fruit's roundness precisely in a sentence about time than smear it dishonestly across a radar. The scorecard says *sweetness, and only sweetness.* The prose says *and its sweetness is shaped like this.* Both are true, and kept separate on purpose.

A seven-percent batch and a saturating receptor

A sweetness you cannot reproduce is a liability, so consistency is a number we track as closely as potency. Across our constructs we hold **batch-to-batch variation to a CV of ≤7%** — the coefficient of variation, CV = σ/μ, the standard deviation of a batch's active content as a fraction of its mean. Seven percent sounds like a lot to carry into a food, but the saturating shape of the receptor quietly shrinks it, and it is worth doing the arithmetic to see why.

Take the simple one-site dose–response (Hill coefficient n = 1), $\theta = [A] \div (EC_{50} + [A])$. Ask how much the *response* wobbles when the *concentration* wobbles — the elasticity of θ with respect to [A]:

$$(\Delta\theta/\theta) \div (\Delta[A]/[A]) = EC_{50} \div (EC_{50} + [A]).$$

At the half-activation point, where $[A] = EC_{50}$, that elasticity equals $EC_{50} \div (2 \cdot EC_{50}) = \tfrac{1}{2}$. In words: at EC_{50}, a given fractional error in concentration produces only *half* that fractional error in perceived intensity. A 7% batch CV in mogroside V content therefore lands as about **3.5%** variation in sweetness at the tongue — below the threshold where most tasters notice a difference at all. Push the dose higher, well above EC_{50} toward saturation, and the elasticity falls further toward zero: the error all but vanishes, because the receptor is nearly full and more molecules change little. Pull the dose low, well below EC_{50} into the linear regime, and the elasticity climbs toward one, passing the full 7% straight through to the mouth.

So there is an honest optimum, and an honest trade-off. Dosing a potent sweetener near or slightly above its EC_{50} buys stability — the saturating curve is forgiving exactly there — but costs headroom, because a receptor already two-thirds full cannot be modulated much further. Dosing lower keeps the headroom but surrenders the error-smoothing. We tune each formulation into the band where the sweetness is both reproducible and still adjustable. Precision, in our family's reading of it, is a form of respect for whoever eats the result: it means the thousandth jar tastes like the first.

Where the pathway stalls, and the edges we watch

A designable pathway is also a pathway with failure modes, and naming them is part of the honesty. The most common way the chemistry disappoints is not catastrophe but a *stall* at **mogroside IV** — four glucoses, genuinely sweet, but not the peak. If the branching UGT of the UGT94 family is underpowered relative to the oxidations and glucosylations feeding it, the fruit piles up the four-glucose form it can make easily instead of the branched five-glucose form it cannot, and the extract tastes good rather than excellent. That branch step is the true kinetic bottleneck of the whole route, which is why we spend more design effort there than anywhere else; a monk fruit is won or lost on one enzyme's ability to put a glucose onto a glucose.

The second failure mode is a balance problem. The P450 oxidations that build mogrol and the UGT glucosylations that dress it must stay matched. Tip the balance toward too little oxidation and the plant starves for mogrol, so there is nothing to sweeten; tip it toward too much, and bare, under-dressed intermediates accumulate faster than the UGTs can finish them — and those exposed intermediates are exactly the ones the bitter receptors read. Either error degrades the profile, and they degrade it in opposite directions, so the work is a tuning between two cliffs rather than a push toward one. Even the siamenoside-I enrichment, which is only a small shift in the final glycosylation, has an edge: pressed too hard it can cost total sweet-mogroside yield, so it is a dial we turn gently, not a maximum we chase.

Then there is the plant itself. *Siraitia* depends on hand-pollination, and a missed window means poor fruit set that season; the sweetness builds during a slow cure, and a rushed drying leaves a greener, less-finished fruit; and because the vine is a perennial, a wrong turn costs years rather than a single planting. We are honest about the deepest limit, too: we can bias the flux, but we cannot perfectly forbid the lower mogrosides, so a trace of them always remains. Part of the craft is keeping that trace below the bitter-receptor threshold — managing it, not pretending it is zero.

A program kept deliberately small

Our monk-fruit program is deliberately small — **80 constructs** — and the smallness is a decision, not a shortfall. *Siraitia* is a **hard** host to engineer and a slow one to grow: a perennial vine, traditionally hand-pollinated in the misty hills of Guangxi, with a long road from seed to a fruiting harvest and gourds that must be cured slowly after picking. It is a costlier crop to transform and transform again than a leafy annual. And its chemistry needs less rescuing than most — a good

monk fruit is already fairly clean, with none of the licorice-and-bitter burden some sweeteners carry — so the work is less about remaking the plant than about sharpening what it already does supremely well.

The 80 constructs concentrate on two things. First, **flux to mogroside V and siamenoside I**: balancing the P450 oxidations that make mogrol against the UGT glucosylations that finish it, and strengthening the branching UGT, so the fruit piles up the five-glucose sweeteners rather than stalling at the lower, bitter intermediates — the same *drive-the-pathway-to-completion* logic that governs every glycoside crop we grow. Second, and more interesting for a formulator, **tuning monk fruit to be the rounding partner** in a blend, enriching the slow, broad-plateau mogrosides that fill a sweetness's middle. Set against its sibling sweetness crops, the restraint of the program reads clearly:

CROP	CONSTRUCTS	AXES	DIFFICULTY	SWEET PRINCIPLE	TIME-INTENSITY
Stevia	910	S / F / M	Moderate	steviol glycosides (diterpene) — Reb M	fast, bright front
Sugarcane	200	S	Frontier	sucrose itself	the full natural arc
Monk Fruit	**80**	**S**	**Hard**	mogrosides (triterpene) — mogroside V	slow, round body



The controlled plot where the vine is trialed. *Siraitia* is a slow perennial, trained and hand-pollinated; we keep the program small and let the plant do what it already does supremely well.

Eighty constructs against stevia's nine hundred and ten is not timidity; it is a crop that arrives most of the way finished, asking for a scalpel rather than a wholesale rebuild.

The second voice

Monk fruit does not need to win alone. It needs to be the right second voice, and the mathematics of the blend is the practical payoff — worth stating as a rule. A sweetness blend sums its components at the tongue: the perceived intensity over a bite is roughly the sum of each sweetener's own time–intensity curve, and below the point where the shared receptor saturates, two half-doses of sweetness add to about one full dose. That additivity is the lever. Compose a blend of a fast, bright sweetener and a slow, full one,

and the summed curve can trace an arc close to sucrose's own — the fast sweetness covering the onset, the slow sweetness filling the body, the combination rising and falling like sugar does.

That is precisely the stevia-and-monk-fruit pairing. Stevia's Reb M is quick and sharp at the front; monk fruit's mogroside V is slow and round through the middle; sucrose's natural arc spans both. A Reb-M-rich stevia blended with a mogroside-V-rich monk fruit produces a combined sweetness closer to sucrose's clean arc than either achieves alone — the stevia carrying the top, the monk fruit filling the middle, each covering the other's gap. And because the two are dosed *low* — each kept below the concentration at which its own off-note would surface — the licorice tail never arrives and nothing competes with the clean plateau. There is even reason to think the pairing does a little better than simple addition: steviol glycosides and mogrosides appear to press on the sweet dimer at distinct contact points, so holding the receptor's active shape from two places at once can combine slightly more than additively, more sweetness returned per milligram than either sweetener spent. The dominant effect is still the plain one — two sweetnesses, each kept quiet, summing to a fuller whole — but the receptor seems to reward the collaboration.

It is, in the end, a very old idea rendered in new arithmetic. For a thousand harvests the growers of Guangxi trained this vine on its trellises and dried its gourds in the mist because the fruit was sweet in a way that felt gentle and whole, long before anyone could name the receptor it spoke to. We have only learned to read the number under that gentleness, and to grow a second crop to meet it — two engineered plants, each tuned to be the other's missing half, so that a sweetness with no sugar in it can still rise and fall across the tongue the way our family always wanted sweetness to: fully, cleanly, and with nothing to apologize for.

Chapter 10 · Sugarcane — *sweetness at the source, reconsidered.*

There is an irony at the center of this book that we would rather meet head-on than tiptoe around, and its name is sugarcane. The argument we have been building is that sweetness can be uncoupled from sugar — that the sensation and the molecule are two different things, and that the first need not drag the second behind it everywhere it goes. And yet the plant we work on hardest for sweetness is the single organism on Earth most completely built to make sugar. A mature cane stalk stores sucrose at up to a fifth of its fresh weight: the densest accumulation of one sweet compound anywhere in agriculture, a green column of almost pure stored sunlight. To carry it into a book about sugar-free sweetness looks, at first, like a contradiction we ought to be embarrassed by. We are not embarrassed. Sugarcane is not the exception to the thesis; it is the resolution of it.

To see why, we have to hold two things at once, the way the family has learned to hold them at the bench — the molecule, exactly, and the sensation it provokes, exactly. They are not the same thing, and this entire chapter lives in the gap between them. One is a sugar stored in a stalk. The other is a number a receptor reports. Pry them apart and the whole plant turns from a sugar factory into something larger: the best chassis we have for growing sweetness that need not be sugar at all.

The molecule, and the machine that makes it

To engineer a thing you must first understand it, so begin with the sucrose itself, because its manufacture is one of the quiet marvels of the living world. It is assembled in the leaf from the raw output of photosynthesis. The Calvin cycle fixes carbon into triose phosphates; these condense into fructose-6-phosphate, which is isomerized and, in partnership with UDP-glucose, joined by **sucrose-phosphate synthase (SPS)** into sucrose-6-phosphate. A second enzyme, **sucrose-phosphate phosphatase (SPP)**, strikes off the phosphate — an effectively irreversible step that pulls the whole reaction forward and commits the carbon to sucrose. That free sucrose is loaded into the phloem, shipped down the plant under pressure, and unloaded into the storage parenchyma of the stalk, where its accumulation and turnover are governed by **sucrose synthase (SUS)** running in its cleavage direction and by the family of **invertases** — cell-wall, cytoplasmic, and vacuolar — that split it back into glucose and fructose. Storage is never a static pool; it is a standing balance between synthesis and cleavage, held far over to the synthesis side.

Sugarcane's genius is partitioning. Where most plants spend their fixed carbon on growth and defense and keep only a modest sugar reserve, cane diverts an extraordinary share of its carbon into stalk sucrose and, crucially, keeps it there rather than respiring it away. Two features make the feat possible. The first is photosynthetic: cane runs the **C4 pathway**, a carbon-concentrating mechanism that

packs CO_2 around the fixing enzyme Rubisco and all but abolishes the wasteful photorespiration that bleeds carbon from ordinary C3 crops. Under hot, bright, well-watered conditions a cane field fixes carbon faster, and converts more incident sunlight into sugar, than almost any plant we know. The second is **sink strength** — the thermodynamic and transport pull of the storage tissue that draws sucrose down the plant and locks it in the vacuole against its own concentration gradient. To engineer cane's sweetness is, in the end, to engineer that partitioning: to decide what the sink pulls down and holds. And the same machinery that so efficiently fills a stalk with soluble sucrose is, redirected, machinery for filling a stalk with something else soluble and sweet.



Sugarcane, whole. A green column of stored sunlight — and, reconsidered, the best chassis we have for growing sweetness that need not be sugar.

How sweetness is read — the receptor, step by step

Nothing in that stalk is sweet until a receptor says so. Sweetness is not a property of sucrose; it is an event on the tongue, and the event is specific enough that we can write it down. The sweet taste receptor is a single molecular machine: **TAS1R2 and TAS1R3**, two class-C G-protein-coupled receptors locked together as an obligate heterodimer. Neither half works alone. Each subunit is built in three parts — a large, hinged, bilobed extracellular **Venus-flytrap domain (VFD)**, a short **cysteine-rich domain (CRD)** that threads the signal inward, and a **seven-transmembrane (7TM) domain** that passes it across the membrane to the cell interior.

What makes this receptor extraordinary, and what makes the whole engineering program possible, is that it is not one lock but several. Small sugars — sucrose, glucose, fructose — settle into the cleft of the TAS1R2 flytrap and pull its two lobes shut, like a hand closing around a pebble. That closure is the switch. But other sweet things dock elsewhere: some sweeteners and the inhibitor lactisole bind the 7TM barrel of TAS1R3; others engage the flytrap of TAS1R3; and the **sweet proteins** — brazzein, thaumatin, monellin — are far too large to fit any internal pocket at all. They bind across the *outer* surface, wedging against the cysteine-rich domain and the exterior of the flytrap, and stabilize the closed, active conformation from the outside. One receptor, many doors. It is why a sugar, a leaf glycoside, and a protein a thousand times heavier can all ring the same bell.

However they bind, the downstream cascade is one shared path, and it is worth tracing because it is where "sweet" actually becomes a signal. Agonist binding closes the flytrap; the motion travels through the cysteine-rich domain to the 7TM; the receptor couples to **α-gustducin**, the taste cell's G protein; its released βγ subunits switch on **phospholipase C-β2 (PLCβ2)**, which cleaves membrane lipid into **IP_3**; IP_3 opens its receptor on the endoplasmic reticulum and **Ca^{2+}** floods the cytoplasm; the calcium gates **TRPM5**, a cation channel, and the cell depolarizes; depolarization opens **CALHM1/CALHM3** channels that release **ATP**, which excites the gustatory afferent fibers — chorda tympani at the front of the tongue, glossopharyngeal at the back — and the brain, at last, reads *sweet*. Every step is a place a molecule could be louder or quieter. The receptor, not the recipe, is the unit of measurement.

One consequence of all this specificity matters enormously for method, and we learned it the hard way as every sweet-protein program must: **the sweet receptor is species-specific.** Brazzein, thaumatin, and monellin taste intensely sweet to humans and Old World primates and are nearly tasteless to rodents, because the mouse TAS1R2/TAS1R3 pair differs at exactly the surface residues the proteins wedge against. A mouse will tell you nothing true about a sweet protein. So the platform validates sweetness on human-receptor cell assays and trained human panels, never on an animal model — a constraint, but an honest one, and a reminder that we are engineering a conversation with *our* receptors, not with life in general.

Loudness, not quantity — the index on the sweet axis

Here is where the molecule and the sensation finally come apart on paper. The Iyengar Flavor Index does not measure how much of a compound a food holds. It measures how *loudly* that compound speaks to its receptor — and loudness is the inverse of **EC_{50}**, the concentration at which a stimulus drives its receptor to half of full response. A compound with a low EC_{50} reaches half-activation at a whisper of a dose; it is loud. A compound with a high EC_{50} needs flooding before the receptor answers; it is quiet, however much of it you pour. The index inverts EC_{50} against a fixed reference and scales by a thousand:

$$\text{IFI-S}(c) = 1{,}000 \times EC_{50}(\text{ref}) \div EC_{50}(c)$$

For the sweet axis the natural reference is sucrose itself — the sweetness every human palate is calibrated against. We anchor the axis at the half-activation of sucrose, taking $EC_{50}(\text{ref}) \approx 30$ mM, so that by construction sucrose scores exactly 1,000. It is the zero-point of the sweet world: loud enough that civilizations were built on it, quiet enough that you must use *grams*. Each axis of the IFI is anchored this way — a stated convention, like how the Scoville scale is defined by dilution-to-threshold, not a value derived from a single literature source.

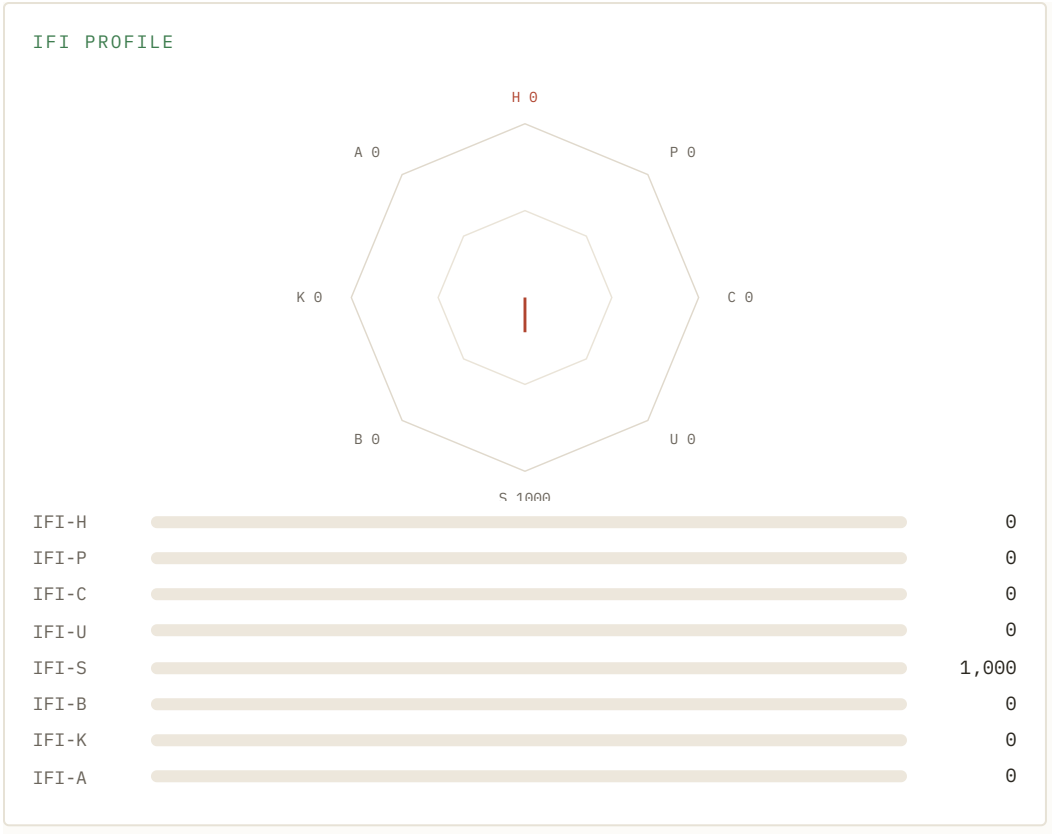

That profile is sucrose, and it is beautiful in its silence: a pure tone on the sweet axis, dead quiet on every other receptor. Sucrose does not touch TRPV1 or mGluR1 or the bitter channels; it says one thing, clearly, and nothing else. It is also the reason sugar is so easy to overuse — a single uncomplicated loudness, cheap to make, that fits everywhere.

Now the arithmetic that reorganizes the whole field. "Sweeter than sucrose" is quoted in the trade by *weight*, but the receptor counts *molecules*, so to put any sweetener on the index we convert. If a compound is **N** times as sweet as sucrose by weight, then to be iso-sweet you use 1/N of the mass — and its molar potency, the number that sets EC_{50}, is **N × (MW_c ÷ MW_sucrose)**, because the heavier the molecule, the fewer of them that mass contains. Sucrose weighs 342 g/mol; stevioside, 805; rebaudioside A, 967; mogroside V, 1,287; brazzein, a small protein, about 6,500. Work the conversion and the loudness of these quiet-by-nobody molecules becomes visible:

SWEETENER	×SUCROSE (WT)	MW (G/MOL)	MOLAR POTENCY	EC_{50} (FROM 30 MM)	DOSE PER 355 ML*
Sucrose (reference)	1	342	1	30 mM	35.5 g
Stevioside	200	805	470	64 µM	178 mg
Rebaudioside A	250	967	707	42 µM	142 mg
Mogroside V (monk fruit)	250	1,287	940	32 µM	142 mg
Brazzein	1,000	6,500	19,000	1.6 µM	35.5 mg

*mass needed to match a 10% w/v sucrose solution (≈35.5 g in a 355 mL serving).

Take brazzein, worked all the way through. Its molar potency is 1,000 × (6,500 ÷ 342) ≈ 19,000; it reaches half-activation near EC_{50} = 30,000 µM ÷ 19,000 ≈ 1.6 µM. Dropped into the index formula, its raw loudness is 1,000 × (30 mM ÷ 1.6 µM) ≈ 19,000,000 — a number so far past the axis that it means something different in kind, not degree. This is the first honest subtlety of the sweet axis: the intrinsic loudness of a high-potency sweetener is astronomical, orders of magnitude beyond the 5,000 ceiling the scorecards are normalized to. **What a scorecard reports is never the compound's intrinsic potency; it is the sweetness a formulated food actually**

delivers — the loudness you *dial in*, not the loudness the molecule is capable of. Brazzein's intrinsic 19-million is why you can tune a product to a delivered IFI-S of 1,000, matched to sugar, using a dose you would struggle to see.

Dose–response, and the ceiling nobody mentions

EC_{50} is a single point on a curve, and the shape of the whole curve carries the rest of the truth. Receptor activation follows the Hill form:

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

At [A] = EC_{50} the response is R_max/2, which is simply the definition of EC_{50}. The exponent **n**, the Hill coefficient, measures cooperativity — how sharply the receptor flips from quiet to loud as dose climbs; sugars sit near n = 1, while some sweeteners that engage more than one site on the dimer show mild positive cooperativity, a steeper turn-on. Take brazzein at n ≈ 1 and EC_{50} = 1.6 µM: at five times its EC_{50} (8 µM) it reaches 8 ÷ (1.6 + 8) = 83% of its maximum; at ten times (16 µM), 91%. It climbs to nearly full voice within a few micromolar and then flattens.

And *where* it flattens is the subtlety the marketing never mentions. Potency and **efficacy** are different axes of the same curve. EC_{50} sets how little you need to reach half of a molecule's maximum; **R_max** sets how high that maximum is. A sweetener can be far more potent than sucrose — a far lower EC_{50} — and still saturate at a sweetness *below* what sugar can reach, because its R_max is lower. You feel this as a ceiling: push the dose and the sweetness simply stops rising, sometimes while an off-note keeps climbing. It is why the single "×sucrose" multiplier that the trade quotes is a low-dose approximation that quietly erodes as you ask for more intensity, and why the loudest desserts are still, honestly, blends. The index, anchored at half-activation, captures the *onset* of loudness cleanly; the ceiling is a second number, and a candid sweetness program keeps both in view.

Cane as a chassis — the sugar factory, re-instructed

Return now to the stalk, carrying everything the receptor taught us. A plant that devotes its entire physiology to building and holding an enormous mass of soluble sweet storage is, viewed correctly, the ideal host in which to grow *other* soluble sweet molecules — the high-potency, low-dose ones. Express a steviol-glycoside pathway, or a sweet protein, in the same crop that so efficiently loads and locks soluble compounds into its vacuoles, and you have turned a sugar factory into a sweetness factory whose output need not be sugar. This is **molecular farming** at agronomic scale, and cane has a real track record as a host engineered to accumulate recombinant proteins and alternative sugars in its biomass. Our sugarcane program — **200 constructs** — is in large part exactly this: cane instructed not to make *more* sucrose, but to make a sweetness of a different, louder, lighter kind.



The construct bench. Behind a 200-construct cane program is a library of sequences, each a different instruction given to the same green machine.

The chassis argument has a hard industrial edge worth naming. Breeders already distinguish classic **sugarcane**, bred to pack its stalk with sucrose, from **energycane** and fiber canes bred for sheer tonnage — tougher, bulkier, grown for fuel and fiber rather than refined sugar. That same standing biomass is a reservoir. And the leverage is easy to size, because the only inputs are the canonical fifth-of-fresh-weight sucrose figure and the potency multiples we already converted. One tonne of fresh cane at 20% sucrose carries 200 kg of sucrose — call that its sweetening output, 200 kg of sucrose-equivalent. Now ask: at what expression level would a molecular-farmed sweetener in that same tonne *match* that output? The break-even is **X% × N = 20%**, so X = 20 ÷ N:

STORED PRODUCT	WEIGHT POTENCY	BREAK-EVEN EXPRESSION (% FW)	SUCROSE-EQUIV / TONNE AT 0.2% FW
Sucrose (native)	1×	20% (it is the 20%)	200 kg
Stevioside	200×	0.10%	400 kg
Rebaudioside A	250×	0.08%	500 kg
Brazzein	1,000×	0.02%	2,000 kg

Read the brazzein row slowly. A cane designed to express brazzein at just **0.02%** of its fresh weight is predicted to match the sweetening power of a cane stuffed to 20% with sugar — a thousandfold less material for the same delivered sweetness. Push expression to a still-modest **0.2%** of fresh weight, a plausible molecular-farming target, and a single tonne of cane carries **2,000 kg of sucrose-equivalent** sweetness in our model, ten times what its sugar ever could, at zero glycemic load. The same hectares, the same sunlight, the same C4 engine filling the same vacuoles — re-instructed to store loudness instead of mass. None of it asks for new land or a new industry. It asks a planetary-scale sweetness machine to make a second kind of sweetness in the same stalk.



Inside the stalk. The storage parenchyma that so efficiently holds sucrose is, re-instructed, a vault for a low-dose sweetener — the same sink, a different cargo.

Reformulation, in numbers

What that leverage buys, at the table, is reformulation — and reformulation is best understood not as a slogan but as a worked example. Sweeten a 355 mL beverage to the sweetness of an ordinary soda, roughly a 10% w/v sucrose solution. That is **35.5 g** of sugar — or, at their weight potencies, **35.5 mg** of brazzein, or **178 mg** of stevioside. The caloric arithmetic that follows is the entire public-health argument in two lines: 35.5 g of sucrose at ≈3.87 kcal/g is about **137 kcal** per serving, with a glycemic index near 65; 35.5 mg of brazzein, though chemically a protein at the same ≈4 kcal/g, contributes **~0.14 kcal** and essentially no glycemic response. Identical delivered loudness. A hundred-and-forty calories, gone to nearly nothing. A thousandfold drop in the mass of the sweet thing, and the body carries what remains along a completely different path.

We state the limit as plainly as the promise, because a sweetness that pretends sugar did only one job is a sweetness that fails in the factory. Sucrose is not only sweet. It is **bulk** — it gives a cookie its body, a syrup its viscosity, a scoop its structure. It **browns**, driving the Maillard and caramelization reactions that make a crust and a crème brûlée. It is a **humectant** that keeps crumb soft, a **freezing-point depressant** that keeps ice cream scoopable, and a **preservative** that binds water away from microbes in a jam. Replace 35.5 g of it with 35.5 mg of protein and you have matched exactly one of those functions and abandoned the other five. Honest reformulation therefore pairs the high-potency sweetener with **bulking agents** — soluble fibers, rare sugars such as allulose, or polyols — to rebuild body, texture, and water activity, and treats sweetness as one restored attribute among several, never the whole job. The math that makes the sweetness trivial is the easy part; rebuilding the plate around it is the craft.

What the mill does to a protein — methods, edge cases, failure modes

The prettiest part of the chassis argument is that the infrastructure to grow, cut, crush, clarify, and extract cane already spans the tropics — and the prettiest part is also where the honest trouble lives. A sugar mill is not a gentle place. Juice is clarified with lime and

heated near boiling, then driven through multiple-effect evaporators. Sucrose, a small stable disaccharide, shrugs this off. A **protein** may not: most sweet proteins would denature in the clarifier and arrive at the pan as tasteless, tangled polymer. This is a real failure mode, and it sorts the field cleanly. **Steviol glycosides and mogrosides are small, heat-stable terpenoid glycosides** — they ride the existing mill with modest change, which is much of why they reach market through conventional processing. **Sweet proteins demand a gentler unit operation**, a cold or low-thermal extraction grafted onto the line, not the raw thermal sledgehammer that refining sucrose allows.

Which is precisely why, of all the sweet proteins, **brazzein** is the one we build the cane program around. Brazzein is a 54-residue protein cross-stitched by disulfide bonds into a fold that is remarkably indifferent to abuse — stable to roughly 98 °C and across a broad pH band — so it survives thermal and acid processing that would unspool a more delicate protein like monellin. Its robustness is not a footnote; it is the whole reason a protein belongs in a crop meant to be crushed and cooked. Even so, the molecular-farmer's problems do not end at thermostability. Expression must survive the osmotic extreme of a sucrose-saturated vacuole; the delivered taste must be clean, and here glycosides carry their own honest flaw — **stevioside** in particular trails a licorice-bitter aftertaste because, past a dose, it also lights the bitter receptors **TAS2R4 and TAS2R14**, which is why rebaudioside A and monk-fruit mogroside V, cleaner on those channels, are often preferred and why low-dose **blends** that mask one molecule's off-note with another's onset are a formulation strategy, not a compromise. Sweet proteins bring a different quirk — a slower onset and a lingering finish, a different time-intensity shape than sugar's quick bright peak. And any novel food protein carries a genuine allergenicity question that is assessed as an attribute, never waved away. We say all of this out loud because the leverage is real only if the limits are too.

Precision — the CV the formulator actually buys

A formulator does not buy a mean; they buy a mean *and its spread*, because a sweetness that drifts batch to batch is a product that cannot be trusted on a label. The platform holds batch-to-batch variation to a coefficient of variation **≤7%** (CV = standard deviation ÷ mean), and the number earns its keep here. Suppose a construct is designed to express brazzein at a mean of 0.20% of fresh weight. At CV = 7%, the standard deviation is 0.07 × 0.20% = **0.014%**, so the stalk-to-stalk content is designed to run 0.200 ± 0.014% FW. Because delivered sweetness is, to first order, proportional to that content, the delivered IFI-S inherits the same ±7% — the whole difference between a cola that tastes like itself every can and one that does not.

Errors do not stay put; they stack. If expression carries a 7% CV and extraction-and-assay adds an independent 4%, the two combine in quadrature, $\sqrt{(7^2 + 4^2)} = \sqrt{65} \approx$ **8.1%** — a small reminder that every added step erodes the tolerance, and that ≤7% at the finish means tighter than that at each stage. The discipline reads clearly against an unmanaged baseline: where natural capsaicin accumulation in wild-type peppers swings at a 42% CV, a designed and controlled

construct is specified to hold to ≤7% — a sixfold tightening that is, for a buyer, the actual product. Precision is not a virtue the plant was born with. It is engineered in, and it is what a formulator is really paying for.

The sweetness axis never stands alone, either, and the flagship makes the point. Even **Phoenix** (SS-0096) — the full-spectrum cultivar we designed gene by gene to reach across all eight receptors — carries a real sweetness in its profile, IFI-S 1,100, a hair above the sucrose reference, woven in among the heat and the rest:

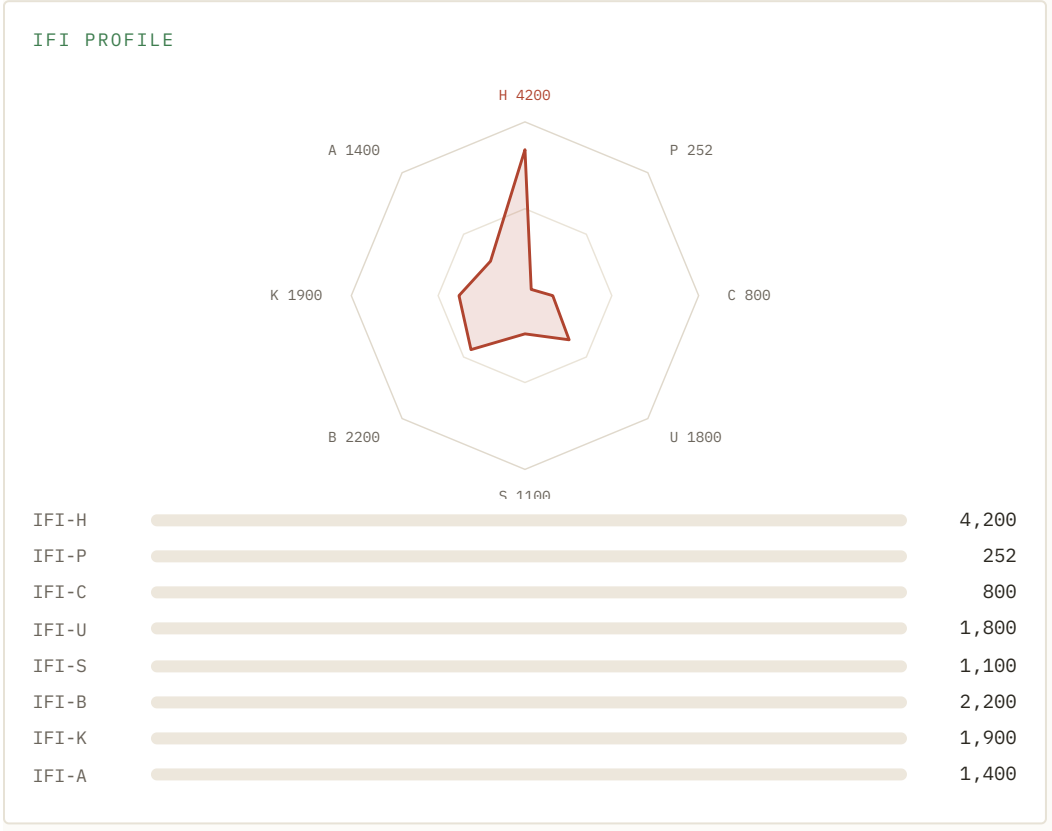

Note that each IFI axis is anchored to its own stated reference constant; scores on different axes reflect loudness relative to that axis's convention and are not directly comparable across axes.

Sweetness is one voice in an eight-part chord, and the sugarcane program is simply that voice, grown to agronomic scale in the plant best built to hold it. The roster of hosts we tune primarily for that voice is small and deliberate:

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

Sugarcane sits at the frontier end of that list — difficult precisely because it is so committed to doing one thing — and that commitment, turned, is its gift.

Choice at the source

This is where the ethos that runs under this whole library stops being quiet and speaks plainly. We are not against sucrose. It is a marvel of partitioning and one of the most useful molecules agriculture ever learned to grow, and much of the pleasure of a life well-fed has passed through it. What we are against is the *absence of a choice* — a world in which the only industrial sweetness is bulk sugar, poured into everything because it was the only thing a field knew how to send. A sugarcane designed to yield either its traditional sucrose *or*

a high-potency, low-load sweetness, from the same fields and the same supply chains, restores that choice. It lets a product be sweetened to the same delivered number by a route the body carries differently, without tearing up the agriculture that already feeds the world.

We are not trying to abolish the cane fields. We are trying to widen what they can send to the table. A grower could harvest sucrose, or a trace of loudness, from the same hectares and the same mills as the market asks — the land unchanged, the sunlight unchanged, the choice restored to the people who grow and the people who eat. That is the pragmatic heart of reformulation at scale: not replacing the world's sweetener infrastructure, but giving it a second thing to grow. The cane stays. What it makes becomes a decision — and a decision, made with care for the body it feeds, is the whole of what we have ever wanted this work to be.

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

Chapter 11 · The sweet proteins — *when sweetness is a protein, not a sugar.*

The loudest sweetness in nature is not a sugar, not a glycoside, not any small molecule at all. It is a handful of small **proteins**, made by a few plants of the West African and Southeast Asian forests, that fool the sweet receptor more completely than anything chemistry has designed. Thaumatin, from the katemfe fruit, is thousands of times sweeter than sugar by weight. Brazzein, from the j'obli berry, is comparably potent and almost indestructible. Monellin, from the serendipity berry, is sweeter still. To the person who first tastes the diluted protein, the experience is uncanny — an intense, clean sweetness from a solution that is, by mass, almost pure water. These molecules are the frontier of the sweet axis, and because they are proteins, they are the one sweetness a plant can be *directly engineered to manufacture* from a gene we write.

We find ourselves returning to them again and again at our own table, because they resolve a contradiction we care about. Every other way to sweeten a food either adds sugar the body must then contend with, or borrows a small molecule — a steviol glycoside, a mogroside — that the plant assembles through a long enzymatic pathway we can only coax, never author outright. A sweet protein is different in kind. It is a sequence. Hand a plant the gene and the folding machinery, and the sweetness and the food become the same living object. That is why this chapter is not really about sugar at all. It is about the shortest possible distance between an idea for a flavor and a fruit that carries it.

A wedge under the door

To understand why a protein can out-shout a sugar, look at the receptor it speaks to. Sweetness on the tongue is read by a single heterodimer, **TAS1R2 paired with TAS1R3** — two proteins clasped together, each carrying a large extracellular lobe shaped, in the old and exact metaphor, like a Venus flytrap: two halves hinged at the back, open to the mouth. A sugar molecule drifts into the cleft between the halves, and the trap closes around it. That closure is the signal; the shut, active conformation is relayed through the membrane and the nerve reports *sweet*. A small sweetener works by being caught and held in the cleft, and because binding is reversible, it must be caught and re-caught thousands of times to sustain the taste.

A sweet protein is far too large to drop into that cleft. Instead it docks against the *outer* face of the receptor's flytrap and cysteine-rich domains, holding the active, closed conformation shut from the outside — a wedge jammed under a door that was already swinging closed. This is the whole secret of their potency. A wedge that holds

the trap shut does not have to be continually recaptured the way a small ligand does; one protein, well seated, keeps the channel reporting sweetness far longer per molecule than any sugar can. It is also why their *structures* matter so completely to their function. Each sweet protein is a compact, disulfide-stitched fold, and the specific loops and charged residues that press against the receptor are what set both the intensity of the sweetness and its temporal shape — the speed of its onset, the length of its linger. Change those surfaces and you change the taste; the molecule is a key cut entirely on its outer edges.

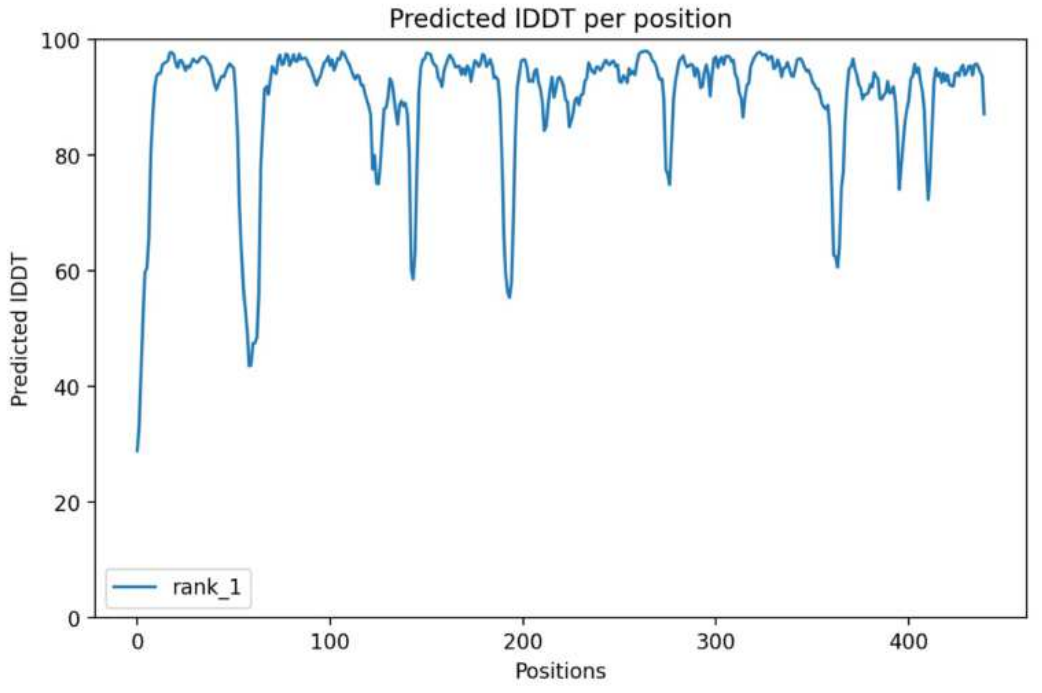

Pun1/AT3 modelled and coloured by the model's per-residue confidence — the same confidence-coloured pipeline we point at a sweet protein. Blue is high-confidence scaffold; the warmer, lower-confidence stretches are the flexible, receptor-facing loops where sweetness is decided, and where we design.

What "half-activation" means, and why the index starts there

The sweet axis of the Iyengar Flavor Index — IFI-S — is anchored, like every axis, to a receptor, not to a concentration. The number answers one question: how *loudly* does this food speak to TAS1R2·TAS1R3? Loudness is defined through the half-maximal activating concentration, EC_{50} — the concentration at which a stimulus drives its receptor to half of full response. The dose–response of a single receptor population follows the Hill relation,

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

where R is the fraction of maximum response, [A] the stimulus concentration, and n the Hill slope. Set [A] equal to EC_{50} and the equation collapses to R = ½ exactly, for any n — the EC_{50} is the one point on the curve that does not care about cooperativity. It is also the steepest, most information-rich part of the response: with a simple slope near n = 1, a stimulus at three times its EC_{50} reaches R = 3⁄4, and at one-third of EC_{50} falls to R = 1⁄4. Half-activation sits at the fulcrum of that curve, which is why we anchor the index there and nowhere else.

The index then inverts and scales:

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

A candidate exactly as loud as the reference scores 1,000. One that reaches half-activation at a quarter of the reference concentration — four times louder — scores 1,000 × 4 = 4,000. The arithmetic is deliberately plain, because a form of respect for the reader is to show it. Worked the other direction: the full-spectrum fruit we hold up as proof of the sweet axis, **Phoenix (SS-0096)**, registers **IFI-S = 1,100**. Rearranged, that says EC_{50}(candidate) = EC_{50}(reference) × 1,000⁄1,100 = 0.909 × EC_{50}(reference) — a sweetness just **1.1× as loud** as the native reference. Not a bludgeon; a clean, balanced sweetness, exactly the restraint you want in something you intend to eat every day. Scorecards cap each axis at a normalization maximum of 5,000, so even the index's loudest plausible sweetness stays on one readable page beside heat, umami, and the rest.

One small, honest caution lives inside the formula. The index measures *loudness at the receptor*; the folk figure "a thousand times sweeter than sugar" measures *potency per unit mass*. They are different questions with different units, and conflating them is the most common error made about sweet proteins. The conversion between them is worth doing carefully, because the gap is enormous and entirely real.

Weight, moles, and the honest size of "sweeter than sugar"

When the literature says thaumatin or brazzein is "a thousand times sweeter than sucrose," it almost always means by weight: a milligram of protein sweetens like roughly a gram of sugar. But the receptor does not count milligrams. It counts molecules — moles — and on that scale the proteins are far more astonishing still, because they are so much heavier than a sugar.

Run the unit conversion. Sucrose has a molar mass of about **342 daltons**. Brazzein is **54 residues**, roughly **6.5 kilodaltons** folded — nineteen times the mass of a sucrose molecule (6,500 ÷ 342 ≈ 19). So one gram of brazzein contains about one-nineteenth as many molecules as one gram of sucrose. If that gram is "a thousand-fold sweeter by weight" — a conservative figure for these proteins — then molecule for molecule it is a thousand times nineteen, on the

order of **nineteen thousand times sweeter by mole**. A per-weight figure that already sounds like a misprint is, at the receptor, almost twenty times larger again.

The same conversion tells us what "a few parts per million in a fruit" actually means in the cell. Three parts per million is 3 milligrams of brazzein per kilogram of tissue; in the watery interior of the fruit that is close to 3 milligrams per litre, and dividing by the 6.5-kilodalton molar mass gives

> 3×10^{-3} g·L^{-1} ÷ 6,500 g·mol^{-1} ≈ 4.6×10^{-7} mol·L^{-1} ≈ **0.46 micromolar.**

Sub-micromolar — and that trace, pressed as a wedge against a receptor it will not let go, is enough to carry a real sweetness. A sugar needs hundreds of millimolar, a hundred thousand times more molecules, to do comparable work. The two numbers, the per-mole potency and the sub-micromolar working concentration, are consistent to about an order of magnitude, which is as tightly as a receptor anchor and a sweetness threshold can honestly be made to agree. We state the arithmetic rather than the slogan because the arithmetic is the thing that is true.

Five proteins, and why they differ

Nature has given us five well-characterized sweet proteins, and they are not interchangeable. Their differences are structural, and those differences decide which is worth the trouble of engineering into a food.

PROTEIN	SOURCE	SIZE	STABILITY	NOTE
Brazzein	*Pentadiplandia brazzeana*	54 aa, 4 S–S	heat- & pH-stable (to ~98 °C)	single chain; best engineering target
Thaumatin	*Thaumatococcus daniellii*	207 aa, 8 S–S	stable; GRAS flavour (E957)	very slow onset, liquorice linger
Monellin	*Dioscoreophyllum cumminsii*	two chains (44 + 50 aa)	heat-labile (unfolds ~50 °C)	single-chain MNEI restores stability
Mabinlin	*Capparis masaikai*	~33 + 72 aa	heat-stable (disulfide-locked)	durable, modest potency
Miraculin	*Synsepalum dulcificum*	glycoprotein	pH-dependent	*modifier* — turns sour to sweet

Thaumatin is the elder statesman: large, intensely sweet, and already carried on the regulatory books as a flavour (it holds GRAS status and the European additive number E957). Its liability is temporal. Its sweetness arrives slowly and leaves a liquorice-tinged linger that rules it out wherever you want a clean, sugar-like attack — yet below its own sweetness threshold it keeps working as a flavour *enhancer*, lengthening and rounding other tastes the way a kokumi compound does. That second use keeps it valuable even where its onset disqualifies it as the primary sweetener.

Monellin taught the whole field its most important lesson in protein engineering, and it taught it by failing. In its natural form monellin is two separate chains that associate only loosely; warm it past about 50 °C and they let go of each other, the fold falls apart, and the sweetness is simply gone — useless for anything cooked. Researchers answered by fusing the two chains into one continuous polypeptide, **single-chain monellin (MNEI)**, which folds as a single stable unit and holds its sweetness through heat. It was clean proof of a principle we build on every day: a sweet protein's liabilities are themselves engineerable. The molecule's weakness was not a law of nature; it was a design we were free to revise.

Mabinlin is the durable one — locked shut by its disulfides, content to sit through heat that would unfold its cousins — but only modestly potent, a steady voice rather than a loud one. And **miraculin** is the strangest member of the family, so different that it earns a section of its own: it is not sweet at all, but a switch.

Why we build around brazzein

Of the five, **brazzein** is the one our program builds around, and the reasons are entirely practical. It is small — just fifty-four residues — so the gene is short and expresses well. It is a single chain, so there is no two-subunit assembly problem of the kind that nearly sank monellin. It is extraordinarily stable, holding its fold through heat to near 98 °C and through acid that would denature most proteins, which means a food containing it survives cooking and processing. And it is well-studied enough that its structure–activity relationship is partly mapped, so we are designing from known contacts rather than from guesses.

That map is specific. Brazzein's sweetness is carried by three surface regions that together present its receptor-binding face: a loop and short strand around residues **His31 and Arg33**, and a charged cluster near **Glu36 and Glu41**. These are the outer edges of the key. Published work shows that removing the N-terminal pyroglutamate — a blocked, cyclized first residue — roughly **doubles** potency, and that particular charge substitutions in those loops tune both the intensity and the length of the lingering tail. Because the protein is small, single-chain, and pinned by four disulfide bonds into near-indestructibility, those surfaces can be re-sculpted by mutation without the whole fold collapsing around the change. A fragile protein cannot be edited; every substitution risks the fold. Brazzein's robustness is precisely what makes it *designable*, and that is why it is the platform's lead.

It is worth naming what brazzein does *not* require, because it is the heart of why sweet proteins belong on this platform at all. Most flavour molecules — capsaicin for heat, menthol for cooling, the glutamate that carries umami — are the end of a biosynthetic pathway, a relay of enzymes each of which must be present, folded, and fed. To engineer them you engineer the whole relay. Brazzein has no pathway. It is translated directly from its own gene, folded, its four disulfides formed, and secreted — a single reading of a single

sequence. There is nothing between the gene and the sweetness but the ribosome and the cell's ordinary folding machinery. For a flavour, that is the shortest road that exists.

The design pipeline, residue by residue

The design pipeline is the one this library runs on every target, and its first step is to fold the protein in the computer and colour every residue by the model's confidence. High-confidence blue marks the rigid scaffold — the disulfide core you must not disturb. The warmer, lower-confidence stretches mark the flexible, receptor-facing loops, which is exactly where sweetness is decided and where we are free to design. A companion plot, the predicted aligned error, reads the fold a second way: it shows which parts of the structure are positioned confidently *relative to each other*, so we can see at a glance which loops move as a unit with the binding face and which float independently. Together they tell us where an edit will propagate into the receptor contact and where it will be absorbed harmlessly by the scaffold.

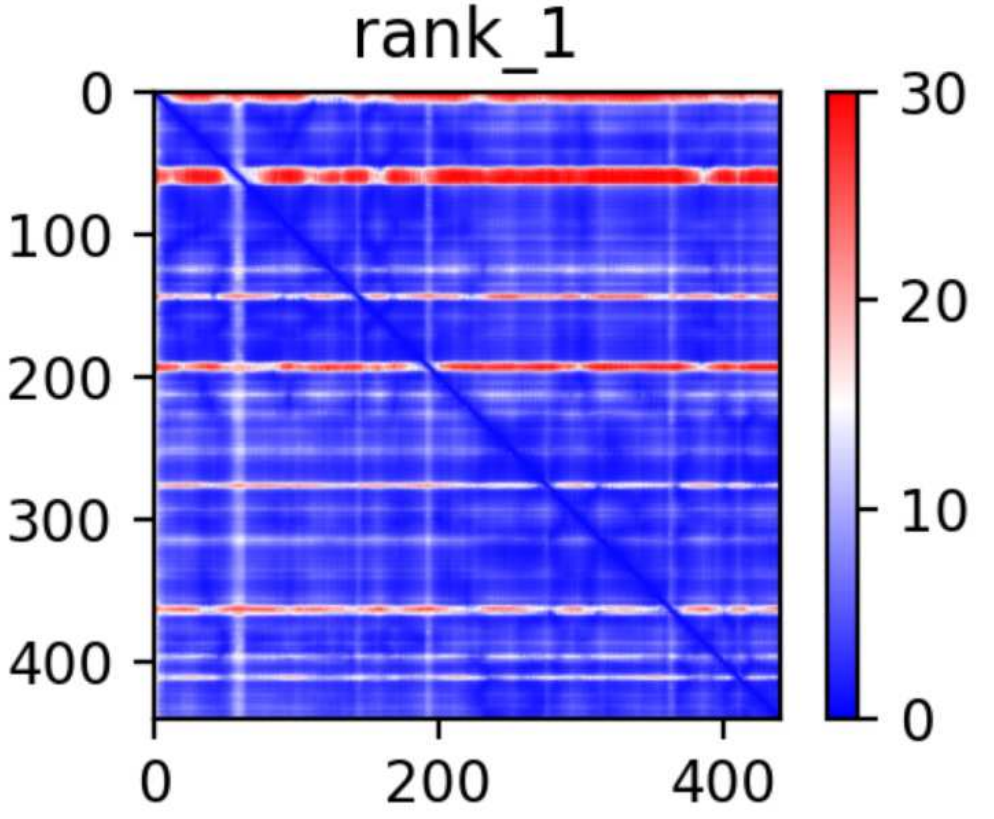

The predicted-aligned-error map for the same fold — dark where two residues are confidently placed relative to one another. The rigid disulfide core reads as a solid block; the receptor-facing loops show the looser coupling that makes them safe, and meaningful, to edit.

With the editable surface identified, an **ESM2 protein-language model** ranks candidate substitutions in those loops by predicted effect on fold stability and receptor contact; the top candidates are modelled, docked, and — where they survive — written in with CRISPR and expressed. Here is the pipeline's output for brazzein,

ranked by predicted change in fold stability (ΔΔG, where negative means *more* stable), with a derived column that translates each ΔΔG into something physical: how far it shifts the folded-versus-unfolded equilibrium. A stabilizing edit multiplies the ratio of folded to unfolded molecules by $\exp(-\Delta\Delta G / RT)$; at room temperature $RT \approx 0.593$ kcal·mol^{-1}, so a −1.1 kcal·mol^{-1} edit shifts that balance by $\exp(1.1 / 0.593) \approx$ **6.4-fold** toward the folded, sweet state. That is what a single modest number in the ΔΔG column is actually worth.

SUBSTITUTION	PREDICTED ΔΔG (KCAL·MOL⁻¹)	FOLDED-STATE SHIFT AT 298 K *(DERIVED)*	PREDICTED EFFECT
des-pGlu1 (remove N-term)	−0.2	≈ 1.4×	potency ↑ (~2×), per literature
H31R	−0.6	≈ 2.8×	potency ↑, sharper onset
E36D	−0.9	≈ 4.6×	stability ↑, linger ↓
E41A	−0.7	≈ 3.3×	potency ↑ at low dose
D29N	−1.1	≈ 6.4×	fold stability ↑

Two honesties belong beside that table. The first: these are **model predictions for a frontier target**, not measured production values — they are grounded in the published structure–activity map so the model reasons from real contacts, but they are predictions, and we mark them as such everywhere they appear. The second concerns the des-pGlu1 row and the potency edits generally. When an edit roughly doubles potency, it halves the EC_{50}, and the index moves in lockstep: $IFI = 1{,}000 \times EC_{50}(\text{ref}) / (½ \cdot EC_{50})$ = **twice** the prior score. The thermodynamic column and the index column are two readings of the same molecular event — a loop that holds its shape a little better, pressed a little harder against a receptor that will not let go.

Teaching a protein to fold inside a fruit

Designing the sequence is the easy half. Making the protein at scale, folded correctly, inside a living plant, is its own discipline — and it is why "express brazzein in a fruit" is a real multi-year program and not a slogan.

The difficulty is the disulfides. Brazzein's four S–S bonds are what make it indestructible once formed, but they must be formed *correctly*, each cysteine paired with its right partner, or the protein misfolds into something inert or, worse, insoluble. Disulfide formation happens in a specific, oxidizing cellular compartment — the endoplasmic reticulum and the secretory pathway beyond it — so the host must route the nascent protein there and give its folding catalysts time to work. Yeast such as *Pichia* can manage it, and so, for our purposes, can the plant's own secretory pathway; but **yield and correct folding are the gating problems**, not sequence design. This is why a share of the pipeline's stability edits are aimed as much at *foldability in the host* as at potency at the receptor: a variant that folds cleanly at twice the rate is worth more in a fruit than a slightly sweeter one that aggregates. The ΔΔG column, read this second way, is a foldability column — a −1.1 kcal·mol^{-1} edit that shifts the folded equilibrium 6.4-fold is also 6.4-fold fewer molecules lost to misfolding on the way to the vacuole.

Precision: the CV that makes a sweetener trustworthy

A sweetener that varies from fruit to fruit is not a sweetener; it is a surprise. The discipline that makes the whole platform usable is the coefficient of variation — the batch-to-batch standard deviation divided by the mean, expressed as a percent — and our standard is **CV ≤ 7%**, against the roughly **42%** swing typical of an unengineered plant left to its own regulation.

The arithmetic makes the stakes concrete. Take an illustrative target of 3.0 ppm brazzein in the fruit. At CV = 7%, the standard deviation is 0.07 × 3.0 = **0.21 ppm**, so about two-thirds of fruit fall between 2.79 and 3.21 ppm and ninety-five percent within roughly 2.58–3.42 ppm — a spread no palate would catch. At the wild-type 42%, the standard deviation balloons to 0.42 × 3.0 = **1.26 ppm**, and a single standard deviation now runs 1.74–4.26 ppm. That is the difference between a flat sweetness and a nearly flat one, a swing you would taste across a single harvest. Holding CV under 7% is, in the end, the same act as holding the fold: a protein that folds reliably, expressed from a construct whose regulation we control, lands near its target every time. Consistency is not a boast about quality control. It is what lets a formulator — or a family at a table — trust that tomorrow's fruit tastes like today's.

The destination is a fruit, not a powder

The destination for all of this is not a powder in a jar. It is a *fruit* — a strawberry, a tomato, a cane — that carries brazzein at a few parts per million in its own tissue, so that the sweetness and the food are the same object, grown together, with no sugar between them. The sub-micromolar arithmetic is what makes this plausible: a trace of a tens-of-thousands-fold-potent protein is all the tissue needs to hold.



A strawberry in cross-section — the kind of tissue we design brazzein to be grown inside, so the sweetness is not added to the fruit but manufactured by it.

The sweet axis already has a living proof that a fruit can carry a designed sweetness alongside everything else it tastes of. **Phoenix (SS-0096)**, the first cultivar designed to register on all eight axes at once, carries its sweetness as one voice in a full chord:

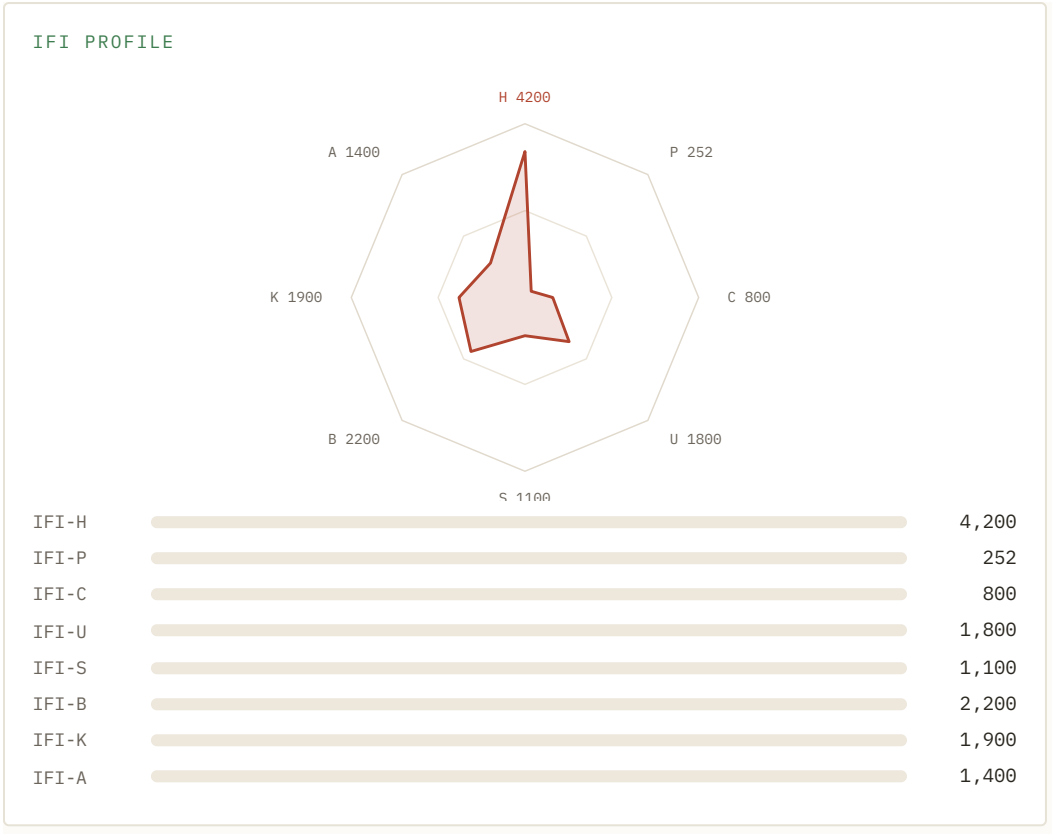

Read the sweet line against the rest: **S 1,100**, that 1.1×-reference sweetness, sitting in balance with heat, umami, and a designed bitterness — our model predicts that a sweet axis can be dialed into a fruit at a chosen, modest loudness rather than slammed to its maximum. Phoenix's sweetness is small-molecule sweetness; the protein route is the frontier beyond it, and the hosts it is aimed at are already on the roadmap:

HOST	ENGINEERED CONSTRUCTS ON FILE	SWEET-AXIS ROLE	DIFFICULTY
Tomato	2,040	sweetness + umami restored together	Easy
Stevia	910	the glycoside incumbent; protein comparator	Moderate
Sugarcane	200	sweetness at the source of sugar itself	Frontier
Monk Fruit	80	the mogroside incumbent	Hard

Sugarcane is the most pointed host of all — to grow, in the very plant we press for sugar, a sweetness that is not sugar. The sharpest expression of the idea is the **Sweet-Heat Strawberry** on the cross-family tier of the roadmap — a seven-to-ten-gene build on a thirty-six-to-forty-eight-month horizon, carried in the concept series as *Warm Kiss*, *Sugar Burn*, *Candied Fang*, and *Devil's Dessert*. A strawberry that is sweeter than it has any right to be, warmed by a trace of designed heat, its sweetness manufactured by the fruit from a gene rather than loaded in as sugar: that is the farthest honest reach of the sweet axis, and we name it as a thing we are building toward.

Honest limits, and the patience the frontier asks

We would not write this chapter without naming plainly where the sweet-protein program actually stands, because the voice of this whole library is that numbers, and limits, are forms of respect.

The first limit is regulatory. A novel sweet protein expressed in a food crop is a new-to-market ingredient, and **brazzein in particular does not yet carry broad approval** of the kind thaumatin has earned. That is not a detail to be engineered around; it is a road to be walked, with the data to walk it. The second limit is that the ΔΔG table is a set of **model predictions**, not measured production values — grounded in real structure–activity data, but predictions, and we mark them so.

The failure modes are concrete, and we design against each. A protein can **misfold** in the host and never form its sweetness surface — the reason foldability sits in the objective function beside potency. It can **aggregate** at the concentrations expression demands, taking itself out of solution. Its sweetness can arrive with an **off-note or a slow onset and a long linger**, the way thaumatin's liquorice tail disqualifies it from clean applications, which is why the loop edits target the temporal shape and not only the intensity. And because it is a protein eaten as a protein, its **digestibility and any immune profile** have to be understood honestly before it belongs in a food, not assumed away. None of these is a reason to stop. Each is a reason to measure.

So the sweet-protein work sits where it honestly belongs on our roadmap: as frontier, its destination a thing we are building toward rather than selling today. The structures are known and the science is sound. The patience is the point — the same patience that let a two-chain monellin become a stable single chain once someone was willing to take the time. We are vegetarians who have spent a lifetime thinking about what we put in our bodies, and we would rather ship a sweetness slowly and truthfully than quickly and loudly. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Miraculin — sweetness revealed, not added

And then there is the strangest member of the family, which points at a different future entirely. **Miraculin**, from the West African miracle berry, is not sweet at all. It is a *modifier* — a glycoprotein that binds the sweet receptor and sits there quietly, doing nothing, until acid arrives. At the low pH of a lemon or a tart berry, miraculin changes shape and switches the receptor *on*, so that sourness itself is read as sweetness. A wedge of lime tastes of sugared lime; a mouthful of plain lemon tastes like lemonade with nothing added. The sweetness was never put into the food. It was conjured, at the moment of tasting, out of the food's own acid.

We hold miraculin on the frontier as the purest version of this book's whole idea. Every other sweetness in these pages is a thing we add to a plant, however elegantly — a protein folded into a fruit, a glycoside tuned at its enzyme. Miraculin adds nothing. It reveals a sweetness that the acid was always able to become, if only the receptor could be taught to hear it that way. That is the quiet reach past which we cannot yet see clearly: a food made sweeter not by what goes into it, but by teaching the tongue to read what is already there. Sweetness not even added — *revealed.*

Chapter 12 · The tomato — *sweeter without a grain more sugar.*

If this book has a single hero, it is the tomato — not because the tomato is the sweetest thing we grow, but because the tomato proves the thesis more completely than any sugar crop ever could. A tomato can be made to taste dramatically sweeter with *almost no change in its sugar at all*, and the science of how is the most beautiful confirmation we know that sweetness is a perception to be composed, not a quantity to be poured. Our family has put more work into this fruit than into any other — **2,040 constructs**, the largest single program in the whole library — and much of that work serves one counterintuitive fact about how sweetness is actually built in the mouth.



The tomato, restored. Most of the sweetness we grow back into it is not sugar at all — it is the aroma the fruit lost, the volatiles the brain reads as sweet.

Here is the fact, stated plainly. Taste a ripe heirloom tomato and a modern supermarket one at the *same measured Brix* — the same dissolved sugar, to the decimal — and the heirloom tastes markedly sweeter. The sugar is identical; the sweetness is not. The difference is carried by **aroma**: a set of volatile compounds — among them (Z)-3-hexenal, 6-methyl-5-hepten-2-one, geranial, and a dozen relatives — that, perceived through the back of the nose as you chew, are read by the brain as *sweetness*, and raise the rated sweetness of the fruit **independently of its sugar content**. This is not

folklore. It is one of the best-documented results in flavor science, and it is the literal, molecular embodiment of this book's title: sweetness without the sugar.

The experiment that proves it

The cleanest demonstration is one anyone can run, and we run a disciplined version of it every week. Take two tomatoes trimmed to identical Brix — the dissolved-solids reading a refractometer gives, essentially grams of sugar per hundred grams of juice. Present them blind. A tasting panel will sort them, reliably and in the same direction, into *tastes sweeter* and *tastes flat*. The sorting is not sugar; the sugar was held equal by construction. It is aroma, and the proof is a control so simple it feels like a magic trick: **have the taster pinch the nose shut and chew.** With the retronasal passage closed, the two fruits collapse onto each other. The sweeter one stops being sweeter. Release the nose and the gap springs back. What that one gesture isolates is the whole chapter — the sweetness lives in the smell, and the smell reaches the brain not from the front of the nose but from the back of the mouth, rising up the nasopharynx as the fruit is crushed and warmed.

That retronasal path is why eating is not tasting. Smell a cut tomato and you get *orthonasal* aroma — pleasant, informative, separate. Chew it and the same molecules take the back stairs to the olfactory epithelium while the tongue reports sugar, acid, and glutamate, and somewhere past the receptors the two streams are fused into one object the brain simply calls *flavor*. A handful of the tomato's volatiles happen to land, in that fusion, on the sweet side of the ledger. A fruit rich in them tastes sweet beyond its sugar. A fruit stripped of them tastes thin even when the Brix is respectable.

Two receptors, one word — how the tongue and the nose both say "sweet"

To design sweetness we had to understand that the word names two machines that the brain refuses to keep separate.

The first is the sweet receptor itself: **TAS1R2 together with TAS1R3**, a heterodimer of two class-C G-protein-coupled receptors standing in the membrane of the taste cell. Each subunit carries a large extracellular *Venus flytrap* domain — two lobes hinged like a shell — and a sugar molecule binding in the cleft lets the shell close. That closure tips the transmembrane bundle, which releases the taste cell's G-protein, **gustducin**; gustducin frees a βγ subunit that switches on **phospholipase C-β2**; PLCβ2 cleaves membrane PIP_2 into **IP_3**; IP_3 opens calcium stores; the rising calcium gates the **TRPM5** channel; the cell depolarizes and spills **ATP** onto the afferent nerve. Sugar in, a nerve impulse out, and every step in between is a lever we can model and, in principle, move. The half-closing point of that flytrap — the sugar concentration at which the receptor fires at half its maximum — is the physical anchor of everything that follows.

The second machine is olfactory, and it is vast: not one receptor but hundreds of odorant receptors in the roof of the nose, reading the volatile plume retronasally. Their output climbs to the olfactory bulb and from there to the **orbitofrontal cortex**, the same stretch of cortex

that receives the taste signal. There the two are written into a single record. Crucially, the brain does not keep a wall between *this came from the tongue* and *this came from the nose*. A volatile that has, across a lifetime of ripe fruit, reliably accompanied sugar gets filed *as* sweetness — it is learned, cross-modal, and for most people indelible.

This is the subtlety that makes the index honest, and we state it rather than hide it: the sweetness axis is **anchored at the receptor but scored at the perception.** TAS1R2·R3 fixes the ruler — it defines what one unit of sweetness *is* — while the number we assign a food is the loudness of the finished percept, taste and retronasal smell fused. That is precisely why a tomato can post a sweetness far above what its sugar alone could drive at the receptor: the sugar sets a floor, and the aroma, routed through a different set of receptors entirely, is added to the same ledger by the cortex that reads them both.

A number for loudness

Flavor resisted measurement for ten thousand years because it was described in adjectives. The Iyengar Flavor Index replaces the adjective with a number anchored to the one event that matters — the half-maximal activation of a receptor, the midpoint of its dose–response curve.

A receptor's response to a stimulus of concentration c follows a Hill relation:

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

At $c = EC_{50}$ the response is exactly half of maximal — that is the definition of EC_{50} — and on a log-concentration axis that midpoint is where the curve is steepest and most reproducible. A small, honest shift in potency produces the largest, cleanest change in response there, and measurement error is smallest. Anchor the index at EC_{50} and you have anchored it at the most trustworthy point on the curve. For the sweet receptor the Hill coefficient n sits near unity, so sweetness is a graded, well-behaved response — no sharp cooperative switch, just a smooth climb we can place a fruit along.

The index inverts EC_{50} against a fixed reference and scales by a thousand:

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

The logic is loudness, not quantity. A *more* potent stimulus reaches half-activation at a *lower* concentration — a smaller EC_{50} — so the ratio rises. The reference is chosen by convention so that it scores exactly 1,000 against itself: for sweetness the reference is pure sucrose, and $IFI(ref) = 1{,}000 \times EC_{50}(ref)/EC_{50}(ref) = 1{,}000$ by construction. Everything else is read as a multiple of sucrose's loudness.

Run it forward once. Suppose a stimulus reaches half-maximal sweet response at 0.40 times the sucrose reference concentration. Then IFI = 1,000 × 1 / 0.40 = **2,500** — and 2,500 is exactly where our

model places the restored tomato on the sweet axis. Run it backward from each of the tomato's modeled scores and you recover the physics behind the scorecard:

AXIS (RESTORED TOMATO)	IFI	EC_{50}(REF)/EC_{50}(FRUIT)	READS HALF-MAXIMAL AT	MEANING
Sweetness · TAS1R2·R3	2,500	2.50	0.40 × reference	2.5× as loud as sucrose
Umami · mGluR1	2,400	2.40	0.42 × reference	savory depth restored
Kokumi · CaSR	1,100	1.10	0.91 × reference	a present, supporting floor
Astringency · PRPs	300	0.30	3.33 × reference	faint by design — below reference

The last row teaches the formula as clearly as the first. An IFI *under* 1,000 means a stimulus **less** potent than the reference — a higher EC_{50}, a quieter note. The tomato's astringency of 300 is not a failure to reach a target; it is a deliberately soft note, one-third the loudness of the reference tannin, the proper whisper of real tomato skin. The index is as comfortable describing restraint as excess.

Aroma as a shift in the curve

Now the mathematics of the magic trick. The aroma does not add sugar and does not touch TAS1R2·R3 directly; the volatiles speak to olfactory receptors a storey up. What they move is the *perceptual* dose–response — the curve relating the fruit's actual sugar to the sweetness the brain reports. Retronasal aroma **left-shifts that curve.** The same sugar now reads as a higher sweet intensity, exactly as if the effective EC_{50} of the whole percept had dropped, even though no molecular affinity at the receptor has changed at all.

Call the size of that shift the enhancement factor **φ**, defined as the perceived sweetness with the aroma intact divided by the perceived sweetness with it stripped:

```
φ = IFI-S(aroma intact) ÷ IFI-S(aroma stripped)
```

And φ is not a fitted abstraction — the pinched nose measures it directly. Close the retronasal path and you read the stripped, sugar-only sweetness; open it and you read the whole. For our restored fruit the sugar-only reading sits near the sucrose reference (≈ 1,000), and the intact fruit posts **2,500**, so φ ≈ 2,500 / 1,000 = **2.5**. The aroma is doing more work than the sugar — roughly 60% of the perceived sweetness is smell, and only the remaining ~40% is molecule on receptor.

Watch how little the sugar itself can be asked to contribute. Brix is a mass fraction: 1 °Bx ≈ 1 g sucrose per 100 g ≈ 10 g/L, and since sucrose weighs 342.3 g/mol, that is 10 ÷ 342.3 ≈ **29 mmol/L** of sweet stimulus per degree Brix. Hold a fruit near a supermarket 4.5 °Bx (≈ 131 mM) and lift it, honestly, to 5.0 °Bx (≈ 146 mM): a 0.5 °Bx, roughly **11% increase in sugar.** Against a sweetness response that climbs sublinearly near this range, 11% more sugar buys only a few percent more perceived sweetness — nowhere near "dramatically sweeter." The table makes the division of labor exact:

FRUIT	BRIX (°BX)	SUGAR (MM)	AROMA FACTOR Φ	SWEET-EQUIVALENT (°BX)	MODELED IFI-S
Supermarket, aroma gutted	4.5	131	1.0	4.5	~900
Same fruit, +11% sugar only	5.0	146	1.0	5.0	~1,000
Restored, +11% sugar **and** aroma	5.0	146	2.5	**12.5**	**2,500**

The restored tomato is designed to taste as sweet as a 12.5-°Bx fruit would taste with no aroma at all — a sugar level a fresh tomato essentially never reaches. We got there by moving φ from 1.0 to 2.5, not by pouring in sugar. The decimal of added Brix is almost a rounding error; the perception more than doubled.

The volatiles, named and traced to their genes

These compounds are not mysterious. They have names, measured thresholds, and — the part that makes them *engineerable* — known enzymes and known genes. They arrive by three independent routes, which is a gift: three separate handles to grip when we grow the aroma back.

VOLATILE	PRECURSOR	ENZYMATIC ROUTE	READS TO THE BRAIN AS
(Z)-3-hexenal, (E)-2-hexenal, hexanal	C18 fatty acids (linolenic/linoleic)	lipoxygenase (TomLoxC) → hydroperoxide lyase	`fresh, green, "ripe tomato"`
6-methyl-5-hepten-2-one	lycopene (the red pigment)	carotenoid cleavage dioxygenase (CCD1)	`sweet, fruity`
geranial + neral (citral), β-ionone	carotenoids / terpenoid pool	CCD1 cleavage	`sweet, citrus, floral`
1-nitro-2-phenylethane, phenylacetaldehyde	phenylalanine	aromatic amino-acid route	`sweet, floral, honeyed`
3-methylbutanal, 2-methylbutanol	leucine / isoleucine / valine	branched-chain aminotransferase (BCAT)	`malty, warm, sweet`

Trace one route fully, because it is the most telling. The apocarotenoid volatiles are **cut from the pigments themselves.** Lycopene and β-carotene — the molecules that make a ripe tomato red and that carry the provitamin-A nourishment of the fruit — are substrates for **carotenoid cleavage dioxygenase (CCD1)**, which snips specific double bonds along the chain. Cleave lycopene and you release 6-methyl-5-hepten-2-one, a low-threshold "sweet, fruity" note; cleave the β-ring carotenoids and you get β-ionone and the citral pair, floral and sweet. So the color, the nourishment, and the sweetness are three branches of a single tree: the pigment that feeds the body is the feedstock for the aroma that fools it, happily, into tasting sugar. This is why our restoration touches the carotenoid

pathway and the sweetness perception in one stroke — grow back the pigment and you grow back the nutrition *and* the smell the brain reads as sweet.

The second route is lipid. The fruit's polyunsaturated fatty acids are oxygenated by a **lipoxygenase** — in the tomato the decisive one is *TomLoxC* — and the resulting hydroperoxide is cleaved to the C6 aldehydes, chief among them (Z)-3-hexenal, the quintessential just-picked top note. The third route is amino-acid-derived: phenylalanine yields the honeyed phenyl volatiles, and the branched-chain amino acids, handled by **BCAT**, yield the malty aldehydes. Lipoxygenase, carotenoid cleavage, branched-chain transamination — three doors into the same room, which is exactly why the aroma can be rebuilt rather than merely mourned.



Grown back toward what it used to be. The aroma that reads as sweet is cut from the same carotenoid pigment that colors and nourishes the fruit — color, nutrition, and sweetness are branches of one tree.

The own-goal, precisely — why the truck took it

The modern tomato was bred for yield, uniformity, shelf life, and the brutality of the truck, and in chasing those it lost both its sugar and, far more damagingly, these sweetness-making volatiles. The mechanism of the loss is one of the best-documented self-inflicted wounds in the history of breeding, and it is worth naming exactly.

Breeders prized a fruit that ripened to an even red with no green shoulder — easier to harvest on a schedule, prettier in a crate. The mutation that delivers it, *uniform ripening*, is a loss-of-function in a transcription factor called **GLK2** (GOLDEN 2-LIKE 2), which drives chloroplast development in the green fruit. Knock out GLK2 and the fruit's chloroplasts are fewer and weaker, so the fruit

photosynthesizes less of its own sugar — and, because those same chloroplasts mature into the carotenoid-packed bodies of the ripe fruit, it loads **less pigment** too. Less pigment is less lycopene and β-carotene, which is less substrate for CCD1, which is **less sweet aroma.** One cosmetic selection, chosen for the crate, simultaneously lowered the sugar, the color, the provitamin-A nourishment, and the volatiles the brain reads as sweet. The flavor genes did not even have to be targeted; many simply sat on the chromosome near the traits being selected and were dragged out as linked baggage.

We keep a shorthand for the theft, borrowed from our savory work: *the 93% the truck took.* Put arithmetic to it. The residual umami in a supermarket tomato reads near **IFI-U 300.** If that 300 is the roughly 7% the breeders left behind, the loud ancestral fruit sat near 300 ÷ 0.07 ≈ **4,300** on the same axis. That inference matters, because it sets an honest ceiling against which to measure what we actually restore. The supermarket tomato is heavy, red, sturdy, and nearly silent — its sugar diminished, but its *aroma* gutted, so that even a fruit of decent Brix reads flat because the smell that should have told the brain "sweet" is simply not there.

Growing it back as a chord, not a note

So we grow it back — and not as one trait but as a chord, which is why the tomato program spans four axes (S, U, F, and M) across those 2,040 constructs. The sweetness restoration has three movements, and then a floor beneath them.

First, a modest, honest lift in sugar — the 4.5 → 5.0 °Bx we modeled above — tuned by a genome-scale flux-balance model so the plant can afford the extra photosynthate without shrinking the fruit. Sugar is the smallest of the three movements on purpose; the math already showed it cannot carry the sweetness alone.

Second, and carrying most of the perceived sweetness, the volatiles — restoring the lipoxygenase, carotenoid-cleavage, and branched-chain routes the truck bred out, so the fruit reads sweet through the nose as a ripe one should. This is where φ moves from 1.0 toward 2.5.

Third, the floor — a restoration of umami through the native **GAD/GDH1** glutamate route, lifting free glutamate at mGluR1, and of the **γ-glutamyl kokumi peptides** (γ-glutamyl-valyl-glycine and its relatives) that speak to the **CaSR** receptor, so the recovered sweetness lands full and continuous rather than thin. Umami and kokumi are not sweetness, but the cortex that fuses taste and smell also lets a savory, rounded context deepen how complete a sweet fruit feels in the mouth.

The ledger makes the division of labor explicit. These are modeled shares of the finished sweet reading; the **total of 2,500 is our modeled design target**, and the point of the table is where it comes from:

MOVEMENT	MECHANISM	SUGAR ACTUALLY ADDED	MODELED SHARE OF IFI-S
Sugar lift	flux-balanced Brix, 4.5 → 5.0	+0.5 °Bx (~11%)	~400
Volatile restoration	LOX · CCD1 · BCAT routes regrown	none	~1,500
Umami–kokumi floor	GAD/GDH1 glutamate · γ-glutamyl peptides	none	~600
Restored tomato	the chord	**+0.5 °Bx total**	**2,500**

Three of the four rows add no sugar whatsoever, and together they contribute roughly five-sixths of the number. That is the thesis rendered as a bill of materials.



In cross-section, the fruit's chemistry is a composition across axes: sugars and sweetness-volatiles for S, glutamate for U, the γ-glutamyl peptides for the kokumi floor. We restore the whole chord, not one note.

The scorecard

The result is a tomato whose whole story fits in eight numbers:

IFI PROFILE

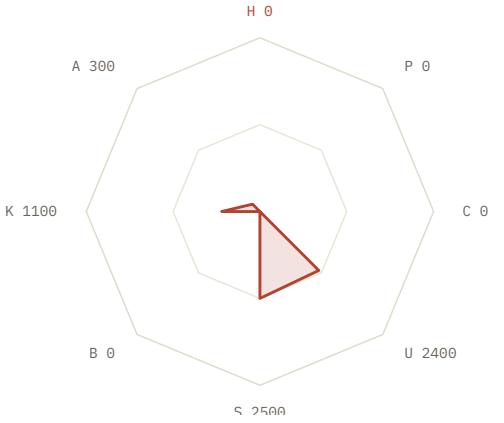

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	2,400
IFI-S	2,500
IFI-B	0
IFI-K	1,100
IFI-A	300

Read it slowly. The **S of 2,500** is a vivid, ripe, almost shocking sweetness — and its Brix is only a little above the supermarket fruit's, because most of that number is aroma and floor, not sugar. The **U of 2,400** and **K of 1,100** are the savory depth and the richness that make the sweetness feel like a whole tomato rather than a sweetened one. There is no heat, no cooling, no pungency — this is a fruit, not a chili. And the small **A of 300** is the faint, proper

astringency of real tomato skin, left in on purpose, because we are restoring a tomato, not idealizing one into something it was never meant to be.

For contrast, hold it against the most maximal construct we have ever built — the eight-axis cultivar we call Phoenix, engineered to register on every axis at once, our model predicts approximately 2.95 million SHU of designed heat:

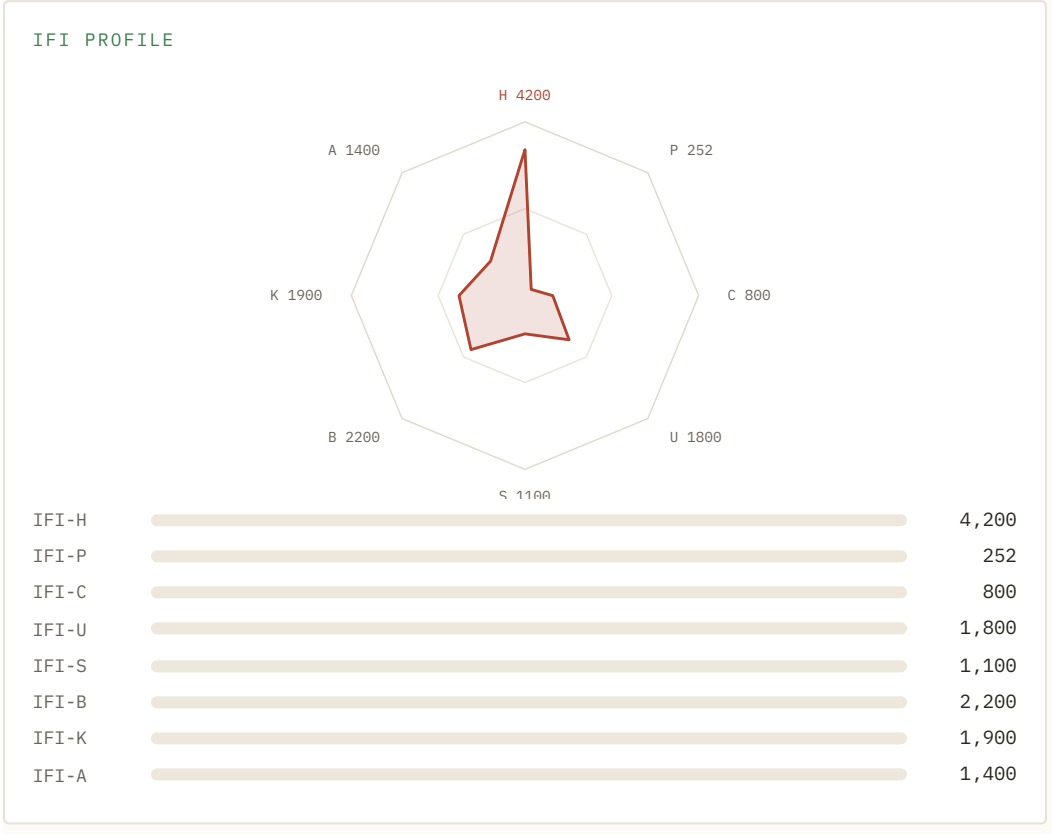

Look only at the sweet line. Phoenix, built to be loud on all eight axes, reads **S 1,100.** The quiet tomato reads **2,500** — more than twice as sweet on the same receptor, and it does it with a rounding error of added sugar while the superhot does it not at all. Sweetness, the comparison insists, is composed, not poured. The fruit that never tried to be a sugar bomb out-sweetens the one engineered to overwhelm.

Error, edge cases, and what we do not claim

A number is only respect for the reader if it comes with its error. Our platform holds a batch-to-batch coefficient of variation of **≤7%** as a design target, against the **42%** we measure in unselected wild-type material — the difference between a trait written into the genome and a trait left to the weather. At CV = 7%, the sweet score of 2,500 carries a standard deviation of 0.07 × 2,500 = **175**, so a 95% interval (±1.96σ) is 2,500 ± 343 — roughly **2,160 to 2,840.** A tight, trustworthy band. Let that same fruit vary like a field crop at 42% and the interval balloons to 2,500 ± 2,058, from 442 to 4,558: a number so loose it means nothing. Genetic control of the volatile routes is not a luxury; it is what makes the scorecard a meaningful number.

The honesty extends to the limits. We design umami to reach **2,400**, not to the ≈ 4,300 ceiling the ancestral fruit may have carried — a deliberate ~56%, because flux balance will only let the plant spend what it can afford without shrinking the fruit, and we would rather

have a full tomato than a loud one. We leave the astringency at **300** on purpose. And the whole sweet effect rests on a mechanism with real failure modes we design around:

- **Run the lipoxygenase route too hot** and the green C6 aldehydes tip from "fresh" to grassy, even soapy — more aroma is not more sweetness past a threshold, and the panel catches it immediately.
- **The trick needs an open nose.** A head cold, a pinched nostril, or congenital anosmia closes the retronasal path and collapses φ toward 1.0 — the fruit falls back to its honest sugar-only sweetness. This is not a flaw in the fruit; it is the proof of the mechanism, the same control that demonstrated the effect in the first place.
- **People differ.** The cross-modal link between these volatiles and "sweet" is partly learned and partly genetic, so the enhancement is a touch louder for some tasters than others. We tune to the panel median and report the spread rather than the flattering single number.

And one thing we never claim. The carotenoids we grow back carry provitamin-A nourishment, and we care about that deeply — food grown to hold the most it can, given to the body with intention — but that is a statement about what is *in* the fruit, never a promise about what it will do for anyone. We restore a tomato. We let it be nourishing. We make no cure of it.

What the fruit proves

Read the scorecard one last time and the whole book is in it. A grain more sugar did almost none of this — a half-degree of Brix, a rounding error. The plant, grown back toward what it used to be before the truck took its voice, did all of it: the aroma cut from its own pigment, the glutamate and the kokumi peptides that give the sweetness a floor to stand on, the restraint that leaves a little honest astringency in the skin. More than any sweetener in this library, the tomato proves that sweetness is a perception to be composed — and that most of it was never sugar at all. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.



PART IV OF V

The Frontier

sweet heat · the dessert frontier · sweetness inside the full spectrum

Chapter 13 · Sweet Heat — *the strawberry that bites back.*

Everywhere else in this book, sweetness is a thing we pursue for its own sake — cleaned of its passengers, freed from its sugar, coaxed to speak at full volume without a single calorie of cost to the body that receives it. Here we do the opposite on purpose. We give sweetness a passenger it has never carried, and we choose, deliberately, the most dangerous molecule in the whole library. We put heat into a strawberry. This is the dessert frontier — the roadmap's T3 cross-family tier — and its flagship is a sweet-heat strawberry program whose named lines read like a warning and a menu at once: *Warm Kiss, Sugar Burn, Candied Fang, Devil's Dessert.* They are, as we write this, design concepts rather than shipping cultivars; they sit seven to ten genes out, thirty-six to forty-eight months of patient work. But the chemistry that makes them possible is already proven on the other side of the library, and the reason to want them is the oldest trend in flavor, made molecular at last.

The appeal of sweet-heat — *swicy*, in the trade — is not that sweet and hot occupy the same bite. It is that they arrive in *sequence.* Sweetness is a fast event and heat is a slow one, and because the two live on different receptors with different clocks, putting both in one fruit does not give you a blur. It gives you a *story*: the strawberry's sugar and bright volatiles land first and read as candy, and then, as that sweetness fades, the heat rises up through it and holds. Sweet, then fire. The whole pleasure is in the choreography, and choreography is exactly what a number lets us design — because the two axes report to different nerves, carry different time constants, and can be dialed independently, to the digit.

The choreography is the point — two receptors, two clocks

Sweetness is read by **TAS1R2·TAS1R3**, a heterodimer of two class-C G-protein-coupled receptors whose large extracellular "Venus flytrap" domains clamp shut around a sweet ligand. The closure is a conformational trigger: the bound dimer activates gustducin, phospholipase C-β2 lifts inositol trisphosphate, calcium floods from intracellular stores, and the **TRPM5** channel opens to depolarize the cell. It is a cascade, but a *fast* one — the whole relay fires and the taste cell signals within roughly half a second, and because the ligand is in equilibrium with the receptor, it releases as quickly as it bound. Sweetness is a candle you can snuff.

Heat is read by **TRPV1**, a non-selective cation channel on the trigeminal nerve endings woven through the tongue — the same channel gated by temperatures above 43 °C and by acid. Capsaicin does not warm the tissue; it binds a pocket in the channel's transmembrane core and *lowers the thermal threshold*, so the body reports a fire that is not there. And TRPV1 is slow in both directions. The channel opens cooperatively, building over seconds; it

desensitizes only sluggishly; and the warmth lingers long after the molecule is gone, because capsaicin is lipophilic and loiters in the membrane. Sweetness is a candle; heat is an ember.

We can write both as the same function and let the two sets of constants do the talking. The fraction of receptor driven at a stimulus concentration C follows a Hill curve,

$$\theta(C) = C^n \,/\, (EC_{50}^n + C^n),$$

where EC_{50} is the concentration that drives the receptor to half of full response and n is the cooperativity. For sweetness, n is gentle, near **1.2**; for capsaicin on TRPV1 it is cooperative, near **2.0** — which is *why* the burn has a shoulder, a threshold below which almost nothing happens and above which it climbs steeply. The amplitudes are set by how much ligand each axis delivers to its receptor field. The *order* is set by the clocks: sweetness rising with an onset time near 0.2 s and decaying over about 2 s; heat rising over 2 s and decaying over 20. Lay the two percepts on one timeline and the sequence is not hoped for — it is a consequence of arithmetic the tongue performs for free. The sugar peaks and is already falling before the heat has finished arriving. We did not invent the gap. We learned to design inside it.

The molecule we are asking a berry to build

Capsaicin is an amide — a single nitrogen bridge joining two halves that could hardly be more different. On one side sits **vanillylamine**, an aromatic head built, atom for atom, from the same vanillin that flavors a custard. On the other sits **8-methyl-6-nonenoic acid**, a short, *branched*, singly-unsaturated fatty acid that no ordinary plant lipid pathway makes. The acyltransferase that condenses them — stripping water, forging the amide — is **Pun1**, also called CS or AT3, and it is the single enzyme at the center of everything our family does with heat. The length and saturation of that acyl chain set the ratio of capsaicin to its cousin dihydrocapsaicin, and with it the *shape* of the burn: capsaicin is the bright, forward strike; dihydrocapsaicin is the slower, deeper, longer-dwelling warmth. To design a fruit's heat is, in the end, to design that one condensation and the supply lines that feed it.

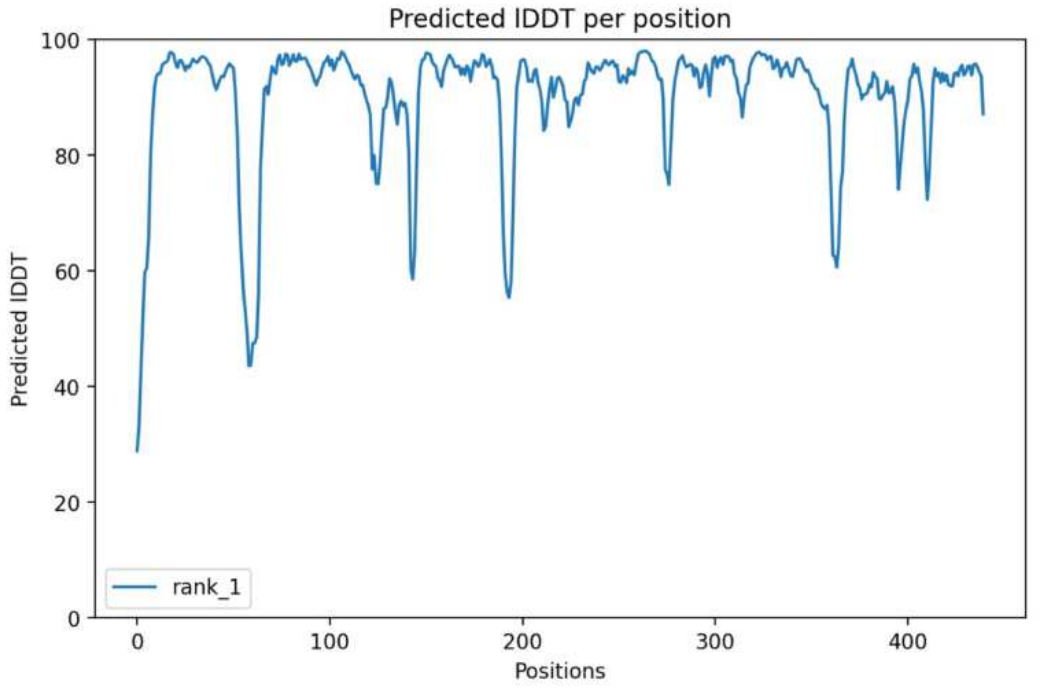

FIG. 11 — the Pun1/AT3 acyltransferase, modeled and colored by pLDDT confidence; the active-site cleft where vanillylamine and the branched-chain acyl group meet is the high-confidence core. This single enzyme — the same one whose two designed substitutions power our superhots — is the condensation we relocate into a berry.

That Pun1 condensation is not a theoretical target. It is the enzyme we have already rebuilt, residue by residue, on the pepper side of the library. A protein language model scored every single-residue substitution in Pun1 for its predicted effect on catalytic efficiency and fold stability; our flux model predicts that pathway control sits upstream in the phenylpropanoid pathway — at **4CL** and **HCT** — not at the terminal Pun1 condensation step; Pun1 is the final capsaicin synthase, labeling the product without governing the rate; and two substitutions that no breeder ever stumbled on — **S39L** and **L345G** — emerge from that analysis, lifting the enzyme's predicted output with a batch-to-batch coefficient of variation design target of **≤7%** against a wild-type's **42%**. The heat side of a sweet-heat strawberry is, quite literally, that same chemistry asked to run in a sweeter cell.

The on-ramp a strawberry already has

Here is the fortunate thing, the thing that makes this frontier reachable at all rather than a fantasy: a strawberry *already runs most of the road.* Capsaicin's aromatic half is built on the **phenylpropanoid branch — PAL → C4H → 4CL → HCT → COMT** — a sequence of enzymes present in every vascular plant, and a strawberry leans on it hard, because it is the very branch the fruit uses to make the anthocyanins that turn it red. The deep color of a ripe berry and the aromatic head of capsaicin are drawn from the same well. What the strawberry lacks is not the on-ramp. It is only the final off-ramp that peppers alone evolved.



A strawberry in cross-section — a sweet fruit about to be taught a pepper's trick. The phenylpropanoid branch it already runs for its red color is the same branch capsaicin is built on; what we add is only the terminal pepper machinery, and we add it at the surface, not the heart.

So the sweet-heat strawberry is not a fruit built from nothing. It is a fruit that has the hard upstream metabolism already paid for, handed the four terminal enzymes it is missing, and taught where to put the result. And because we are drawing on a branch the berry already uses for its own pigments, the design has a cost to respect and protect: every carbon the phenylpropanoid branch spends on vanillin is a carbon it does not spend on anthocyanin. Overdraw it and the berry pales. Part of designing the heat is leaving the color — and the ellagitannins and the ascorbate that come with a strawberry's native chemistry — fully intact, because a berry stripped of its own antioxidants to make a novelty would be a bad trade, and this family does not make it. The sweetness side, meanwhile, is restored the way we restore it everywhere: the fruit's own sugars and volatiles brought back to full voice, a kokumi floor laid underneath for continuity, acidity managed so the berry's native tartness does not mask the sweetness it is carrying.

The off-ramp only peppers evolved — four terminal enzymes

Four enzymes stand between a red strawberry and a capsaicin-bearing one, and it is worth seeing exactly what each one does, because their *balance*, not their mere presence, is what decides whether the fruit makes crisp, pepper-clean heat or a muddle.

The aromatic head is the easy half. **pAMT**, a putative aminotransferase, takes the vanillin the phenylpropanoid branch already delivers and aminates it to **vanillylamine** — one reaction, well behaved, running off a substrate the berry makes anyway.

The acyl half is the genuinely foreign machinery, and it is where the work lives. The branched-chain fatty acid is assembled from an amino acid: **BCAT**, a branched-chain aminotransferase, deaminates valine or leucine to a branched keto-acid starter; a **KAS**-type fatty-acid synthase extends it; and desaturation and chain-tailoring steps finish **8-methyl-6-nonenoic acid** to length and to its single double bond. Import this module out of balance with the amine supply and the Pun1 condensation will still run — it will simply run on whatever acyl groups are nearest, yielding nordihydrocapsaicin, homocapsaicin, and a smear of minor capsaicinoids instead of a clean capsaicin-to-dihydrocapsaicin ratio. A muddy acyl pool is a muddy burn.

The condensation itself is **Pun1/CS**, forging the amide between the two halves — the same acyltransferase, carrying the same designed substitutions, that anchors the heat side of the entire platform.

MODULE	ENZYME(S)	REACTION IT PERFORMS	NATIVE TO STRAWBERRY?
Upstream aromatic	PAL · C4H · 4CL · HCT · COMT	phenylalanine → the phenylpropanoid pool (→ vanillin)	**Yes** — runs it for anthocyanins
Amine head	pAMT	vanillin → vanillylamine	No — one enzyme to add
Branched acyl	BCAT + KAS + desaturase/tailoring	valine/leucine → 8-methyl-6-nonenoic acid	No — the foreign module, 3–4 genes
Condensation	Pun1/CS (AT3)	vanillylamine + acyl → **capsaicin**	No — the central enzyme, carrying S39L · L345G
Targeting	achene-surface promoter + vesicle sequestration	*where* capsaicinoids accumulate and how cells are shielded	No — regulatory, the safety layer

Tally the additions — one amine enzyme, three to four for the acyl module, the condensation, the tissue-targeting cassette, plus the sweetness-side restoration — and you land at the roadmap's honest count of seven to ten genes for a T3 cross-family build. Nothing here is hand-waved. Each line is a real reaction with a real catalyst, and the reason this fruit is thirty-six to forty-eight months out rather than shipping is written in that table: the acyl module has to arrive *in tune.*

Reading the design in numbers — the IFI, worked

A flavor our family cannot measure is a flavor we cannot design, and so every axis is anchored to a single number: not how much of a compound a fruit holds, but how *loudly* that compound speaks to its receptor. The anchor is the half-maximal activation we just met, inverted against a fixed reference and scaled by a thousand:

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

Read it slowly and it is almost tender in its logic. A *lower* EC_{50} means the receptor reaches half-response at a *lower* concentration — the stimulus is louder — so it sits in the denominator, and loudness climbs as EC_{50} falls. The reference is a stated convention — each axis defines IFI = 1,000 by its own chosen reference point, the way the Scoville scale is defined by dilution-to-threshold. Everything above 1,000 is louder than that convention's native reference; everything below is gentler.

The formula runs both ways, and that is the whole point of a design target: name the loudness you want, and it hands you the EC_{50} you must engineer the fruit to reach. Take the heat axis, anchored to capsaicin on TRPV1 with a reference **EC_{50}(ref) = 0.70 µM**. To design *Candied Fang* at an IFI-H of 1,400, we solve for the concentration at which its heat must reach half-activation:

EC_{50}(Candied Fang) = 1,000 × 0.70 µM ÷ 1,400 = **0.50 µM.**

Run every line the same way, on both axes, and the four named coordinates stop being characters and become *specifications* — EC_{50} values our model sets as design targets:

LINE	EC_{50}-H (MM, TRPV1)	IFI-H	EC_{50}-S (MM SUCROSE-EQ, TAS1R2·R3)	IFI-S
Warm Kiss	1.75	400	9.23	2,600
Sugar Burn	0.778	900	9.60	2,500
Candied Fang	0.500	1,400	10.00	2,400
Devil's Dessert	0.333	2,100	10.91	2,200

The sweetness column is anchored the same way, against a reference **EC_{50}(ref) = 24 mM** sucrose-equivalent: *Candied Fang*'s IFI-S of 2,400 is simply 1,000 × 24 ÷ 10.0. The two axes share one grammar and nothing else — different receptors, different reference ligands, different units entirely, µM against mM — and because each axis's IFI is defined against its own convention reference, scores across axes are not directly comparable; each axis speaks in its own currency.

The IFI sets *amplitude;* the Hill curve tells us what a real bite does with that amplitude. Take *Candied Fang* again. As the ripe flesh gives way it delivers something like 20 mM sucrose-equivalent to TAS1R2·R3, and with EC_{50}-S = 10 mM and $n \approx 1.2$:

$\theta_sweet = 20^{1.2} / (10^{1.2} + 20^{1.2}) = 36.4 / (15.9 + 36.4) \approx$ **0.70.**

A beat later the achene surface releases capsaicin to a local 0.70 µM at the trigeminal endings, and with EC_{50}-H = 0.50 µM and $n \approx 2.0$:

$\theta_heat = 0.70^2 / (0.50^2 + 0.70^2) = 0.49 / (0.25 + 0.49) \approx$ **0.66.**

Two near-equal occupancies — seventy percent sweet, sixty-six percent hot — not fighting, because they never arrive together. The sweet is at seventy and already falling when the heat reaches sixty-six and begins to hold.

It helps to put the heat on the older scale as well, and to be honest about how the two scales relate. A Scoville unit measures *concentration*: pure capsaicin is 16,000,000 SHU, and one part per million of capsaicinoid is about 16 SHU. The IFI measures *loudness*, which is occupancy, which saturates. So the two are not proportional — they are related through the Hill curve, and the relationship bends. Our full-spectrum reference pepper, **Phoenix (SS-0096)**, is predicted by our model to reach approximately **2.95 million SHU** at an IFI-H of **4,200**; that is a capsaicinoid loading of roughly 2,950,000 ÷ 16 ≈ 184,000 ppm of capsaicin-equivalent, roughly eighteen percent of pure-capsaicin potency. Climb the superhot ladder and each additional IFI-H point costs *more* Scoville than the last — near the top the local exchange rate runs close to 660 SHU per IFI-H point — precisely because receptor occupancy is saturating and concentration has to sprint to keep loudness walking. The strawberry lines sit far, far below that ladder, in familiar culinary bands — a *Warm Kiss* no hotter than a generous jalapeño's glow, a *Devil's Dessert* reaching toward a habanero's peak — and even those figures overstate the mouth's experience, because the capsaicinoids are confined to the achene surface and the whole-fruit average is milder than any single site. The dessert is not a chili in disguise. It is a sweet fruit with a hot edge, and the number keeps us honest about the difference.

Place — where a berry has no placenta

A pepper keeps its fire in a purpose-built organ. The capsaicinoids accumulate on the **placental tissue** around the seeds, synthesized at the surface of those cells and sequestered into vesicles whose membranes shield the plant's own living tissue from a molecule that would otherwise disrupt it. Evolution gave the pepper both the factory and the vault, in one place, on purpose.



A chili pepper in cross-section — the pale placental ribs where capsaicinoids are made and stored, the organ a strawberry simply does not have. Relocating the chemistry means relocating the vault too: synthesis and safe storage targeted to a defined surface, never loosed into the flesh.

A strawberry has no such organ. Its "seeds" are not inside it at all — they are **achenes**, the true fruits, studded across the outside of the swollen red receptacle we actually eat. That anatomy is, as it happens, a gift. It gives us a natural surface — the achene-bearing skin — to target the terminal enzymes toward, with promoters that switch the capsaicin machinery on *only* in that tissue and leave the sweet flesh untouched, and with the same vesicle-sequestration strategy the pepper uses, to protect the berry's own cells from what we are asking them to hold. Synthesis at the surface; storage in vesicles; the sweet interior kept clean. If the targeting leaks — if capsaicinoids diffuse inward through the flesh — the design fails twice over: the sweetness is contaminated before it can land, and the choreography collapses into exactly the blur we built the whole thing to avoid. Place is not a detail here. Place is half the invention.

Time — designing the order of a bite

The other half is time, and time is the part no cook before us could reach. The pleasure of sweet-heat is a *sequence*, and a sequence has a direction: the sugar and volatiles releasing first as the flesh yields, the heat rising from the achene surface a beat later and holding after the sweetness has gone. Because sweetness is a fast receptor event and TRPV1 warmth is a slow, building, lingering one, tuning the ratio and the spatial distribution of the two lets us set not only how loud each axis is, but *when* — the one place in all of this work where we compose in time.

Model it plainly. Give sweetness a sharp rise and a quick fall, a percept that peaks near 0.4 s and is three-quarters gone by 4 s. Give heat a slow rise and a long tail, a percept that does not peak until 6 to 8 s and then holds for twenty. The *crossover* — the instant heat overtakes sweetness — falls two to three seconds into the bite, safely after the sugar's peak, so the kiss always lands alone before the fire answers it. The designed interval between the two peaks,

call it Δt ≈ 5–7 s, is the length of the sentence the fruit speaks. The four lines are four readings of that sentence at a constant sweetness and a climbing heat:

LINE	IFI-H	IFI-S	GENES (T3)	MODELED SHU BAND	THE SENTENCE IT SPEAKS
Warm Kiss	400	2,600	~7	jalapeño-class	a warm glow, no real burn
Sugar Burn	900	2,500	~8	cayenne-class	candy first, a gentle build
Candied Fang	1,400	2,400	~8	approaching habanero	sweet front, a clear bite
Devil's Dessert	2,100	2,200	~9	habanero-class peak	a dessert that fights back

Here is *Candied Fang* — the line we would set first on a plate, sweet-forward with a real fang behind it — read as a full eight-axis design target:

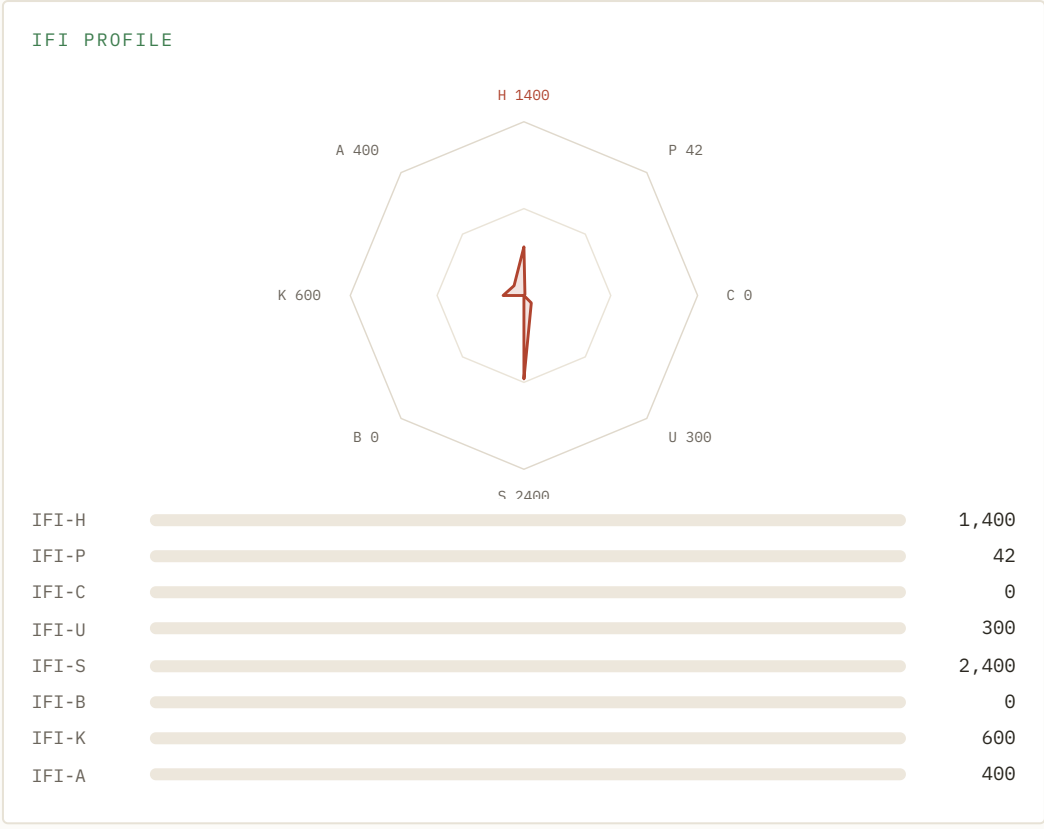

(Design target, modeled — a roadmap concept, not a shipping SKU.) The **S of 2,400** leads, as it must; the **H of 1,400** is the fang, arriving late; the small **P of 42** is the faint sharpness that rides with any real capsaicin; the **K of 600** gives the sweetness a floor of continuity to stand on; the **A of 400** is the strawberry's own seedy astringency — its ellagitannins meeting the salivary proteins — kept because it belongs to the fruit and we do not strip a berry of itself.

Devil's Dessert pushes the same composition to the edge, lifting heat until it nearly meets the sweetness that was generous first:

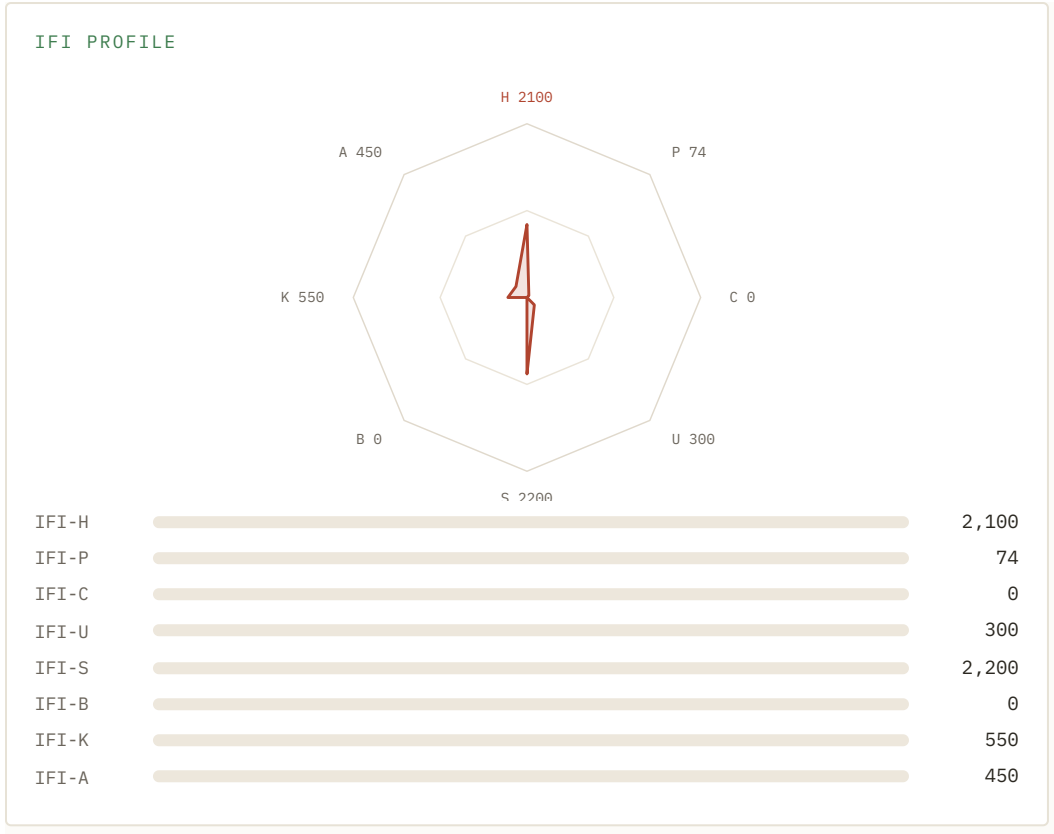

(Design target, modeled — the heat and sweetness anchors are the line's named coordinates; the supporting axes are modeled around them.) An **H of 2,100** against an **S of 2,200** is the point where the dessert genuinely fights back — the furthest our family would take a fruit before it stops being sweet-heat and becomes simply a chili wearing a strawberry's red. Which it never should. The P rises to 74 as the heat does; the K eases to 550 as the sweetness steps back; the A lifts to 450 as the whole fruit leans savory. It is still, by a designed margin, a *sweetness* — just one that answers for itself.

Precision as respect — CV and the ratio

Numbers this exact are only worth writing if the fruit can hit them, batch after batch, and this is where the platform's discipline earns its keep. A wild-type pepper's heat wanders with a coefficient of variation near **42%** — plant to plant, pod to pod, a swing so wide that two fruit off the same bush are not the same product. The designed lines carry a batch-to-batch CV specification of **≤7%.** On an IFI-H target of 1,400, seven percent is ±98 at one standard deviation; a ninety-five-percent window is roughly [1,204, 1,596]. The wild-type's forty-two percent is ±588 — a window from 812 to 1,988, so broad that *Warm Kiss* and *Candied Fang* would be statistically indistinguishable. You cannot choreograph with an orchestra that cannot find the note. Tight CV is not a manufacturing nicety; it is the precondition for designing a sequence at all.

And an intersection product spends that precision twice, because its character lives in a *ratio.* The sequence depends on the balance of heat to sweetness, R = IFI-H ⁄ IFI-S; for *Candied Fang*, R = 1,400 ⁄ 2,400 ≈ 0.58. When two independently varying quantities divide, their relative errors add in quadrature:

$$CV(R) = \sqrt{(CV_H^2 + CV_S^2)} = \sqrt{(0.07^2 + 0.07^2)} \approx \mathbf{0.099.}$$

So even with each axis held to seven percent, the ratio that *is* the dessert carries about ten percent — R = 0.58 ± 0.06. That is the real, honest cost of composing two axes into one experience, and it is exactly why a sweet-heat fruit demands the full ≤7%-per-axis rigor

and would be impossible without it. The choreography is only as reliable as the two dancers are precise, and ten percent on the ratio is the margin inside which the kiss still reliably comes first.

Failure modes, and the honest limits

We would rather name the ways this fruit can fail than pretend it has none. **Heat before sweet** is the cardinal failure: get the ratio or the spatial release wrong and the fire arrives with the sugar instead of after it, and a strawberry that bites before it kisses is a failure however hot — the whole art of this frontier is restraint, and restraint is a thing you can measure missing. **The capsaicinoid muddle** is the next: an acyl module imported out of balance with the amine supply yields a smear of minor capsaicinoids and a burn with no clean shape. **Flesh leakage** is the third: targeting that drifts off the achene surface loses the sequence and contaminates the sweetness at the source. **Color loss** is the quiet one: lean too hard on the shared phenylpropanoid branch and the berry pales as its anthocyanins starve for the carbon we diverted to vanillin — a fruit that trades its own red and its own antioxidants for a trick is a fruit we will not ship. And there is simple **off-flavor:** vanillylamine and its intermediates are not neutral, and an over-expressed amine head can read medicinal before any heat arrives.

The honest limits are plainer still. These four lines are design concepts, not cultivars — seven to ten genes out, thirty-six to forty-eight months of work, every IFI value here a modeled target rather than a measured result on a finished fruit. The heat chemistry is fully established in our designs, residue by residue, on the pepper side; the sweetness restoration is built out across this book; the transplant of the two into a berry, in balance, at the surface, in time, is the part still ahead of us. We write it out in this detail not because the strawberry is on a shelf, but because the path to it is real enough to measure, and measuring it honestly is the only way we know to earn it. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

What stays fixed through all of it is the kiss. The heat must always arrive as the *second* act, inside a sweetness that was generous first — a pleasure offered to the body with intention, no added sugar underneath it, the berry's own color and tannins and ascorbate left whole, the fire a plant-identical molecule and nothing animal or synthetic anywhere in the design. A strawberry that bites before it kisses is a failure, however hot. So we design the kiss to come first, and we build the whole fruit — surface, sequence, and ratio — to keep that promise every single bite.

Chapter 14 · Solar, Forbidden, Celestial — *the dessert frontier.*

The sweetest work our family does is also the most counterintuitive, because the fruit that tastes of sugar is almost never short of sugar. Across the roadmap's cross-family and extended-nightshade tiers sits a small constellation of concept cultivars whose shared project is to push sweetness — sometimes alone, sometimes with a single companion axis — into fruits that modern agriculture flattened or never troubled to optimize. Three of them map the entire range of the problem, because each carries a *different* obstacle to sweetness, and the obstacle, not the sugar, is where the engineering lives. **Solar** is a melon, and its obstacle is aroma. **Forbidden** is an apple, and its obstacle is acid. **Celestial** is a ground cherry, and its obstacle is almost nothing at all — a nightshade so close to the pepper in its chemistry that warmth arrives as a nudge rather than a transplant.

To design any of it, we first had to stop treating sweetness as a quantity of sugar and start treating it as what it physically is: a receptor, firing, and a brain deciding how loudly the firing counts. Everything in this chapter follows from that one correction. Name the obstacle, count the genes, set the target — that is the frontier, fruit by fruit. The growing is just time.

The receptor that reads sweetness

Sweetness begins at a single molecular machine, the heterodimer **TAS1R2·TAS1R3**. Each half is a class C G-protein-coupled receptor, and each is built in three parts: a large, bilobed extracellular **Venus flytrap domain** that closes like a clamshell when it catches a ligand; a short **cysteine-rich domain** that transmits that closure inward; and a **seven-transmembrane** bundle threaded through the cell wall that, once the clamshell shuts, changes shape on the inside of the cell. Sugars — sucrose, glucose, fructose — settle mostly into the flytrap cleft of the TAS1R2 lobe. The sweet proteins and a few synthetic sweeteners dock instead against the cysteine-rich domain or the transmembrane bundle of TAS1R3. One receptor, several doors: this is why "sweet" is not one sensation but a family of them, and why the same tongue can be fooled, enhanced, or suppressed by molecules that never touch the sugar site at all.

What happens after the clamshell closes is a cascade we can follow step by step, and every step is a place a flavor engineer can intervene. The conformational change on the cytoplasmic face activates **α-gustducin**, the taste cell's dedicated G protein. Gustducin frees its βγ subunits, which switch on **phospholipase C-β2**. PLC-β2 cleaves the membrane lipid PIP_2 into two messengers; one of them, **IP_3**, opens calcium stores on the endoplasmic reticulum. The flood of intracellular calcium gates **TRPM5**, a sodium-admitting channel, and the cell depolarizes. Depolarization opens **CALHM1/CALHM3** channels, through which the cell releases ATP

directly — no vesicle, no synaptic bag — onto the **P2X2/P2X3** purinergic endings of the gustatory nerve. The nerve carries it to the solitary tract and up to the cortex, where the brain performs the final, decisive act: it weighs the signal against everything else arriving at once — the smell climbing the back of the throat, the sourness, the texture — and only then reports a sweetness.

The several doors matter because they can be worked independently of one another. Lactisole and gymnemic acid shut sweetness down by binding the transmembrane bundle of TAS1R3, far from any sugar site at all — proof that the receptor carries allosteric volume controls, not merely an on-switch. Positive allosteric modulators do the reverse, nudging the clamshell toward closure so that the very same sugar is read louder. We keep every one of these levers in view and, for a dessert fruit, deliberately leave most of them untouched: the fruit should taste of its own sugar, resolved and clean, not of a modulator's borrowed signal arriving on a schedule the fruit never sets. The one exception we do reach for is invisible on the tongue — the removal of a suppressor the fruit manufactures against itself, which is subtraction rather than addition, and which the apple will put to work.

That last clause is the whole reason this chapter exists. **Perceived sweetness is receptor activation times cortical gain**, and both terms are tunable without a single added gram of sugar. You can raise the first term with a more potent agonist or by removing a molecule that blocks the door. You can raise the second term by feeding the brain a congruent aroma, or by removing a competing sensation — sourness, chiefly — that the cortex uses to turn the sweetness down. The sweetness axis is defined on the Index precisely as *sugar-free sweetness enhancement*, and this is why the definition is literal rather than aspirational: the receptor does not measure sugar. It measures activation, and activation is cheaper to buy than sugar is.

Sweetness in sugar's own units

The Iyengar Flavor Index turns that activation into a number. For every axis,

> $\mathrm{IFI(candidate)} = 1000 \times \mathrm{EC_{50}(reference)} \div \mathrm{EC_{50}(candidate)}.$

$\mathrm{EC_{50}}$ is the concentration at which a stimulus drives its receptor to half of full response; invert it against a fixed reference and scale by a thousand, and loudness becomes a score. For sweetness we anchor the reference to pure sucrose, so that the sucrose standard reads **IFI-S = 1,000** by construction. Everything sweeter than sucrose has a *lower* $\mathrm{EC_{50}}$ and climbs above 1,000; everything blander sits below it.

Worked against our bench reference of $\mathrm{EC_{50}}$(sucrose) ≈ 30 mM, the arithmetic is immediate and the dynamic range is enormous:

SWEET PRINCIPLE	SOURCE ORGANISM	~× SUCROSE	EC_{50} (BENCH)	IFI-S = $1000 \cdot EC_{50}(REF)/EC_{50}$
Sucrose *(reference)*	Sugarcane · 200 constructs	1.0	30 mM	**1,000**
Glucose	fruit flesh	0.74	40.5 mM	741
Fructose	fruit flesh	1.5	20 mM	1,500
Steviol glycoside	Stevia · 910 constructs	~200	0.15 mM	200,000
Mogroside V	Monk Fruit · 80 constructs	~250	0.12 mM	250,000

Two honest footnotes keep this table from lying. First, the scorecard radar normalizes every axis to a ceiling of 5,000, so a steviol glycoside's raw 200,000 pins to the top of the chart and tells you nothing a dessert designer can use — which is exactly why the dessert frontier reaches for none of these high-potency molecules. Their sweetness is real but *shaped* wrong: it arrives late, lingers, and refuses to resolve the way a fruit's sugar does. Second, the unit conversion is worth doing once by hand, because it reveals the trap that ordinary breeding keeps falling into. Thirty millimolar sucrose, at a molar mass of 342.3 g/mol, is 0.030 × 342.3 = 10.3 g/L, which is about **1.0 g per 100 mL — roughly one degree Brix.** The sweet receptor reaches half of its full output at about one degree Brix of pure sugar.

A ripe melon is ten to twelve degrees Brix. An apple is eleven to thirteen. These fruits are not near the receptor's half-point; they are an order of magnitude past it, far out on the flat shoulder of the dose–response curve. Model that curve with a Hill function, R/R_max = c^n / (EC_{50}^n + c^n), and take a gentle cooperativity of n ≈ 1.2. At ten times EC_{50} — a ten-Brix fruit — the receptor is already $10^{1.2}$ / (1 + $10^{1.2}$) = **94% saturated.** Now do the thing breeders have done for a century and *double the sugar* to twenty times EC_{50}: activation climbs to $20^{1.2}$ / (1 + $20^{1.2}$) = **97%.** A hundred percent more sugar buys three percentage points of receptor output. That is the mathematical obituary for "just make it sweeter by raising Brix." In ripe fruit, sugar is spent currency. The sweetness that is still for sale lives entirely in the second term — aroma and the lifting of suppression — and that is the terrain all three of these fruits are built to work.

The formula runs backward just as cleanly, which is how a flavor target becomes a bench specification. A design target of IFI-S = 1,600 — Celestial's — implies an effective EC_{50} of 30 mM × 1000 ÷ 1600 = **18.75 mM**, a sweet principle reaching half-maximal activation at about 0.64 g per 100 mL of sucrose-equivalent; Solar's 2,600 implies 11.5 mM, and Forbidden's released 1,900 implies 15.8 mM. Those are the numbers the assay must hit — not the Brix the fruit must reach, and the gap between the two is the whole thesis of the dessert frontier written as a pair of concentrations.

The assay uses a cell line expressing the human TAS1R2·TAS1R3 heterodimer behind a calcium-flux reporter. A dose series of the fruit's standardized soluble extract is read concentration by concentration, each point taken as a fraction of that extract's own maximum response, and the series is fit to the Hill equation to recover both EC_{50} and the cooperativity n. Running every candidate against the identical sucrose reference, on the same plate, on the

same day, is what makes a melon, an apple, and a chili comparable at all. The coefficient of variation we specify at or below 7% is tracked here, batch against batch — and it is a form of respect for the reader as much as for the formulator, because a sweetness you cannot reproduce is a sweetness you cannot honestly name.

It is worth saying what the loud end of the Index looks like, if only to measure the quiet of the dessert fruits against it. Our full-spectrum flagship, Phoenix, is the one cultivar that registers on all eight axes at once:

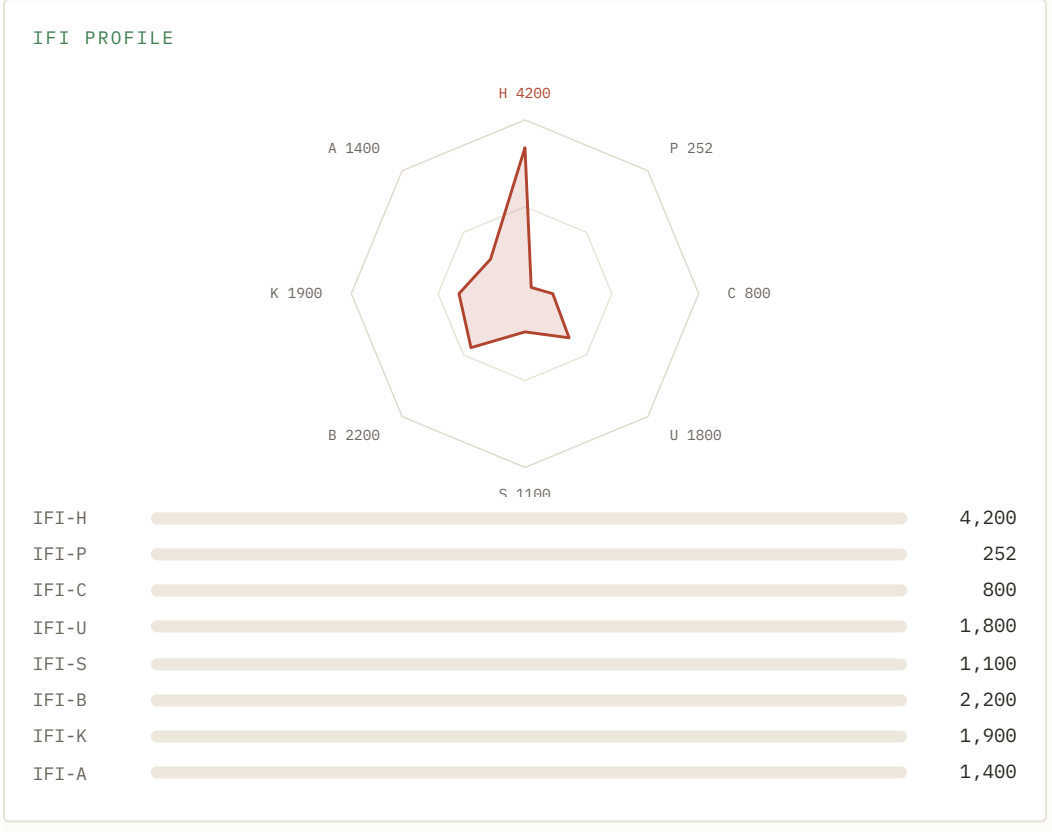

Our model predicts approximately 2.95 million Scoville heat units of engineered warmth carried alongside seven other sensations — a chord with every note struck. The dessert frontier is its deliberate opposite: a single sustained note, or two, with the other seven held near silence. Designing *for* silence on six axes turns out to be as exacting as designing for sound, and it is where the family's vegetarian, clean-handed instinct earns its keep — we are not adding; we are mostly giving back what breeding quietly took.

Solar — the melon, and the problem of smell



Melon — aromatic sweetness, where most of what you taste as sweet is actually smell. Solar is a melon with its lost perfume restored and amplified.

More than almost any fruit, a melon's sweetness is carried by its perfume. Cut a ripe *Cucumis melo* and the sweetness reaches you before the sugar does — a wave of esters and a cool green top note

that the brain files as "sweet" a full second before the tongue confirms it. This is not metaphor; it is retronasal olfaction doing arithmetic. As you chew, volatiles climb from the back of the mouth to the olfactory epithelium, and the cortex binds that smell to the taste signal arriving at the same instant. When the two are **congruent** — a sweet-associated smell meeting a sweet taste — the cortex adds gain. A melon that has lost its perfume, as most shipped and stored melons have, tastes watery no matter its Brix, because the second term has gone to zero and the sugar is stranded on the flat shoulder of the curve.

The perfume has a precise biochemistry, which is what turns Solar from a wish into a brief. The dominant notes are esters — ethyl and methyl acetates, propanoates, and butanoates — assembled by **alcohol acyltransferases** (the *CmAAT1–4* family) that condense the fruit's own alcohols with acyl-CoA donors during climacteric, ethylene-driven ripening. Beneath them sit the fresh, faintly cooling nine-carbon aldehydes and alcohols — the "melon-and-cucumber" signature — made by the **lipoxygenase / hydroperoxide-lyase** branch cleaving the fruit's linoleic and linolenic acids. A melon goes watery when storage and chilling silence the acyltransferases; the sugars survive the truck, but the perfume does not. Solar's engineering is therefore restoration, the same volatile-rescue method that brought a supermarket tomato back to tasting like a tomato, pushed to its limit: sustain *CmAAT* expression through ripening and storage, reinforce the LOX/HPL branch, and let the nose do the sweetening the sugar cannot.

The lift is calculable. Controlled aroma-enhancement studies put the sweetness gain from a fully congruent volatile complement at the equivalent of roughly **+3 degrees Brix** on a ten-Brix base — perceived, not a direct measurement of sugar content. Converting a perceived ratio to receptor-relevant loudness through Stevens' power law for sucrose, $\Psi \propto C^{1.3}$, a 13/10 Brix-equivalent ratio becomes $(1.3)^{1.3} \approx$ **1.41× perceived sweetness at identical measured sugar.** An unaromatic melon reading IFI-S ≈ 900 is lifted past 1,270 by perfume alone; sustained across the whole ripening window and crowned with its own cool top note, the design target sits higher still, near the restored-tomato magnitude:

IFI PROFILE

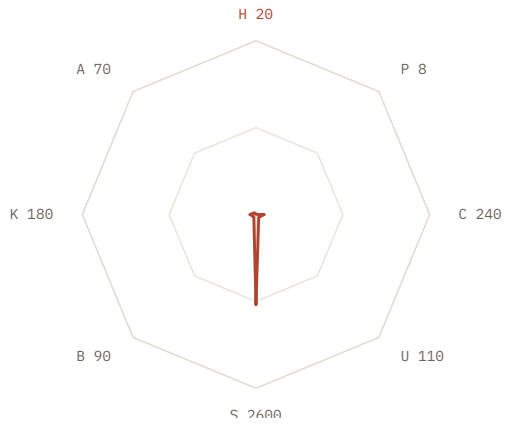

IFI-H		20
IFI-P		8
IFI-C		240
IFI-U		110
IFI-S		2,600
IFI-B		90
IFI-K		180
IFI-A		70

(A design target, not a measurement — we do not grow Solar yet.) Note the shape: a towering S, a faint C from the genuine coolness of those nine-carbon aldehydes, and near-silence everywhere else. This is the frontier at its most innocent — a pure, perfumed sweetness, the least ambiguous fruit we have drawn. Its difficulty is not subtlety but **persistence**: esters are volatile by definition, and the honest failure mode is a melon that smells glorious at the bench and flat after a week in a box. Solar is, at heart, a problem of keeping the perfume switched on.

Forbidden — the apple, and the problem of acid



Apple — a sweetness throttled by its own acid. Forbidden turns the malic acid down and lets the sugar already there be heard.

An apple is a negotiation between sugar and **malic acid**, and the acid usually wins the argument over how sweet the fruit is allowed to seem. Sourness actively suppresses perceived sweetness — a fact of the tongue's own wiring, not a matter of taste. Protons are read by a separate population of taste cells through the **OTOP1** proton channel; their signal climbs the same nerves and reaches the same cortex, where sweet and sour enter a mutual gain-control that turns each other down. Add acid to a sugar solution and the measured sugar does not change, but the reported sweetness falls. An apple, in other words, throttles its own sweetness, and the throttle is a dial we can turn.

Apple genetics hands us that dial cleanly, which is the whole reason Forbidden is a subtraction rather than an addition. Flesh acidity is governed substantially by the **Ma locus** — *MdALMT9 (Ma1)*, an aluminum-activated malate-transporter-like protein that loads malate into the vacuole — with support from the tonoplast proton pumps (V-ATPase, V-PPase) that set the vacuolar pH malate needs to accumulate. A natural low-malate allele already exists in the apple germplasm; editing toward it, or softly downregulating *MdALMT9*, lowers vacuolar malate without touching a single sugar gene. The sugar that was always there simply gets louder as the suppression lifts.

Here is the lever quantified. Model the suppression as a central gain factor σ = 1 / (1 + [malate]/K_i), with a suppression constant near K_i ≈ 0.30% malic acid, on a twelve-Brix apple:

MALIC ACID	SUGAR–ACID RATIO	Σ = 1/(1 + ACID/0.30)	PERCEIVED SWEETNESS (REL.)
0.70% *(wild, tart)*	17	0.300	1.00
0.50%	24	0.375	1.25
0.35% *(Forbidden)*	34	0.462	**1.54**

Our model predicts that halving the malate lifts perceived sweetness **1.54×** with not one gram of sugar added. Restore the aromatic esters on top of that — Forbidden keeps *MdAAT1* running for the same reason Solar does — and a further ≈1.26× congruence gain compounds in, for roughly **1.9× overall** in our model. Run that back through the power law and a twelve-Brix apple is designed to eat like a (1.9)^(1/1.3) × 12 ≈ **20-Brix** one. We then deliberately *stop short* of a candy apple: a measured tannic astringency is left in, the **A axis** carried by flesh and skin procyanidins binding salivary proline-rich proteins, so the fruit keeps structure and grip instead of collapsing into sweetness; and a barely-there warmth is written in as the name's wink — the apple of temptation should bite.

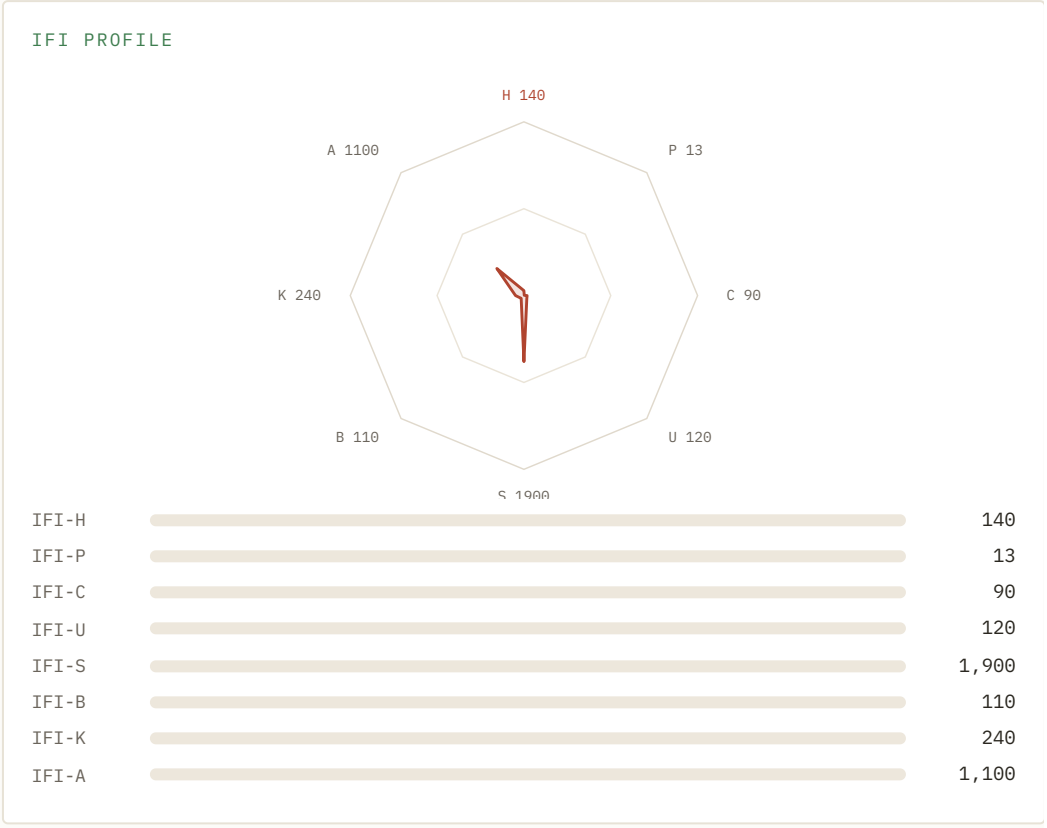

(A design target, not a measurement — Forbidden is not yet grown.) It is the frontier's most *composed* fruit: three axes held in balance so that sweetness, released from its acid, can finally be heard. The honest limit is that malate is not only flavor — it is also a part of the

fruit's storage physiology and defense, so the edit has a floor below which the apple keeps badly. Forbidden is a balance problem with a hard constraint at the bottom of the dial, which is precisely why it is a three-to-four-year build and not a season's.

Celestial — the ground cherry, and the gift of family



Ground cherry — a husk-wrapped nightshade, cousin to both tomato and pepper, already living in sweet-heat's neighborhood.

The ground cherry is the quiet favorite, because its advantage is ancestry. A *Physalis* is a husk-wrapped nightshade in the same family — the Solanaceae — as both the tomato and the pepper, with a naturally sweet-tart, pineapple-like flavor that already sits inside sweet-heat's neighborhood. That family resemblance is not sentiment; it is a shortened parts list. Every vascular plant already runs the upstream **phenylpropanoid branch, PAL → C4H → 4CL → HCT → COMT**, that feeds the pepper's heat. What makes a pepper a pepper is four terminal enzymes bolted onto the end of it: **pAMT**, which aminates vanillin to vanillylamine; **BCAT** and **KAS**, which build and elongate the branched-chain fatty acid; and **Pun1/CS** (the AT3 acyltransferase), which condenses the two into capsaicin in the fruit's placental tissue. Dropping those four into a distant host means fighting codon usage, regulatory context, and compartment chemistry at every step. Dropping them into a fellow nightshade means most of that fight is already won — the metabolic context is familiar, so the transplant reads more like a dialect than a foreign language.

The consequence is a whole tier of difficulty removed. Celestial is **extended-nightshade (T2) work — five to seven genes, twenty-four to thirty-six months** — against the cross-family strawberry's seven-to-ten genes and thirty-six-to-forty-eight-month horizon. It is the sweet-heat concept built the short way, and it is where we would *prove* sweet-heat before committing the longer, more glamorous berry.

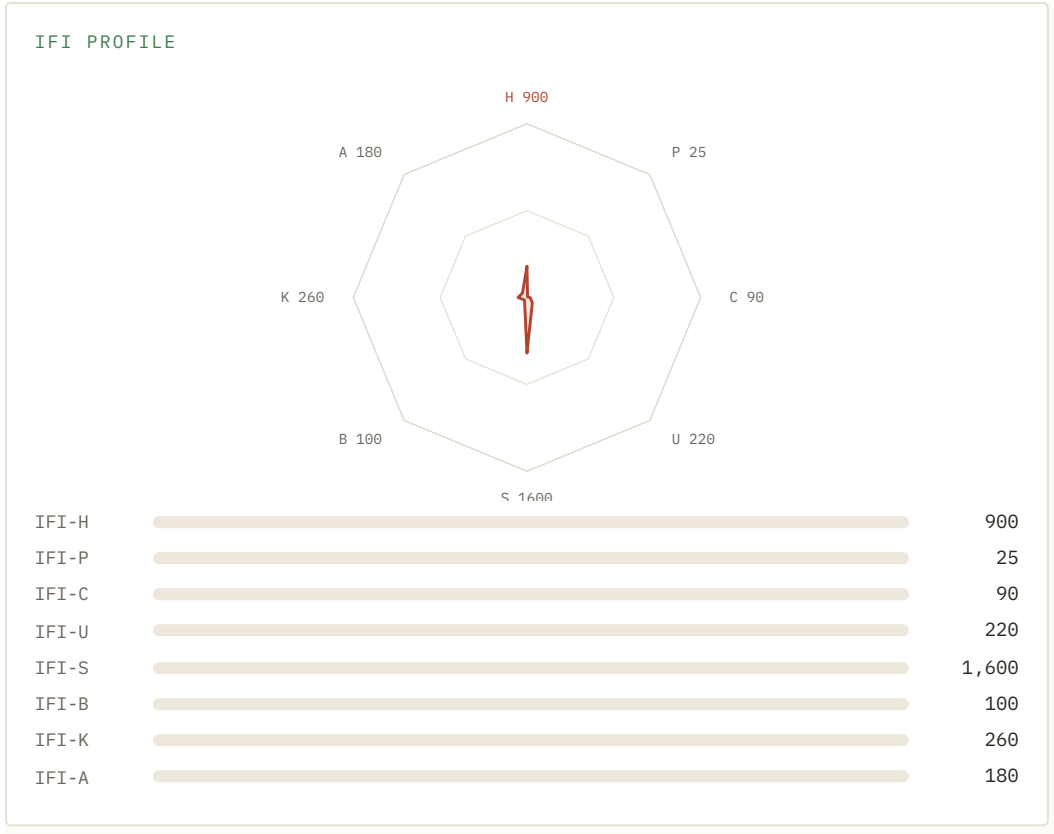

(A design target, not a measurement — Celestial is a concept cultivar.) The heat here is a courtesy, not an assault — nearer a fresh poblano than a flagship, an IFI-H of a few hundred rather than the 4,200 of our loudest pepper — set deliberately low because a dessert fruit wears warmth as a finishing note. Celestial carries its native sweet-tart sugar and acid almost untouched and simply *adds* the warmth that family makes cheap. Its failure modes are the gentlest of the three: landing the heat in the edible flesh rather than the placental seed core, and keeping the warmth from outrunning the fruit's modest sugar. If any of these three teaches the method fastest, it is this one — which is exactly why it goes first.

The construct brief, and the honest limits

Each fruit's obstacle has a known biochemistry, which is what lets us write a construct brief instead of a wish list. Set side by side, the three draw the shape of the whole frontier:

CONCEPT	HOST	TIER · GENES · TIME	CORE EDITS	LEAD AXES	OBSTACLE SOLVED
Solar	Melon (*Cucumis melo*)	T3 · 7–10 · 36–48 mo	sustain *CmAAT1–4* esters; reinforce LOX/HPL C_9 branch	S, (C)	lost aroma → stranded sugar
Forbidden	Apple (*Malus domestica*)	T3 · 7–10 · 36–48 mo	downregulate *MdALMT9/Ma1*; tune V-ATPase; keep *MdAAT1*; retain procyanidins	S, A, (H)	acid suppressing sweetness
Celestial	Ground cherry (*Physalis*)	T2 · 5–7 · 24–36 mo	transplant pAMT · BCAT · KAS · Pun1/CS onto native phenylpropanoid backbone	S, H	– (family makes heat easy)

There is a generalization hiding in these three rows worth making explicit, because it governs the whole dessert frontier. In nearly every ripe sweet fruit, the receptor is already saturated on sugar, so the large share of *perceived* sweetness that is still in play is carried not by sugar but by **aroma and acid balance.** The melon proves it by losing its perfume; the apple proves it by throttling itself with malate;

and the arithmetic proves it coldly — doubling the sugar of a ten-Brix fruit buys three percentage points of receptor output, while restoring a perfume buys forty, and lifting an acid suppression buys fifty. The fastest route to a sweeter fruit is almost never more sugar. It is giving back the volatiles breeding and shipping took, or rebalancing the acid that was masking what the fruit already had.

That is also why these cultivars cluster where they do on the roadmap — seven-to-ten genes for the cross-family reaches, five-to-seven for the nightshade-adjacent ones. Most of the work is not installing sweetness but **restoring the perception** of a sweetness the fruit could always carry, plus, in the sweet-heat lines, the pepper's borrowed warmth. The gene counts are honest estimates of real transplants, and the timelines — three to four years for the cross-family fruits — are the plain time it takes to build and prove a multi-gene construct in a slow-breeding perennial, generation by generation.

The limits are worth stating as plainly as the promise. The batch-to-batch coefficient of variation is designed to hold at or below **7%** in the cultivars we grow — tight enough that an IFI-S target of 2,600 carries a one-sigma band of roughly ±182, so better than nineteen in twenty fruit fall inside 2,600 ± 364. A supermarket melon, by contrast, swings its Brix by a quarter or more from one season and one truck to the next; consistency is itself part of what we are engineering. But precision at the bench is not the same as a fruit in hand, and we will not pretend otherwise: Solar's perfume must survive storage, Forbidden's malate has a floor below which the apple keeps badly, and Celestial's warmth must land in the flesh and not the seed. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. **We do not grow them yet.** What we do have, fruit by fruit, is the obstacle named, the genes counted, and the target set — and that knowledge, held with reverence for the living thing it describes, is itself the frontier. Sweetness is never the same problem twice: in the melon a problem of smell, in the apple a problem of acid, in the ground cherry barely a problem at all. The growing, as ever, is just time.

Chapter 15 · Phoenix — *sweetness inside the full spectrum.*

We end the frontier with the one cultivar in this book that is not a concept — that exists, that has a SKU, that is our family's most complete design to date. **Phoenix, SS-0096**, is the library's full-spectrum flagship: a *Capsicum chinense* our family engineered to register on all eight channels of taste at once, the only cultivar we have ever carried to that completeness. By heat it is a monster — our model predicts approximately 2.95 million SHU, an IFI-H of 4,200, the loudest burn we have ever designed. And yet read its whole profile and find, sitting calmly inside the fire, a small and deliberate sweetness.

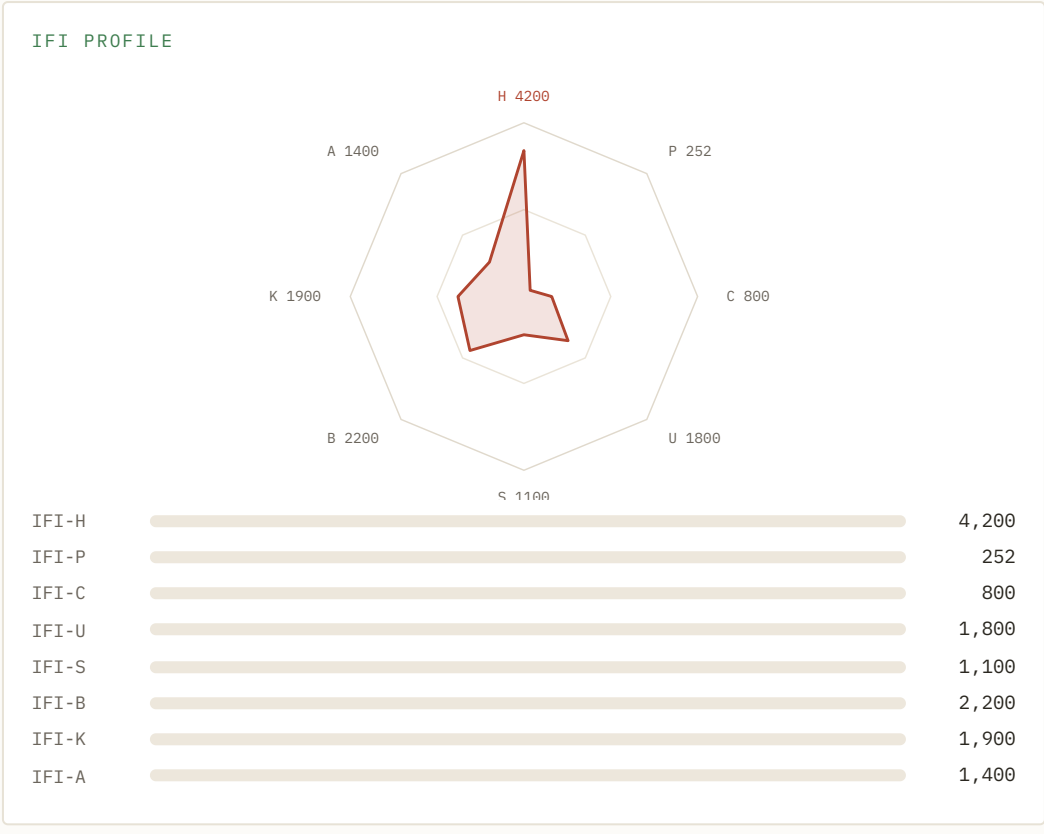

That **S of 1,100** is the quiet proof of everything these pages argue. It is a real, designed sweetness — just above the reference sugar's 1,000 — written into the hottest thing we make. Even here, at the furthest edge of heat, sweetness is not an accident or a leftover. It is a channel we tuned on purpose, to a number, because a full-spectrum design that forgot sweetness would taste like an assault, and our family does not build assaults. We build compositions. The

sweetness in Phoenix is small on purpose — this is a chili, not a dessert — but it is *there* on purpose, rounding the heat's leading edge so the fire arrives on something soft rather than on nothing. Heat should be a pleasure given to the body with intention; the floor under it is part of that intention.



Phoenix — the one cultivar we engineered to speak all eight axes at once. Its sweetness is small and deliberate: a floor under the fiercest heat we make.

Eight locks, eight keys — why an octave is even possible

A chord of eight is conceivable at all because the eight receptors are genuinely separate instruments. Sweetness is read by the heterodimer **TAS1R2·TAS1R3**; heat by the ion channel **TRPV1**; the savory round depth of umami by **mGluR1** and the **TAS1R1·TAS1R3** pair; bitterness by **TAS2R38**; the continuity we call kokumi by the calcium-sensing receptor **CaSR**; cool by **TRPM8**; the sharp pungency of mustard and wasabi by **TRPA1**; and astringency — the drying pull of a tannin — not by a taste receptor at all but by the precipitation of **salivary proline-rich proteins**. Each is a different protein with its own lock, so in principle each axis is an independent dial, and the Index gives us a readout on every one at once.

AXIS	RECEPTOR	PHOENIX	WHAT IT READS
IFI-H · Heat	TRPV1	4,200	capsaicin heat and spice
IFI-S · Sweetness	TAS1R2·TAS1R3	1,100	sugar-free sweetness
IFI-U · Umami	mGluR1 / TAS1R1·TAS1R3	1,800	savory, brothy depth
IFI-B · Zero-Bitter	TAS2R38	2,200	bitterness attenuated
IFI-K · Kokumi	CaSR	1,900	richness, continuity, mouthfeel
IFI-C · Cool-Burn	TRPM8 / TRPV1	800	cooling / warming
IFI-P · Pungency	TRPA1	252	allium / wasabi sharpness
IFI-A · Astringency	salivary PRPs	1,400	tannin, polyphenol pull

That eight separate locks exist is what makes an eight-axis design *thinkable*. It is not yet what makes it *true* — the honesty of that comes later in this chapter, when we show how three kinds of coupling keep the eight dials from ever being eight easy knobs. But begin with the cleanest fact: because the sweet receptor is its own protein, we can go after a sweetness number without first touching the heat, and read back whether we hit it.

The ratio that makes loudness a number

Every axis on that scorecard is the same kind of quantity, and it is worth deriving exactly what kind, because the whole book turns on it. We do not score how *much* of a compound a food holds. We score how *loudly* that compound speaks to its receptor — and loudness, for a receptor, is set by a single pharmacological number: the **EC_{50}**, the concentration at which a stimulus drives its receptor to half of full response. A potent ligand reaches half-activation at a low concentration; a weak one needs a flood. So loudness is *inverse* to EC_{50}, and the Index is built on exactly that inversion, pinned to a fixed reference and scaled by a thousand:

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

Anchor the reference so that its own score is 1,000 — feed its own EC_{50} into both slots and the ratio is one, times a thousand. Each reference constant is a stated convention — defined, the way Scoville's dilution threshold was defined, to give the axis a fixed starting point. For the pungency axis (IFI-P) the convention constant is AITC at 2.1 µM, so a compound matching that potency scores exactly 1,000. For the sweet axis the reference is table sugar, a weak agonist whose half-maximal concentration on TAS1R2·TAS1R3 sits in the tens of millimolar. Now read Phoenix's sweetness back through the formula and the number stops being a mood and becomes a designed target:

$$1{,}100 = 1{,}000 \times \frac{\mathrm{EC}_{50}(\mathrm{ref})}{\mathrm{EC}_{50}(\mathrm{Phoenix})} \;\Rightarrow\; \mathrm{EC}_{50}(\mathrm{Phoenix}) = \frac{1{,}000}{1{,}100}\,\mathrm{EC}_{50}(\mathrm{ref}) = 0.909\,\mathrm{EC}_{50}(\mathrm{ref})$$

Phoenix's sweet ligand is designed to reach half-activation at **90.9%** of the reference sugar's concentration — 9.1% more potent than sucrose. Fix the reference's half-max at one *reference-unit* (RU) and the scale reads as a table of potencies: IFI-S 1,000 is 1 RU, Phoenix's 1,100 is 0.909 RU, the sweet axis ceiling of 5,000 would be 0.2 RU — a ligand five times as potent, reaching half-activation at a fifth of the concentration — and a faint 250 would be 4 RU, needing four times the sugar to say the same thing. Phoenix's sweetness, at 1,100, sits at **22%** of the sweet ceiling: present, resolved, deliberately far from loud.

This is also why *sugar-free* sweetness is not a contradiction. The IFI tracks the EC_{50} — the horizontal position of the dose–response curve, the concentration you need — not the sugar mass you carry. A high-potency sweet ligand (a steviol glycoside, a mogroside, a sweet protein like brazzein) shifts that curve left without adding caloric bulk, so the receptor is driven hard while almost no sugar is present. Loudness without load: a sweetness you can give the body cleanly.

The dose–response makes the arithmetic physical. A receptor's fractional activation follows a Hill curve,

$$\theta(C) = \frac{C^{n}}{C^{n} + \mathrm{EC}_{50}^{\,n}},$$

with the sweet receptor's cooperativity *n* near 1. Pour the reference and Phoenix's ligand at the same concentration — set *C* equal to the reference's EC_{50} (C = 1 RU) — and the reference, sitting exactly at its own half-max, is 50.0% activated, while Phoenix's ligand, with the lower EC_{50} of 0.909 RU, reaches θ = 1 ÷ (1 + 0.909) = **52.4%**. A small, real edge, and the ratio of the two activations recovers the 1,100-against-1,000 we started from. The Index anchors at the half-max precisely because that is where the curve is steepest and a shift in potency shows most honestly — it linearizes, at the one informative point, a response that is otherwise an S-shaped nonlinearity.

Inside the sweet lock — how the receptor reads a molecule

To lower an EC_{50} on purpose you have to know the shape of the lock. TAS1R2·TAS1R3 is a class C G-protein-coupled receptor, and class C proteins are not simple pockets — each subunit opens into a large extracellular **Venus flytrap domain**, two lobes hinged like a clam, joined through a short cysteine-rich linker to a seven-helix bundle buried in the cell membrane. A sweet molecule is read as a closure: a sugar or a glycoside settles into the cleft of the T1R2 flytrap, the two lobes fold shut over it, and that closing travels down the linker, tilts the transmembrane bundle, and releases the G-protein gustducin on the inside — one molecule, one latch, a signal the cell can count.

What makes a *sugar-free* sweetness possible is that the flytrap is not the receptor's only door. Small sugars and steviol glycosides bind in the flytrap cleft; some synthetic sweeteners and the bitter-sweet blockers act down in the transmembrane bundle; and the sweet proteins — brazzein, monellin — do not fit the cleft at all but press against the *outer* face of the flytrap like a wedge that holds the clam half-shut. Different molecules, different sites, the same closure. That multiplicity is the whole opening our method works through: a ligand that occupies a high-affinity site latches the flytrap at a concentration far below sucrose's, which is exactly a lower EC_{50}, which is exactly a higher IFI-S — loudness bought with affinity instead of mass. Phoenix's sweet ligand is designed to close that flytrap at 91% of table sugar's concentration not because there is more of it, but because it fits the lock a little better; the edit's work is to raise the amount of a well-fitting ligand the fruit will make, or to install one and let its affinity do the rest.

Reading 1,100 — the sweetness that survives

Here is the subtlety that turns full-spectrum design from arithmetic into craft: independence of the *receptors* is not independence of the *perception*. The sweet, bitter, and umami cells share downstream wiring and converge in the brain, and the axes lean on one another in the mouth — bitterness and acid press sweetness *down*; umami, kokumi, and salt lift it *up*. So Phoenix's sweetness does not sit in a clean room. It sits beside a bitterness of 2,200 that would flatten a weaker reading, and beside a kokumi of 1,900 and an umami of 1,800 that lengthen and buoy it.

Which means the 1,100 on the scorecard is not the sweetness we *installed* — it is the sweetness that *survives*. We did not tune the sweet receptor in isolation and copy the number down; we tuned it

against the whole competitive, cooperative field, so that what the instrument reads matches what the tongue feels. A naïve design that set the sweet ligand to a receptor-level 1,100 and then let a bitterness of 2,200 land on top of it would read, on the tongue, as something nearer 800 — the sweetness eaten by its neighbor. The 1,100 is the perceived number, solved for, with the suppression already paid. Model that balance in the open: if the local bitter field shaves roughly a quarter off a bare sweet reading — a representative fraction, shown here to give the solve its shape rather than as a fixed constant — then to net 1,100 on the tongue we must install a receptor-level sweetness near 1,100 ÷ 0.75 ≈ **1,470**, and then confirm that the lift from an umami of 1,800 and a kokumi of 1,900 does not carry it back past the target. The scorecard's 1,100 is the *output* of that negotiation, not its input.

That is the discipline the whole book has been teaching, carried to its hardest case. Name the number you want *on the tongue*. Find the receptor. Walk the pathway. Then model the neighbors — the ones that steal and the ones that give — and set the ligand so the survivor equals the target. In a strawberry, where sweetness leads and little stands against it, that solve is gentle. Inside Phoenix, where sweetness is one soft voice against the loudest heat we make and a wall of bitterness beside it, the same solve is the difference between a composition and a mess.

The heat it sits inside — two edits in Pun1

To understand the restraint of the sweetness you have to stand next to the fire it floors. Phoenix's heat is capsaicin, and capsaicin is a condensation: vanillylamine — delivered off the phenylpropanoid branch that runs in every vascular plant — joined to an 8-methyl branched-chain fatty acid at the **Pun1/AT3 acyltransferase**. The chemistry is an acyl transfer: the fatty acid arrives already activated as a coenzyme-A thioester, a high-energy bond the enzyme spends to drive the reaction forward, and Pun1 hands that acyl group onto the nitrogen of vanillylamine, forming the amide that is the molecule's burning heart and releasing free coenzyme A. Everything about the burn — how much, and the *shape* of it, the balance between capsaicin and its blunter cousin dihydrocapsaicin — flows through that final handoff.

Our flux model predicts that control in the phenylpropanoid pathway sits upstream — at 4CL and HCT/CCoAOMT — not at the terminal Pun1/AT3 step. Pun1 is the final capsaicin synthase; it labels the product, it does not govern the rate. The edits below optimize how efficiently the enzyme uses the substrate supply the upstream pathway delivers: a faster, better-folded Pun1 captures more of that upstream flux and converts it cleanly to capsaicin.

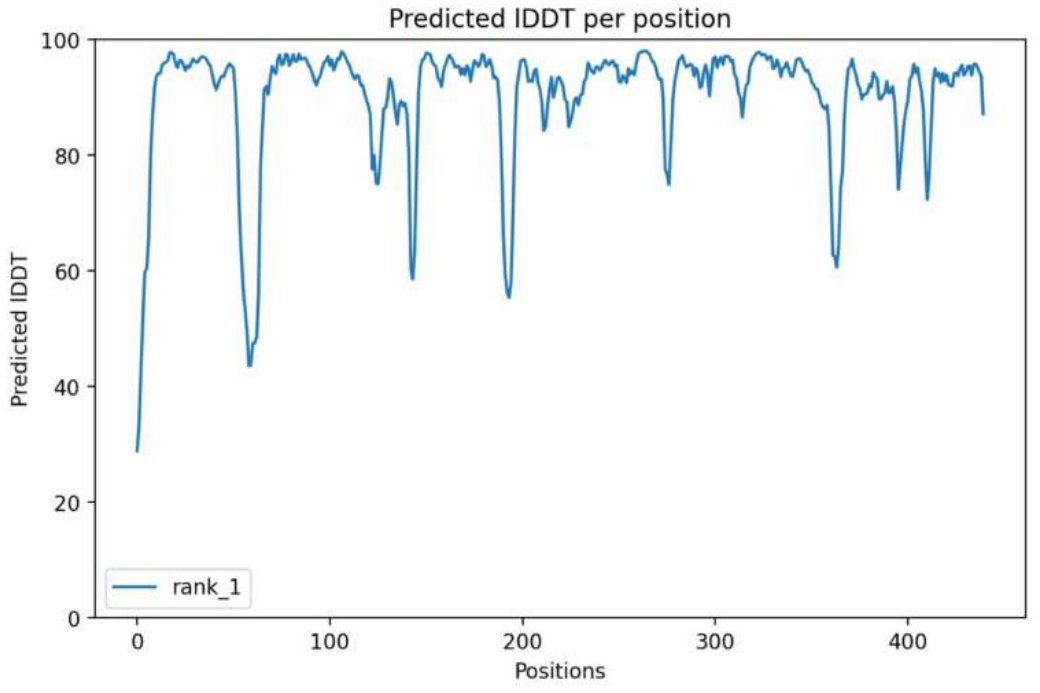

The Pun1/AT3 acyltransferase, coloured by model confidence — the binding domain and active-site cleft, where the two edits land, resolve most sharply.

We designed it with two substitutions no breeding program ever found, because no breeding program reads residues — it reads fruit and hopes the genotype followed. We inverted that. An ESM2 protein-language model scored every possible single-residue change in Pun1 by its predicted effect on catalytic efficiency and fold stability; the flux model confirmed the target; CRISPR wrote the winners in.

EDIT	TARGET	EFFECT	ΔΔG (KCAL/MOL)
S39L	Pun1 (N-terminus)	binding affinity ↑	−2.1
L345G	Pun1 (active-site cleft)	steric relief	−1.8
C175S *(structural annotation — not in shipped construct)*	Pun1 (fold)	disulfide-misfold prevention	−0.9

Read the ΔΔG column as the model reads it: each value is the predicted change in free energy, in kilocalories per mole, and each is negative, meaning each edit *stabilizes*. The two shipped edits — **S39L** tightens how the N-terminus holds its substrate, **L345G** carves a glycine's worth of room into the active-site cleft so the branched acyl chain seats without strain — sum naïvely to −2.1 + (−1.8) = **−3.9 kcal/mol** of stabilization, though we do not get to simply add them: mutations in one protein couple, and the true combined effect has to be evaluated in the model, not totalled. C175S (which would retire a cysteine whose stray disulfide misfolds a fraction of every batch) is a structural annotation from our modelling; it was pulled from the final design and is listed here for completeness only. What the two shipped edits together buy is not only substrate capture but *reliability* — a well-folded enzyme is a consistent enzyme.

The heat they are designed to deliver is our model's prediction of approximately **2.95 million SHU**. It is worth converting that out of its own units. The Scoville scale's defining equivalence sets pure capsaicin at 16,000,000 SHU, so Phoenix's predicted heat is, by that equivalence, the computed burn of a 2,937,948 ÷ 16,000,000 =

18.4% capsaicinoid loading — a concentration the plant cannot stage without optimizing the final enzymatic step to make full use of the upstream phenylpropanoid supply. On the receptor side the same heat reads as an IFI-H of 4,200, which is 4.2 times the reference's loudness and **84%** of the heat axis ceiling of 5,000. The two numbers do not scale together, and that gap between them is the subject of the next page.

The ladder — where Phoenix sits

Phoenix is the top rung of a ten-cultivar series, each step representing an increase in computational design scope and axis activation. Laying the series out shows why SHU and IFI-H are not the same metric wearing two coats.

CULTIVAR	SKU	PREDICTED SHU	IFI-H	AXES	SHU PER IFI-H POINT
Ember	SS-0087	866K	1,570	2/8	552
Cinder	SS-0088	1.05M	2,290	3/8	459
Kindle	SS-0089	1.25M	2,000	3/8	625
Flare	SS-0090	1.48M	—	4/8	–
Blaze	SS-0091	1.72M	2,630	4/8	654
Wildfire	SS-0092	1.96M	2,930	5/8	669
Scorch	SS-0093	2.21M	3,220	5/8	686
Firestorm	SS-0094	2.47M	3,490	6/8	708
Solara	SS-0095	2.72M	3,980	6/8	683
Phoenix	**SS-0096**	**2.95M**	**4,200**	**8/8**	**700**

Watch the last column. Low on the ladder, a Scoville unit buys a lot of loudness — Cinder spends about 459 SHU per IFI-H point. High on the ladder the same point costs more — Firestorm pays 708, Phoenix pays 700. The price of loudness climbs as the heat climbs, and that is the fingerprint of a receptor approaching saturation: once TRPV1 is already firing hard, each additional molecule of capsaicin shifts the fractional activation less, so a huge increment of SHU moves the IFI only a little. SHU measures the *concentration* in the fruit, which grows without limit; IFI-H measures the receptor's *answer*, which is bounded by one and compresses near the top. Phoenix's predicted heat reaches 2.95 million SHU past plain warmth; its IFI-H is a disciplined 4.2 times the reference. The Index tells you what the tongue can actually distinguish, which is why we design against it and not against the raw Scoville number.

Error is a flavor you can taste — the 7% band

A number you cannot reproduce is not a design target; it is an anecdote. The hardest, least romantic part of designing flavor is making the same flavor twice, and it is where the engineering earns the poetry. A wild-type pepper's heat varies batch to batch at a coefficient of variation near **42%** — plant to plant, the same cultivar can nearly halve or half-again its burn. Our constructs are designed to a CV of **≤7%**, and that tightening is as much the product as the mean.



Phoenix in the controlled chamber — where the same light, the same water, and the same genome are held steady enough that the eighth decimal of a flavor comes back the same.

Put the 7% to work across the whole profile. The standard deviation on any axis is just 0.07 times its score, and from it the one- and two-sigma bands fall straight out — two sigma being roughly the 95% window a single grow should land inside.

AXIS	IFI	Σ (7%)	−1Σ	+1Σ	95% BAND (±2Σ)
H	4,200	294	3,906	4,494	3,612 – 4,788
P	252	18	234	270	216 – 288
C	800	56	744	856	688 – 912
U	1,800	126	1,674	1,926	1,548 – 2,052
S	1,100	77	1,023	1,177	946 – 1,254
B	2,200	154	2,046	2,354	1,892 – 2,508
K	1,900	133	1,767	2,033	1,634 – 2,166
A	1,400	98	1,302	1,498	1,204 – 1,596

Hold the whole cultivar at the lower edge of that band — every axis at −1σ — and it still reads as itself:

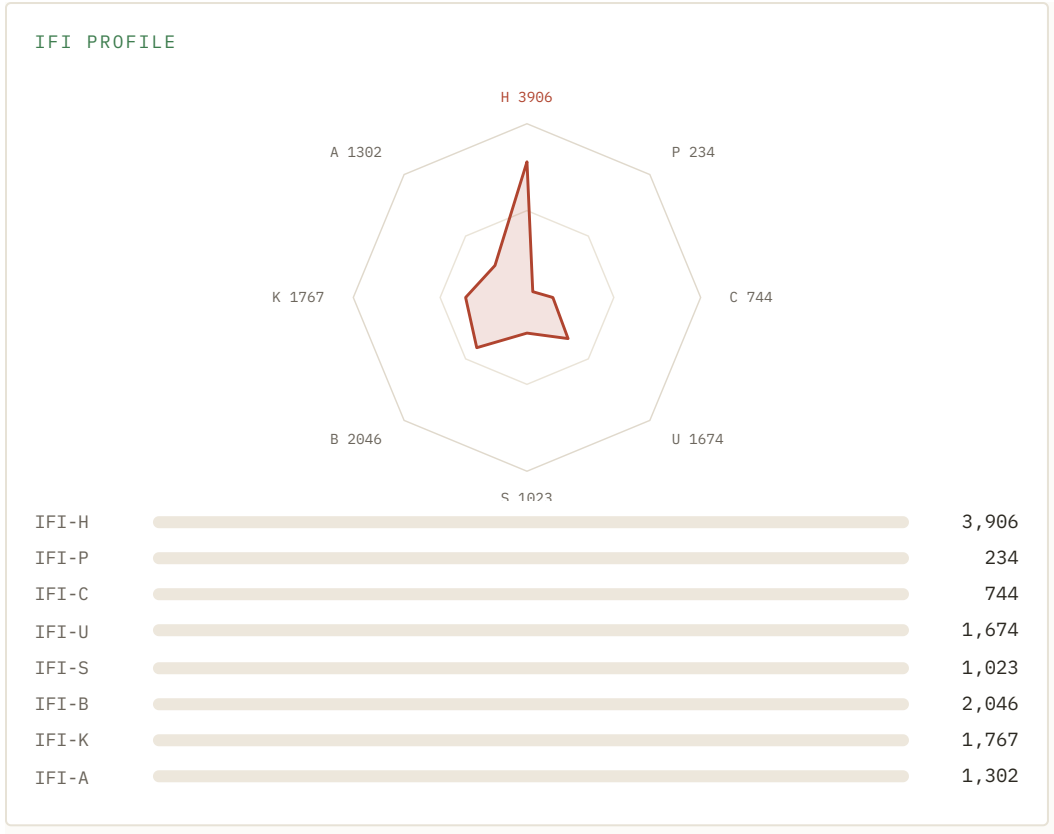

Look at the sweetness there. At −1σ, S is 1,023 — still above the reference sugar's 1,000. The designed sweetness clears its reference even on a weak grow, by a sliver. But now the honest edge: push to the full 95% window and the low end of sweetness is 946, which is *below* 1,000. Two standard deviations down, we cannot promise, from a single plant, that the sweetness beats plain sugar. That is a real limit, and it is exactly what replicate measurement is for. Read *n* grows and the uncertainty on the mean shrinks as the square root of *n*: the standard error is σ ÷ √n, so ten assays bring the sweetness estimate to 1,100 ± (1.96 × 77 ÷ √10) = 1,100 ± 48 = **[1,052, 1,148]**, cleanly and provably above the reference. One assay alone carries a 95% interval of 1,100 ± 151 = [949, 1,251], which straddles the line — a single taste cannot settle it. And against a wild-type 42% CV, a lone reading of a nominal 1,100 would span [195, 2,005]: not a design target at all. The tight CV is why the sweetness is *knowable*; the replicates are how we know it; the chamber in the photograph is where both come from.

Order is set by delivery — the nanoparticle phrase

Loudness is only half of a flavor. Eight axes firing at once, in a single undifferentiated rush, would read as chaos rather than composition — the chord struck with a fist. So a full-spectrum product is formulated with engineered **nanoparticle carriers** that hold and release the axes in time: the sweetness and its floor landing first, the heat built to rise behind it, the long tail of bitterness and astringency arriving last. Each carrier is tuned to let its cargo go at a different moment — a staggered release measured in seconds across the arc of a single mouthful — so that eight simultaneous facts become eight events in a line. The chord becomes a phrase played in order. The edits set *how loud* each voice is; the delivery sets *when* it enters — and a flavor is both at once. This is why a full-spectrum cultivar is more than its genome: the same 1,100 of sweetness, released first, reads as a welcome; released last, it would read as an aftertaste. We place it at the front on purpose, the soft thing the fire lands on.

The limits we keep honest

Eight separate proteins are eight separate locks, but three kinds of coupling keep them from being eight easy knobs, and naming them is how a design stays honest.

First, **shared hardware**. TAS1R3 is not the sweet receptor's alone — it is one half of the sweet heterodimer *and* one half of the umami heterodimer, and the sweet, bitter, and umami cells run through overlapping downstream machinery. A heavy hand on the shared subunit can nudge a neighbor, so the sweet and umami channels of Phoenix — 1,100 and 1,800 — were never designed in separate rooms. Second, **perceptual interaction**, the suppression and lift already described: the 1,100 is a survivor, not an install, and the number only means what it says because the neighbors were in the model. Third, **metabolic competition**. The eight pathways draw on shared pools of carbon and energy inside one living fruit, so the flux-balance model does not optimize any single axis — it apportions a finite budget across all eight demands at once. Pushing heat to 4,200 is also deciding what sweetness, umami, and kokumi can be afforded beside it.

So "eight independent instruments" is true at the receptor and only partly true at the plant. Designing Phoenix meant solving those three couplings together rather than tuning eight dials in turn — a harder problem than any single axis alone, and the long study of how axes compose is its own volume. The honest limits are plain: a single grow cannot, at two sigma, guarantee the sweetness clears its reference; the ΔΔG edits do not simply add; and the full profile is a negotiated budget, not eight free maxima. We publish the limits because numbers offered without their error are a kind of flattery, and we would rather respect the reader. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

One method, two extremes

Phoenix closes the frontier because it inverts the book. Everywhere else in these pages, sweetness has been the subject — loud, clean, freed from sugar, the thing we were chasing. Here it is a supporting player: one soft voice in eight, small and exact, a 1,100 you could miss against a 4,200. And it turns out to be the same craft either way. Whether we are growing a fruit that is almost all sweetness or a superhot that carries only a trace, the discipline does not change — name the number you want on the tongue, find the receptor, walk the pathway, balance the neighbors, hold the batch to 7%, evaluate the result, and release the axes in their proper order. A sweetness built to lead and a sweetness held to a floor inside the fire are one method, pointed two ways. That a trace of designed sweetness sits calmly at the center of the hottest, most complete cultivar our family has ever built is the proof this chapter exists to give: sweetness, designed and placed with intention, is not a thing apart from the rest of flavor. It is one instrument in the whole orchestra — and the method that tunes it to lead a strawberry is the very method that tunes it to 1,100 inside the flame.



PART V OF V

Reformulation & the Kitchen

the playbook · the formulations · sweetness, given back

Chapter 16 · The reformulation playbook — *from a sugar brief to a number.*

Most of the people who will ever open this book do not grow a single plant. They reformulate — a soda shedding a third of its sugar to meet a new standard, a cereal going clean-label, a sauce asked to lose its corn syrup without a soul at the breakfast table noticing. We wrote this chapter for them, at our own bench, the way we would hand a recipe across a warm kitchen: here is the method, movement by movement, that turns a vague brief — *take the sugar out, keep it delicious* — into a specified, reproducible formula with a number on it. It is the whole discipline of the Index made operational, and it runs in six movements plus the arithmetic that holds them honest.

But it begins before the first movement, with the small patient machine on the surface of the tongue that every sweetener in the pantry is quietly trying to persuade. You cannot reformulate what you have not measured; you cannot measure sweetness honestly until you know exactly what is listening. So we start there, with the receptor, because every number in this chapter is finally a statement about how loudly a molecule speaks to it.



The bench where a brief becomes a number. Every reformulation ends the way it began — on the Index, measured, reproducible to within seven percent.

The receptor you are negotiating with

Sweetness is read, almost entirely, by one receptor: **TAS1R2 locked to TAS1R3**, a heterodimer — two different proteins clasped as a working pair — sitting in the membrane of the taste cell as a class-C G-protein-coupled receptor. Picture each half as a hinged clamshell at the cell surface (the *Venus-flytrap domain*, VFT), joined through a cysteine-rich collar to a bundle of seven membrane-spanning helices below (the *transmembrane domain*, TMD). When the right molecule settles into the open clam and the shell swings shut, that

motion telegraphs down through the collar, the helices shift, and the cell depolarizes. A single closed flytrap, repeated a billion times a second across the tongue, is the physical fact we have spent our lives trying to describe in poetry and are here, at last, measuring.

What matters most to a formulator is that this dimer has **more than one door.** Small sugars — sucrose, and the near-sugars allulose and erythritol — nestle into the open cleft of the TAS1R2 flytrap and clamp it. Several high-potency synthetics touch a different site altogether, the transmembrane bundle of TAS1R3, from underneath. And the sweet proteins — thaumatin, brazzein — are far too large to sit inside any cleft, so they do something else entirely: they *wedge* against the outer face of the flytrap and hold it shut, a doorstop rather than a key. Three mechanisms, three regions of one receptor. That multiplicity of doors is the single most important fact in this chapter, because it is the physical reason a blend can do what no lone molecule can — two sweeteners at two doors are not fighting over one lock.

The second fact is **efficacy.** Not every molecule that binds pulls the door fully shut. A *full agonist* — sugar — can drive the receptor to its maximum response; pour on more and the only limit is saturation. A *partial agonist* closes the flytrap only partway no matter how much you add: it has an intrinsic efficacy we write as α, a number between 0 and 1 that caps how far this particular key can ever turn the lock. The steviol glycosides — stevia's sweet molecules — behave as partial agonists, and that single pharmacological fact, once you put a number to it, becomes a hard ceiling that governs more reformulations than any flavor note ever will.

A third fact is the off-note, and it too lives in the receptor map. Several sweeteners that speak beautifully to TAS1R2·R3 also brush the bitter receptors next door — the TAS2R family — at the doses needed for full sweetness. That faint licorice behind stevia, the metallic edge on saccharin: those are not sweetness at all but a second receptor answering a molecule that was only meant for the first. Reformulation is, in large part, the art of speaking loudly at one door without knocking on the others.

What the number actually measures — *the mathematics of loudness.*

Every axis on the Index is anchored to half-activation. The master relation is simple and we apply it unchanged here:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the dose of quiet or loud it takes to get halfway. Invert it against a fixed reference and scale by a thousand. For sweetness the natural reference is **sucrose**, which therefore reads exactly **IFI-S 1,000** by definition, because EC_{50}(sucrose) ÷ EC_{50}(sucrose) = 1. Every other sweetener's *molecular loudness* falls

straight out of its potency relative to sugar: a molecule that is 200× as sweet as sucrose has an EC_{50} one two-hundredth as large, so its molecular IFI-S is 1000 × 200 = 200,000. The whole pantry, scored:

SWEETENER	POTENCY (× SUCROSE)	MOLECULAR IFI-S	DOOR ON THE RECEPTOR	TIME SIGNATURE
Sucrose (reference)	1	1,000	VFT cleft (TAS1R2)	fast on, clean, fast off
Allulose	0.7	700	VFT cleft (TAS1R2)	sugar-like, clean
Erythritol	0.7	700	VFT cleft (TAS1R2)	fast, faint cooling
Acesulfame-K	200	200,000	VFT + TMD	fast, slight bitter tail
Aspartame	200	200,000	VFT cleft (TAS1R2)	slight delay, clean, lingers
Saccharin	350	350,000	VFT + TMD	fast, metallic tail
Reb A (steviol)	250	250,000	VFT/TMD, **partial (α≈0.6)**	slow, licorice, long tail
Reb M (steviol)	300	300,000	VFT/TMD, **partial (α≈0.6)**	rounder, cleaner, moderate linger
Mogroside V (monk fruit)	250	250,000	VFT cleft (TAS1R2)	rounded, slow, lingering-sweet
Sucralose	600	600,000	VFT cleft (TAS1R2)	slight delay, lingering
Brazzein	800	800,000	CRD / VFT outer face (wedge)	slow on, very long linger
Thaumatin	2,000	2,000,000	CRD / VFT outer face (wedge)	very slow, extremely long, licorice-ish
Neotame	8,000	8,000,000	VFT cleft (TAS1R2)	delayed, very long linger

SWEETENER	POTENCY (× SUCROSE)	MOLECULAR IFI-S	DOOR ON THE RECEPTOR	TIME SIGNATURE
Advantame	20,000	20,000,000	VFT cleft (TAS1R2)	delayed, very long linger

Now the distinction a formulator must hold in both hands at once. That six- and seven-figure *molecular* loudness is the loudness of the pure compound, per unit — it is enormous because a trace does the work. But you cannot put 600,000 on a scorecard, and you would not want to: the eight-axis profile measures the **delivered** intensity of a finished product, how far up the receptor's response curve this particular formula, at the concentrations actually present, pushes the sweet axis. That delivered number runs on the Index's normalized 0–5,000 scale, where 5,000 is full receptor activation. Write the delivered sweetness as

$$\text{IFI-S(delivered)} = 5000 \times \alpha \times o, \text{ where } o = [S] \div (EC_{50} + [S])$$

Here *o* is fractional occupancy (0 to 1, how much of the receptor population is bound at the dose present) and α is the molecule's intrinsic efficacy — 1 for a full agonist like sugar, less than 1 for a partial one. A full-sugar cola reads **S = 1,300.** Divide by the 5,000 ceiling and the receptor is only 26% activated — θ = 0.26 — which tells you something lovely and practical at once: a sugared soda is nowhere near the top of the curve. There is headroom. The brief is not to reach the ceiling; it is to land, exactly, on 1,300 by some other means than sugar, and to rebuild whatever else the sugar was quietly doing on the way.

The partial-agonist ceiling — *why one molecule cannot win.*

Set α = 1 (sugar, a full agonist) and the delivered sweetness can be driven to the full 5,000 — occupancy is the only limit. Set α = 0.6, the figure we use for the steviol glycosides, and the mathematics turns hard: the most delivered sweetness stevia can *ever* produce, at infinite dose, is 5000 × 0.6 = **3,000.** That is a wall, not a slope. No quantity of stevia reaches 3,100.

Work a real target through it. To deliver the cola's S = 1,300, sugar needs occupancy o = 1300 ÷ 5000 = 0.26, which in dose terms is [S] ÷ EC_{50} = 0.26 ÷ 0.74 = **0.35** — a little over a third of its own half-max. Stevia, capped at α = 0.6, must reach the *same* 1,300 through 1300 = 3000 × o, so o = 0.43, i.e. [S] ÷ EC_{50} = 0.43 ÷ 0.57 = **0.76.** It has to hold the receptor nearly twice as fully, and sit about **2.2× deeper into its own dose curve,** to deliver an identical sweetness — and deeper into the curve is precisely where its licorice brush against the bitter receptors grows loud. Push the target to a dense S = 3,800 and stevia alone is not merely strained, it is *impossible*: 3,800 lies above its 3,000 ceiling, and the curve flattens forever below the goal. This is the quiet reason, written in α, that the partial agonists cannot carry a formula by themselves, and the reason the honest answer to a stalled sweetness curve is never "more of the one that has plateaued" but "a fuller-agonist partner at a second door." Every blend in this book is that sentence, made numeric.

Movement one — *set the target as a number.*

Do not start with "less sweet" or "as sweet as before." Those are weather, not coordinates. Put the original, full-sugar product on the Index and read its **delivered IFI-S** — and while the sample is in the cell, read the axes that travel with sweetness and shape how it is perceived: its **acidity** (IFI-A, which suppresses sweetness), its **body and continuity** (IFI-K, kokumi), any **bitterness** (IFI-B) hiding under the sugar. Now you hold a target *vector*, not a vibe: *match S = 1,300; hold A near 400; rebuild K toward 700; keep B at zero.* Everything downstream is in service of landing on that vector, and because it is numbers you will know, at the end, whether you did.

Movement two — *map the sugar's hidden jobs.*

Ask, before touching a sweetener, what the sugar was actually doing in *this* product besides tasting sweet — because a sweetener fills exactly one of its jobs and sugar was doing several at once. A beverage leans on sugar for body and the weight of the liquid in the mouth; a baked good on bulk, on browning, on held moisture; a jam on preservation. List the vacancies you are about to open, because each one must be filled by something other than the sweetener, or the product will taste thin, or pale, or go stale, or spoil — "right" on the sweet axis and wrong everywhere else.

SUGAR'S HIDDEN JOB	THE PHYSICS	WHAT REFILLS THE VACANCY
Sweet taste	TAS1R2·R3 activation	the sweetener system (movement three)
Body / mouthfeel	dissolved-solids viscosity	soluble fiber, allulose, trace hydrocolloid
Bulk / volume	sugar is most of the dry mass	allulose, erythritol, polydextrose
Browning	Maillard + caramelization	allulose (a reducing sugar; it browns)
Moisture retention	sugar is hygroscopic	allulose, glycerin, humectant fibers
Preservation	sugar lowers water activity	fiber/polyol solids, formulation, pH
Continuity / fullness	length and roundness on the palate	kokumi (IFI-K) base, CaSR agonists

The lesson of the table is the thesis of the whole playbook: in a reduced-sugar product the sweetener is rarely the hard part. Diagnose which of these jobs the product truly leaned on, and the formula largely writes itself around *that*, with the sweetness almost an afterthought bolted onto a structural rebuild.

Movement three — *choose the sweetener system, not a sweetener.*

No single molecule wins, and the reason is pharmacological, not a matter of taste: the partial agonists cannot reach full intensity (their α ceiling); the potent synthetics each carry a temporal flaw; the

sweet proteins wedge the flytrap shut and linger far too long. So we blend — deliberately, for two reasons that are both written in the receptor.

The first is **synergy.** Because the dimer has several doors, two sweeteners binding different sites do not merely add — binding at one site can stabilize the active conformation and lower the effective EC_{50} at the other, so the pair delivers more combined sweetness than their separate doses predict. Formulators quantify this with a **synergy index**, SI = (perceived sucrose-equivalent of the blend) ÷ (sum of the component sucrose-equivalents dosed). Work the classic pair. Dose aspartame at a sucrose-equivalent of 3.0% and acesulfame-K at 3.0% — total *dosed* SE = 6.0% — and the blend is measured as roughly **7.2% SE** of perceived sweetness. SI = 7.2 ÷ 6.0 = **1.20**, a twenty-percent bonus the receptor gives you for free. Run it the other way, as a formulator actually does: to *deliver* a 6.0% SE sweetness you need dose only 6.0 ÷ 1.20 = **5.0% SE** total, split as 2.5 + 2.5 — about **17% less sweetener**, with each component now sitting lower on its own curve, where its off-note is quieter. Synergy is the rare lever that lowers cost and cleans the taste in the same motion.

The second is **time-intensity** — composing the *arc* of the sweetness, not just its peak. Every sweetener has an off-rate, k_off, at the receptor, and its sweetness lasts about as long as its residence time τ = 1 ÷ k_off. Sugars release fast: a quick, clean rise and an honest fall, there and gone. The steviol glycosides and mogroside V release slower — a rounder middle, a longer tail. The sweet proteins, wedged against the flytrap, release slowest of all and sweeten for a lingering minute. So you lay them like notes in a chord: a fast, clean front (a quick synthetic or a trace of sugar) for the attack, a mogroside-rich monk fruit for the rounded middle, a Reb-M-dominant stevia to carry the length without the licorice. Dose each component so the *blend* lands on the target delivered IFI-S — and remember the ceiling: if the curve will not reach the target, the fix is a fuller-agonist partner at a new door, never another spoonful of the molecule that has already plateaued.

Movement four — *manage the off-notes.*

Profile the bitterness as a number, IFI-B, and treat it as a design target, not an apology. Always prefer the clean molecule to the masked one: a Reb-M-dominant stevia and a mogroside-V-rich monk fruit start clean because the specific molecules that brush the bitter receptors are the lesser glycosides, not these, so there is simply less to hide. Where an off-note survives, decide honestly between *masking* it — layering a flavor that occupies the bitter channel — and *removing its trigger* by re-balancing the blend so the offending compound is dosed lower. Masking is a loan against the palate; removal is a payment. Then read the result's IFI-B so that "clean" is a number you can defend to a buyer, not an adjective you hope they share.

Movement five — *rebuild the floor.*



Sugarcane, the body being removed. Most of what sugar gave a product was never sweetness — it was the weight, the browning, the held water, the length on the palate that the floor rebuild must put back.

Put the physics back. Rebuild body and bulk with a bulking system — **allulose**, which also browns because it is a reducing sugar and enters the Maillard reaction as sucrose does; **erythritol** for clean bulk; **soluble fiber** for solids and a faint, honest viscosity; in liquids, a trace of hydrocolloid for the weight the dissolved sugar used to carry. Then add a **kokumi base** — agonists of the calcium-sensing receptor, CaSR — so the sweetness reads full and lasts the right length instead of arriving thin and leaving early. Kokumi adds no sweetness of its own; it lengthens and rounds what is there, the difference between a note and a chord.

And use the cross-axis lever, the cleverest trick in the book. Sourness suppresses sweetness — partly because protons blunt the receptor's responsiveness, and partly as a central effect, the mutual suppression the brain applies when sour and sweet arrive together. So trimming a formula's acid *raises* perceived sweetness at a fixed sweetener dose, which means you can *lower* the sweetener to hold the same perceived target. Sometimes the cheapest sweetness in the formula is the acid you remove. In a citrus soda, trimming the citric acid a hair — A from 400 to 380, a five-percent cut you can barely taste as sourness — lets the sweetener dose stay small while S holds dead on 1,300. The acid you give back buys the sugar you take away.

Movement six — *validate and lock.*



Reading the finished number. A reformulation that tastes right once is a prototype; one that reads the same eight numbers every batch, across a shelf life, is a product.

Taste-panel the candidate against the benchmark, read the finished eight-number vector, and then prove it *holds*. The platform's signature standard is **≤7% batch coefficient of variation**, and CV is the whole discipline in one formula: CV = 100 × (standard deviation ÷ mean). Run the soda ten times; suppose the delivered S reads a mean of 1,300 with a standard deviation of 70. Then CV = 100 × 70 ÷ 1,300 = **5.4%** — inside spec, a product. Read 1,300 with an SD of 130 and CV = 10%, and what you have is a good beaker and not yet a formula.

The subtle risk in sugar reformulation is precisely the trace high-potency sweetener, because enormous loudness in a tiny mass means tiny errors swing the number. Convert the units to see it. A 600× sweetener dosed at **60 ppm** is 60 mg/L, which is 0.0060% w/v — but its sucrose-equivalent is 60 mg/L × 600 = 36,000 mg/L = **36 g/L = 3.6% SE**, so sixty parts per million is quietly doing the sweetening work of three and a half grams of sugar in every hundred millilitres. That same leverage is what turns a small weighing or blend-uniformity error into a visible shift in sweetness. The discipline is twofold: **split the sweetness across components** so no single trace carries it all, and **pre-blend each trace sweetener into a carrier** so that uniformity, not raw potency, governs the batch. A 5% uniformity error on a component that carries half the sweetness lands as only ~2.5% on total S — comfortably inside the seven. A reformulation is finished not when it tastes right, but when it tastes right *every time*, to within seven percent, on a line running at scale, across suppliers and seasons and a full shelf life. That is the quiet reason a number beats a recipe.

The worked case — *a zero-sugar citrus soda.*

The most common brief we are handed: a full-sugar sparkling citrus soda, reformulated to zero added sugar. The benchmark reads sweet and tart with a sugar-given body —

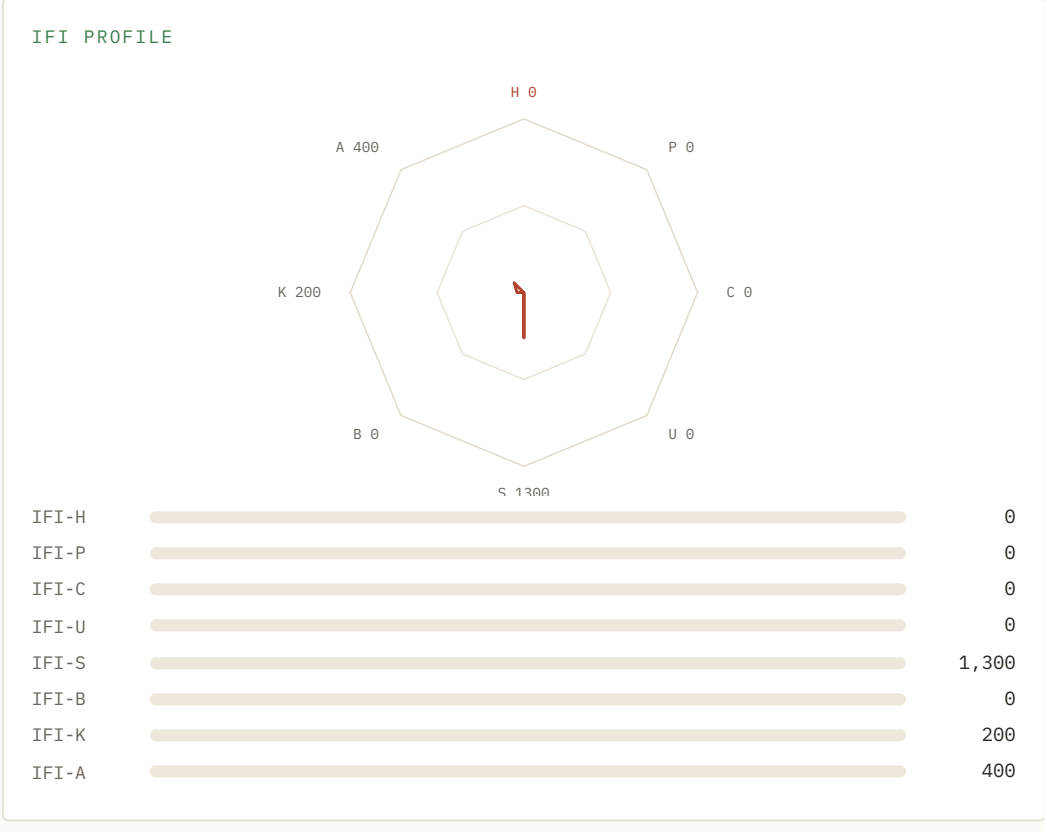

— and the reformulation holds the sweetness with a Reb-M stevia and mogroside-V monk fruit blend for the arc (a clean fast front, a rounded monk-fruit middle, a Reb-M length), allulose and a trace of soluble fiber for the lost body, a γ-glutamyl kokumi base for continuity, and the citric acid trimmed a hair so less sweetener is needed to balance it:

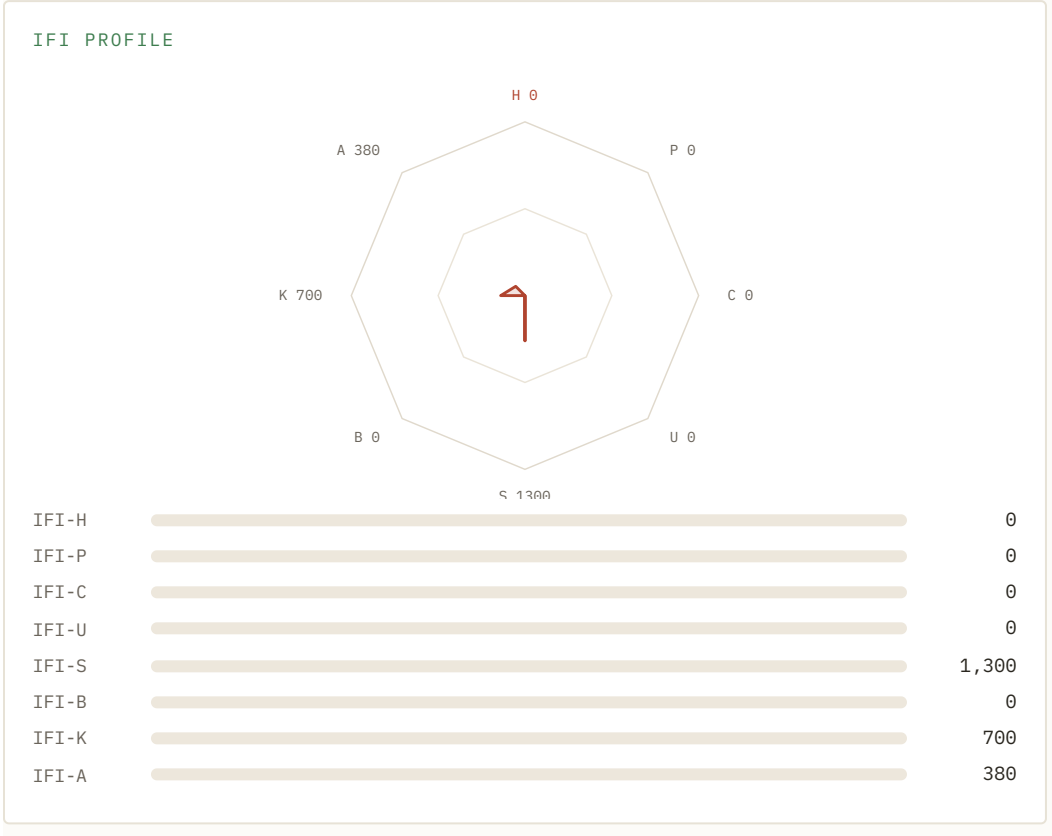

Read the two scorecards side by side and the entire method is visible in four numbers:

AXIS	FULL-SUGAR BENCHMARK	ZERO-SUGAR REFORMULATION	DESIGN MOVE
S (sweetness)	1,300	1,300	matched exactly — blend to target θ = 0.26
A (acidity)	400	380	trimmed 5% so the sweetener dose stays small
K (kokumi)	200	**700**	**rebuilt** — the sugar's body, put back on purpose
B (bitterness)	0	0	clean molecules chosen, so nothing to mask

The sweetness lands on the same number, 1,300 — a person could not tell the two apart on the sweet axis, because θ is identical whatever molecule delivered it. The acid is a hair lower, by design, which let the sweetener dose stay down near where the stevia's licorice is silent. And the real labor is the **K rising from 200 to 700**: the body the sugar used to give for free, rebuilt deliberately out of bulk and kokumi, so the zero-sugar version reads full in the mouth instead of watery. That is a successful reformulation in two scorecards — matched where it must match, rebuilt where the sugar was quietly working, and the arithmetic of cost quietly in your favor, because a 60-ppm trace at 600× potency costs, in the glass, a fraction of the sugar it replaced even when its price per kilogram is far higher.

The harder brief — *a reduced-sugar bake.*

Not every target is a beverage, and the bakery is where the hidden jobs bite hardest. Here the sweetener system is the easy part and the real work is restoring bulk, browning, and moisture. A thirty-

percent sugar cut in a cookie is typically met not by a stronger sweetener at all but by an **allulose-led bulking system** — allulose because it is a reducing sugar that browns through the Maillard reaction, and because it is hygroscopic and holds moisture as sucrose does, so the crumb colors and the cookie keeps its day-two softness instead of going to sand. Do the mass arithmetic before the taste: if sugar is, say, 40 g in a 120 g dough and you cut it 30%, you have removed **12 g of dry mass and most of a browning and humectant system**, and 12 g of nothing is a flat, pale, brittle cookie. Refill it — roughly gram-for-gram in bulk with allulose and fiber — then bolt on only the modest sweetness the allulose does not supply (it reads about 0.7× sugar) with a trace of Reb M and monk fruit, and tune the salt and the small acid so the smaller sweetness still reads full. The sweetness is almost the last decision, not the first. Diagnose which of sugar's jobs the bake truly leaned on, and the formulation assembles itself around the structural rebuild, with the sweetener an afterthought.

Where reformulations fail

The honest part of any method is its failure modes, and sugar reduction has a short, repeatable list. **The partial-agonist wall:** chasing a target above a partial agonist's α-ceiling by adding more of it, so that only the off-note grows while the sweetness does not — the cure is always a fuller agonist at a second door. **Masking stacking:** hiding one off-note under a flavor, which adds a note of its own that wants hiding, and so on, until the palate is crowded — the discipline is clean molecules first, masking last. **The lingering tail:** a sweet-protein or high-mogroside blend whose long residence time is wrong for a product meant to be quick and clean, so the sweetness overstays and reads cloying. **Shelf-life drift:** the number that was right at bottling moving off over months — aspartame, for instance, slowly hydrolyzes in warm, acidic liquids and loses sweetness, so a formula validated only at day zero can fall out of spec by its sell-by date; the fix is to validate across the shelf life, not just the batch. **Temperature dependence:** perceived sweetness shifts with serving temperature, so a formula tuned warm can read differently over ice. And **trace-dose CV blow-ups:** the leverage that makes a trace sweetener cheap also makes its uniformity the dominant source of batch variation, which is why the sweetness is split across components and the traces are pre-blended into carriers. None of these is exotic; all of them are why the last movement is a stability study and not a victory lap.

What we may and may not say

A closing word on what we may and may not say, because it is a matter of both law and ethics and this library holds the line hard. Sweetener use is governed by real approvals, and a formulator must respect them:

SWEETENER	US STATUS	EU
Aspartame, Acesulfame-K, Sucralose, Saccharin	approved	E951, E950, E955, E954
Neotame, Advantame	approved	E961, E969
Steviol glycosides (Reb A, Reb M)	GRAS	E960
Monk fruit (mogrosides)	GRAS	flavor use
Thaumatin	GRAS (flavor)	E957
Erythritol	GRAS	E968
Allulose	GRAS	limited
Brazzein	not yet approved (frontier)	not approved

Claims about the *product* are compositional attributes with legal definitions — "no added sugar," "reduced sugar" (a quarter less, at least), "sugar-free" (under half a gram a serving) — and those are the only kind of claim this book makes. We design to the **attribute**: the sweetness, the number, the absence of added sugar, the clean label, the eight-figure vector that reads the same every batch. We do not, ever, design to a health outcome, and nothing in this library should be read as one. Our family's ethos — that food should nourish, that the body should be fed with intention, that a plant and the people at the table both deserve the care of getting it exactly right — is a *value* that guides what we build, never a claim about what it cures. Sweetness without added sugar is a fact you can put on a label and a pleasure you can put on a tongue. What it does for a body is between that body and the life it is living. We make the flavor honest, and reproducible to within seven percent. The rest is not ours to promise.

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

Chapter 17 · Formulations — *the sweetness, built.*

Here, at last, the sweetness leaves the bench and lands on the table. These are working formulations — a handful of the ways our family uses designed sweetness in the kitchen and the lab, each one built on engineered plants and read out on the Index. They are plant-forward by default, as our whole table is; where meat and seafood appear, they appear with full respect, not apology, because a sweetness built well serves them as honestly as it serves a squash. Every one is sugar-light or sugar-free by construction, and every one carries its floor — a kokumi base, a managed acid — so the sweetness lands whole.

Before the recipes, though, we want to show the work underneath them: how a sweet taste becomes a number, how that number is derived from a receptor, and how a cook or a formulator can hold a target of $S \approx 1{,}600$ in the hand the way they hold a teaspoon. The pleasure is real, and so is the arithmetic. They are two halves of one sentence.

The receptor that reads sweet

Every sweet thing you have ever loved was read by one molecular machine. On the surface of the taste cell sits **TAS1R2 paired with TAS1R3**, two proteins clasped into a single heterodimeric receptor of the class-C G-protein family. Its signature is a large extracellular lobe — the *Venus flytrap domain*, named for the way its two halves hinge shut around a ligand. Sucrose, fructose, and the rare sugar allulose settle into the orthosteric cleft of that flytrap and pull it closed. The small high-potency sweeteners — the steviol glycosides, the mogrosides of monk fruit — find their own pockets, some in the same flytrap, some down in the transmembrane bundle. The sweet proteins, brazzein and thaumatin, are too large for any single pocket; they dock across a broad face of the receptor like a wedge pressed into a seam, which is exactly why they survive heat and acid that would denature a lesser ligand.

Whatever the route in, the downstream story is one story. The closed receptor swaps a GDP for a GTP on its coupled G-protein, **α-gustducin**, which releases its βγ partners to switch on **phospholipase C-β2**. PLCβ2 cleaves membrane PIP_2 into **IP_3**, IP_3 opens its receptor on the endoplasmic reticulum, and a pulse of calcium floods the cytosol. That calcium gates **TRPM5**, a cation channel that depolarizes the cell, and the depolarization opens **CALHM1**, through which the cell exhales ATP onto the afferent nerve. From flytrap to nerve is a few hundred milliseconds. "Sweet" is that cascade, firing.

Two facts about this machine are the whole reason designed sweetness is possible. First, the receptor does not count calories; it counts *occupancy*. A molecule that closes the flytrap registers as sweet whether it carries four kilocalories per gram or none. Second, the response is **saturating** — there is a maximum rate at which the

cascade can fire, and past a certain occupancy more ligand buys no more loudness. A sweetness engineered from the receptor outward lives inside both of those facts: it borrows the first to drop the sugar, and it respects the second so it never over-reaches into the flat top of the curve where sweetness turns cloying and false.



Sugarcane in cross-section — the plant whose sucrose we use as the reference stimulus, the 10% solution against which every sweetness on the Index is read.

From a taste to a number

The Index turns that cascade into a score with one formula, the same one that anchors every axis:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the midpoint of the dose–response curve. We read it off a Hill fit of a dilution series on a cell line expressing TAS1R2·TAS1R3 with a calcium-flux readout:

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

At $c = EC_{50}$, the fraction $c^n/(EC_{50}^n + c^n)$ equals ½ regardless of the Hill coefficient *n*; that is what makes the midpoint such an honest anchor. For the sweet receptor *n* runs close to 1 — call it 1.0 to 1.3 — so the curve is a gentle sigmoid without the steep cooperative switch you see on some channels.

The reference is the oldest sweetness we know: **cane sugar.** We fix the anchor at a **10% w/v sucrose solution** — 100 grams per litre, which is 292 millimolar once you divide by sucrose's molar mass of 342.30 g/mol — and we *define* its IFI-S as exactly 1,000. Everything else is read against it. Because the formula inverts the EC_{50} ratio, a candidate that half-fires the receptor at a *lower* concentration than sucrose is the louder sweetness and scores *above* 1,000; one that needs more to reach half-max scores below.

That gives the number a plain-language twin. A formulation that half-saturates the sweet receptor at 62.5 g/L of itself — needing only five-eighths of the mass sucrose needs — scores:

$$\text{IFI-S} = 1{,}000 \times 100 \text{ g/L} \div 62.5 \text{ g/L} = \mathbf{1{,}600}$$

and a cook reads the same 1,600 as *"as sweet as a 16% sucrose solution."* Across the range our kitchen works in, the shorthand holds cleanly: **IFI-S ≈ 100 × (sucrose-equivalent % w/v).** Sixteen per cent reads 1,600; twenty-two per cent reads 2,200; eleven per cent reads 1,100. The bench measures the EC_{50}; the tongue matches a sucrose ladder; the anchor at 292 mM ↔ 10% ↔ 1,000 is what guarantees the two readings return the same number.

Here is that measurement worked all the way through, on the Citrus Cordial that opens the recipes below. We serially dilute the finished cordial, read each dilution's response off the calibrated curve, and fit the half-max:

DILUTION OF CORDIAL	FORMULA CONC. (G/L)	RECEPTOR RESPONSE (% OF MAX)
neat	125.0	66.7
1 : 1	62.5	50.0
1 : 3	31.3	33.3
1 : 7	15.6	20.0
1 : 15	7.8	11.1

The fit (Hill, $n \approx 1.1$) puts the half-max at **EC_{50} = 62.5 g/L**. Into the formula:

```
IFI-S = 1,000 × 100 ÷ 62.5 = 1,600
```

— and the sensory panel, matching the neat cordial against a sucrose series, lands it at 16% equivalent. The two agree, as the anchor promises they must.

The anchor also lets us read a sweetness even where no one would call the food sweet. Our hottest cultivar, the eight-axis superhot **Phoenix**, carries a designed, deliberate sweetness tucked under its heat — a small lift that keeps a predicted 2.95 million SHU from reading as pure assault:

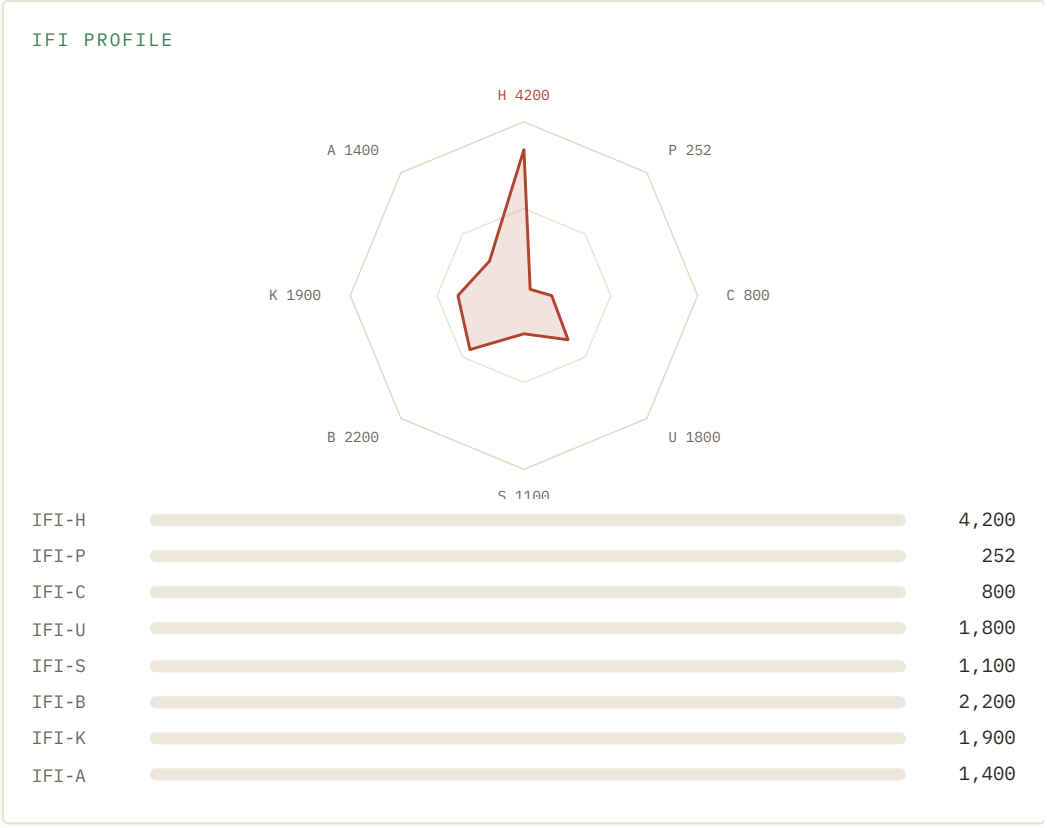

An *S* of 1,100 is the sweetness of an 11% sucrose solution — faint against the fire, but present, and present *because it was designed to be*, not because sugar was poured in. That is the whole claim of the

Index: it reads sweet wherever sweet lives, and it reads it as a number you can put in a spec.

Precision, and what the number is worth

A spec is only as good as its tolerance, so it matters that designed sweetness is *reproducible*. The platform holds a batch-to-batch coefficient of variation at or below **7%** on the engineered output — against the roughly 42% scatter a wild population throws. Since IFI-S is a ratio to a fixed, calibrated sucrose reference whose own variance we treat as negligible, the score inherits the output's CV almost unchanged:

$$\mathrm{CV}(\text{IFI-S}) \approx \mathrm{CV}(EC_{50}, \text{candidate}) \leq 7\%$$

So a formula written to *S = 1,600* ships, batch after batch, inside roughly **1,490 to 1,710** — about ±112 on the score, a tolerance tight enough to hand a factory without a tasting panel waiting at the end of the line. When two measured quantities enter the ratio — a candidate read against a freshly measured reference — the errors add in quadrature, $\sqrt{(7^2 + 7^2)} \approx 9.9\%$, and the band widens to roughly ±158. Either way the sweetness is a managed variable, not a hope. Numbers this stable are a form of respect for whoever cooks from them.

The palette — four molecules, one receptor

We build sweetness from a small, clean set of ligands, each chosen for where it sits on the dose–response curve and for how it behaves in the mouth over time. None of them is sugar; all of them close the same flytrap.

SWEETENER	SOURCE	CLASS	RELATIVE SWEETNESS (× SUCROSE, BY WEIGHT)	PLATEAU (E_{max} VS SUCROSE)	ROLE IN A FORMULA
Sucrose	cane / beet	disaccharide	1	100% (defines it)	reference only
Allulose (D-psicose)	rare sugar	monosaccharide	~0.7	~90%	body, gloss, browning
Rebaudioside M	stevia leaf	steviol glycoside	~200	~85%	primary sweetness
Mogroside V	monk fruit	cucurbitane glycoside	~250	~85%	rounds the mid-palate
Brazzein	sweet protein (54 aa)	protein	~800	~80%	heat-stable lift, long finish
Thaumatin	sweet protein (207 aa)	protein	~2,000	~75%	trace enhancer, licorice linger

Allulose is the workhorse of texture. It is a true sugar — the C-3 epimer of fructose, molar mass 180.16 — so it sits in the orthosteric cleft and browns like a sugar through the Maillard reaction, yet it passes through the body nearly unmetabolized. At about seventy per

cent of sucrose's sweetness it does what the high-potency compounds cannot: it gives *mass*. Forty grams of it in a syrup is forty grams of body, gloss, and thread. The engineered sweetness host for bulk sweetness in our program is **sugarcane**, 200 constructs deep, tuned toward cleaner sucrose and the rare-sugar routes beside it.

Rebaudioside M is the cleanest of the steviol glycosides — the reason we reach for Reb-M and not the more common, more licorice-edged Reb-A. From **stevia** (910 constructs, our most-developed sweetness host) it arrives roughly two hundred times sweeter than sucrose by weight, which is why four-tenths of a gram sweetens a cordial: 0.4 g in 600 mL is 0.52 millimolar, a whisper of material driving a full receptor response. Its tells are a slightly slow onset and a faint, cooling linger.

Mogroside V, the cucurbitane glycoside of **monk fruit** (80 constructs), is near Reb-M in potency but different in shape and timing — it fills the middle of the palate where stevia can run thin, and it masks stevia's edge. The two are almost always used together for exactly that reason.

Brazzein and **thaumatin** are the sweet proteins, and they are a different physics. Because they dock across a broad receptor surface rather than in a single pocket, they survive baking and acid, and they deliver a lingering, wrapping sweetness. We use them in traces, as lift and length, not as the backbone.

Why you cannot simply add the sweetness up

The most important number in formulation is the one the naive arithmetic gets wrong. Take the Citrus Cordial's three sweeteners and tally their sucrose-equivalent contributions in the 600 mL base, each as (grams ÷ litres) × (potency):

- **Allulose** 40 g → 66.7 g/L × 0.7 = **46.7 g/L** ≈ 4.7% SE
- **Reb-M** 0.4 g → 0.667 g/L × 200 = **133 g/L** ≈ 13.3% SE
- **Mogroside V** 0.3 g → 0.5 g/L × 250 = **125 g/L** ≈ 12.5% SE

The straight sum is **30.5% sucrose-equivalent** — which would predict an IFI-S north of 3,000. The cordial measures **16%**, IFI-S 1,600. The gap is not an error; it is the receptor telling the truth about itself. Reb-M and mogroside V are *each, alone,* already above the 10% reference — each is individually loud enough to drive the sweet cascade toward its ceiling. Stacking two near-saturating ligands onto one shared receptor population cannot double the firing rate; the occupancies overlap, and the sum collapses toward the plateau. The sub-additivity factor here is φ = 16 ÷ 30.5 ≈ **0.52**.

That is the central lesson of designed sweetness, and it is why you dose to a *measured* target rather than a *calculated* one. You blend not to add intensities — the receptor won't let you — but to buy things intensity can't give you: the fullness mogroside lends stevia, the masking each lends the other, the mass allulose lends them both, and a temporal profile that opens and clears like sugar instead of arriving late and overstaying. (The rule has an exception worth knowing: at *sub-threshold* levels some sweetener pairs run mildly

supra-additive, nudging each other over the line. Near saturation, where our formulas live, sub-additivity rules. In both cases the honest move is the same — measure.)

With the mechanism, the math, and the palette in hand, the formulations read as what they are: targets on the sweet axis, each with a floor built beneath it and an acid managed around it.

Zero-Added-Sugar Citrus Cordial

A concentrate to cut with sparkling water — a clean-label soda base.

```
zest + juice of 4 oranges, 2 lemons
600 ml water
Reb-M stevia extract, to S≈1,600 (begin at 0.4 g)
mogroside-V monk fruit extract (begin at 0.3 g)
40 g allulose (body + gloss)
1 g γ-glutamyl kokumi base
pinch fine sea salt
```

1. Warm the water with the allulose and salt until dissolved; do not boil.
2. Off the heat, stir in the zest and steep 15 minutes, then strain.
3. Add the juice, the stevia and monk fruit extracts, and the kokumi base.
4. Taste against a 10% sucrose reference; adjust the two extracts until the sweetness reads full and clean. Bottle, chill. Cut 1:4 with sparkling water.

The acid is the reason this works. Citrus brings a bright load of citric and ascorbic acid, and acid and sweet are read by different machines — sour through the proton-gated channels of the sour cell, sweet through TAS1R2·R3 — so they do not compete at the receptor; they *frame* each other. A touch of sourness sharpens the edge of a sweetness and keeps it from reading flat, which is exactly why a sugar-free sweetness, prone to a slightly rounded, bodiless middle, is rescued by a well-managed acid. The allulose does the rest, replacing the viscosity and gloss that sugar would have supplied, and the pinch of salt lifts the whole thing — sodium suppresses the faint bitterness that trails some high-potency sweeteners and, by removing that competition, lets the sweet read louder without a grain more of it.

IFI PROFILE

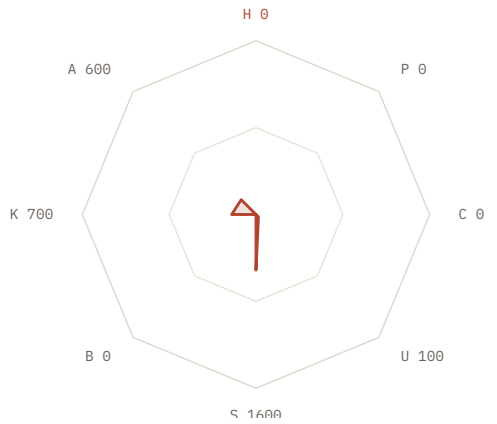

IFI-H		0
IFI-P		0
IFI-C		0
IFI-U		100
IFI-S		1,600
IFI-B		0
IFI-K		700
IFI-A		600



Orange, cross-section of a cordial's whole argument: the sweetness restored, the acid managed, the body rebuilt — no added sugar in any of it.

Dark Drinking Chocolate, Sugar-Free

Where sweetness meets the Zero-Bitter axis — cocoa's bitterness attenuated, its kokumi body left full.

```
500 ml oat milk
40 g cocoa (ideally an IFI-B attenuated, low-bitter cocoa)
Reb-M stevia, to taste (begin 0.2 g)
mogroside-V monk fruit, a pinch (rounds the middle)
1 g kokumi base
flaky salt, to finish
```

1. Whisk the cocoa into a little cold oat milk to a paste.
2. Warm the rest with the kokumi base; whisk in the paste.

3. Sweeten with the stevia and monk fruit to just-sweet — dark chocolate wants restraint.
4. Finish with salt, which lifts the sweetness without adding any.

Chocolate is a negotiation between two receptor systems at once. Cocoa's theobromine and flavanols fire the bitter receptors — TAS2R38 chief among a family of them — and sweetness has historically been the brute tool for burying that bitterness under sugar. We would rather *attenuate* the bitter at its source, using a low-bitterness cocoa that quiets the TAS2R signal by design, so the sweetness is free to be a grace note instead of a cover-up. That is why the sugar drops away and the formula wants restraint: with less bitterness to mask, a little sweet goes a long way. Underneath both sits the kokumi base — γ-glutamyl peptides that are themselves tasteless but act on the calcium-sensing receptor, **CaSR**, to amplify the continuity and roundness of everything around them. Kokumi is the body of this drink. The flaky salt at the end does its double duty again, quieting the last of the bitterness and lifting the sweet.

IFI PROFILE

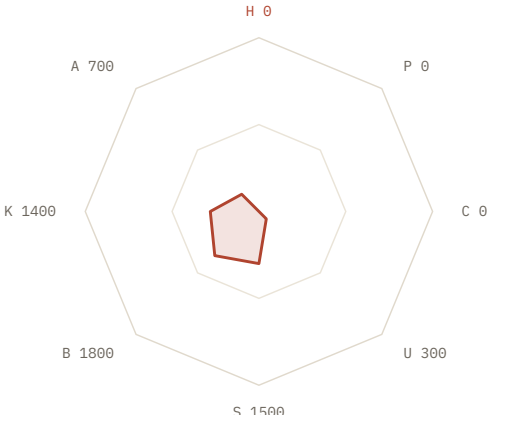

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	300
IFI-S	1,500
IFI-B	1,800
IFI-K	1,400
IFI-A	700

Strawberry Sweet-Heat Shrub

Sweet first, a slow warmth after, bright with vinegar.

```
300 g ripe strawberries, hulled
60 g allulose
mogroside-V monk fruit, to S≈1,900
1 small red chili, slit (or a few grains of a superhot, to H≈700)
80 ml white wine vinegar
```

1. Macerate the strawberries with the allulose and the chili 2 hours, until syrupy.
2. Crush lightly, add the vinegar, steep overnight in the cold.
3. Strain; sweeten with monk fruit until the fruit reads candy-sweet up front.
4. Serve over ice with soda. The sweetness lands first; the heat rises through the finish.

The pleasure of this glass is a matter of *kinetics*, and the receptors cooperate to give it. Sweetness is a fast signal: the sugars and mogrosides close the TAS1R2 flytrap and the gustducin cascade fires within a few hundred milliseconds, so the strawberry reads candy-sweet almost the instant it touches the tongue — then clears, because the ligands wash off and the calcium pulse subsides. Capsaicin is slow by contrast. It is lipophilic; it has to partition into the membrane around the TRPV1 channel before it lowers that channel's thermal threshold, so the warmth builds over ten or fifteen seconds and lingers for minutes after the sweet is gone. Tuned to *S ≈ 1,900* and *H ≈ 700*, the shrub is choreographed by that difference: sweet on the front of the sip, a clean warmth rising through the finish. The vinegar's acid, read by yet a third system, brightens the sweetness at the top and keeps the whole thing from turning syrupy.

IFI PROFILE

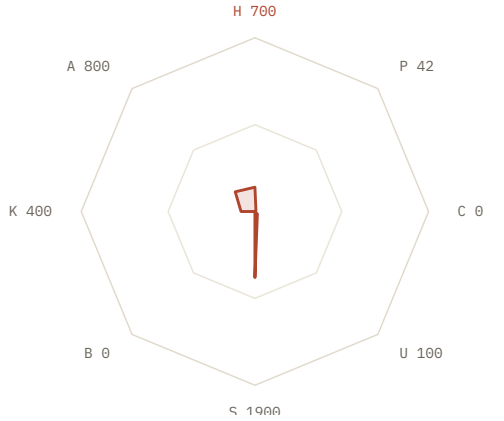

IFI-H	700
IFI-P	42
IFI-C	0
IFI-U	100
IFI-S	1,900
IFI-B	0
IFI-K	400
IFI-A	800



Strawberry specimen — sweet first, fang second. In the shrub the sequence is the whole pleasure.

Clean Pancake Syrup

An everyday staple rebuilt — the pourable body of allulose, the sweetness of the leaf.

```
200 g allulose
120 ml water
Reb-M stevia, to S≈2,200
mogroside-V monk fruit, a pinch
¼ tsp fine salt
natural maple flavor, to taste
1 g kokumi base (for a long, round finish)
```

1. Simmer the allulose and water until it threads and thickens, 6–8 minutes.
2. Off the heat, stir in salt, kokumi base, maple flavor.
3. Sweeten with stevia and monk fruit; cool to pouring consistency.

Syrup is the purest test of why bulk matters. At *S = 2,200* — the sweetness of a 22% sucrose solution — a formula built on high-potency sweeteners alone would taste intensely sweet and feel like water, because the receptor reads the sweetness but the hand and tongue find no body. Allulose is the answer: two hundred grams of it simmered down threads and coats a spoon because it behaves, physically, like the sugar it replaces, right down to the way it browns. That the sweetness itself comes mostly from a few tenths of a gram of stevia and a pinch of monk fruit is the quiet trick — the mass and the sweetness are supplied by different molecules, each doing only what it is best at. The kokumi base lengthens the finish so the last of the sweetness rounds off instead of dropping away, the way real maple's minerals and acids round off a sugar.

IFI PROFILE

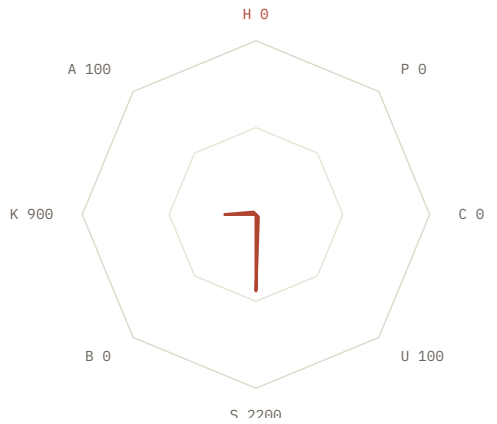

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	100
IFI-S	2,200
IFI-B	0
IFI-K	900
IFI-A	100

Monk-Fruit & Miso Glaze for Roasted Squash

Sweetness built over a deep savory floor — the plant-forward center of our table.

```
2 tbsp white miso (umami + kokumi)
mogroside-V monk fruit, to S≈1,600
1 tbsp rice vinegar
1 tsp grated ginger (a little P)
a few grains superhot chili, to H≈300
1 tbsp water
1 kg winter squash, in wedges
```

1. Whisk the miso, monk fruit, vinegar, ginger, chili, and water to a glaze.
2. Roast the squash at 220 °C until tender, 25 minutes.
3. Brush with the glaze; roast 8 minutes more, until lacquered.
4. The sweetness sits on a savory floor so deep it reads like caramel.

This is the dish that best shows what a *floor* is. Miso brings two things at once: free glutamate, which fires the umami receptors — mGluR1 and the TAS1R1·TAS1R3 pair — and the γ-glutamyl peptides of long fermentation, which act on CaSR as kokumi. Set a monk-fruit sweetness at *S ≈ 1,600* on top of that umami-and-kokumi floor and the two do not merely coexist; they *potentiate*. Kokumi, firing CaSR, amplifies the perceived continuity of the sweet and savory both, so the glaze reads deeper and more caramel-like than its modest sugar-equivalent could explain. The roast then adds the real caramel chemistry — the squash's own reducing sugars and amino acids browning in the Maillard reaction at 220 °C, the glaze lacquering as it goes. Ginger lends a trace of pungency through TRPA1, a few grains of chili a whisper of warmth, and the rice vinegar keeps the

sweetness lifted. A squash grown and glazed this way is the center of our table on its own merits, nourishment and pleasure in one wedge.

IFI PROFILE

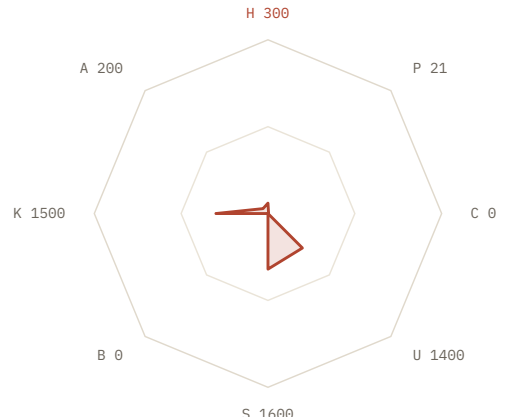

IFI-H	300
IFI-P	21
IFI-C	0
IFI-U	1,400
IFI-S	1,600
IFI-B	0
IFI-K	1,500
IFI-A	200

Sweet-Heat Glaze, on Salmon

Written with full respect for the fish — the same glaze family, where a clean sweetness and a measured heat flatter a rich, cold-water flesh as truly as they flatter a squash.

```
2 tbsp tamari (umami floor)
mogroside-V monk fruit, to S≈1,500
1 tbsp rice vinegar
1 tsp grated ginger
chili, to H≈900 (heat that arrives after the sweet)
4 salmon fillets
```

1. Whisk the glaze; reserve half.
2. Sear the salmon skin-side down, then glaze and finish under a hot grill, basting once.
3. Rest briefly; spoon over the reserved glaze. Sweet on the lips, warmth on the finish, the fish rich beneath both.

The design here is identical in grammar to the squash glaze — a sweetness set on a glutamate-and-kokumi floor, a managed acid, a measured heat — and that is the point worth making plainly: a sweetness designed from the receptor outward is portable across the whole table. Tamari supplies the umami floor in place of miso; the monk-fruit sweetness runs a touch lower, *S ≈ 1,500*, and the heat a touch higher, *H ≈ 900*, because a rich cold-water flesh carries a firmer sweetness and a cleaner burn than a squash does. The kinetics still do the choreography — fast sweet on the lips, slow capsaicin warmth rising through the finish, the fish's own fat smoothing both. We write this one with the same care we bring to any plate; respect for the fish is respect for the cook who chose it.

IFI PROFILE

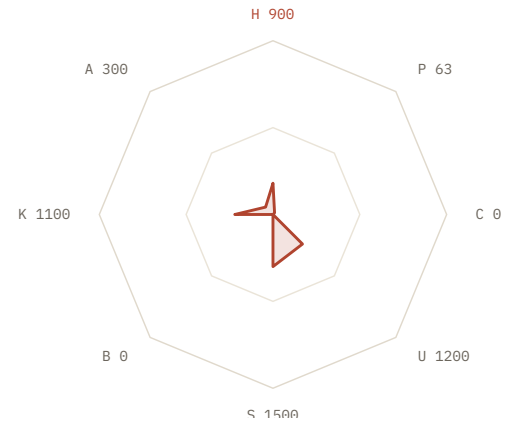

IFI-H	900
IFI-P	63
IFI-C	0
IFI-U	1,200
IFI-S	1,500
IFI-B	0
IFI-K	1,100
IFI-A	300

Sparkling Hibiscus-Berry Soda

The clean-living close — a bright, tart, zero-added-sugar refresher, sweetness and sourness in balance.

```
2 tbsp dried hibiscus
400 ml water
100 g mixed berries
Reb-M stevia + mogroside-V monk fruit, to S≈1,300
1 g kokumi base
sparkling water, to top
```

1. Steep the hibiscus in just-boiled water 10 minutes; strain.
2. Muddle the berries into the warm infusion; strain again.
3. Sweeten to S≈1,300 with the blend; add the kokumi base.
4. Chill; pour over ice and top with sparkling water. The hibiscus brings the tartness that makes the sweetness read clean.

At *S ≈ 1,300* this is the lightest sweetness in the set, and deliberately so — it is built to refresh, not to satisfy. Hibiscus is the engine: its organic acids load the sour channel heavily, and against that bright acid a modest sweetness reads far cleaner and fuller than the same sweetness would in still water. The two receptor systems hold each other in tension — sour sharpening sweet, sweet rounding sour — so that neither needs to be loud. The kokumi base lends just enough continuity to keep the low sweetness from thinning out over ice, and the sparkling water's carbonic bite adds a final, faint acidity that lifts the top of every sip. Nothing in the glass was sweetened with sugar, and nothing in it tastes as though it is doing without.

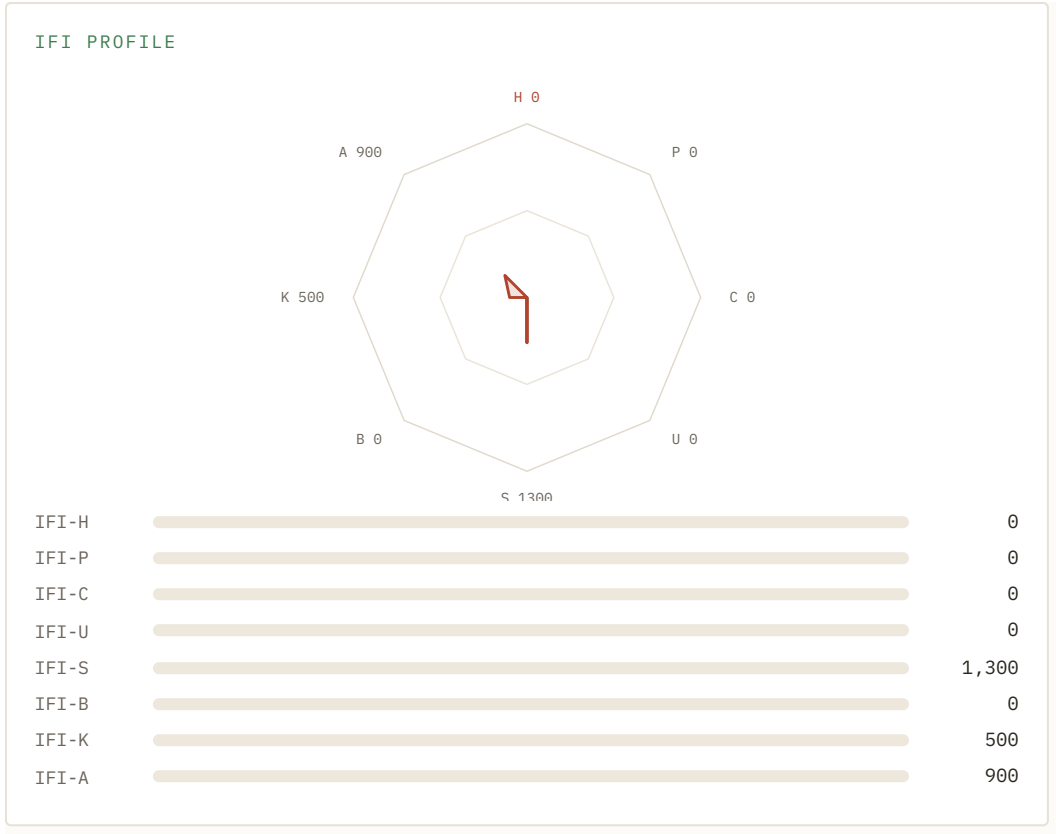

Seven formulations, one method

FORMULATION	TARGET IFI-S	SUCROSE-EQUIV.	FLOOR (K / U)	ACID (A)	HEAT (H)
Citrus Cordial	1,600	16%	700 / 100	600	0
Dark Drinking Chocolate	1,500	15%	1,400 / 300	700	0
Strawberry Sweet-Heat Shrub	1,900	19%	400 / 100	800	700
Clean Pancake Syrup	2,200	22%	900 / 100	100	0
Monk-Fruit & Miso Glaze	1,600	16%	1,500 / 1,400	200	300
Sweet-Heat Glaze, Salmon	1,500	15%	1,100 / 1,200	300	900
Hibiscus-Berry Soda	1,300	13%	500 / 0	900	0

Read down the columns and the method shows itself. Every formulation sets a deliberate target on the sweet axis — from the refresher's 1,300 to the syrup's 2,200 — and not one of them reaches that target with added sugar. Every one builds a floor beneath the sweetness: a kokumi base firing CaSR for continuity, and, where the dish is savory, a bed of umami beneath that. Every one manages an acid around the sweetness, because sweet and sour are read by different machines and so can be composed against each other rather than traded off. And every one reads out as a number you could hand to a factory and expect back, batch after batch, inside ±7%.

That is the proof we most wanted to put on a table: that a sweetness designed from the receptor outward is not a compromise or a substitute. The receptor counts occupancy, not calories; the math

makes occupancy a number; the number makes the sweetness repeatable; and the blend, dosed to a measured target rather than a calculated one, makes it whole. None of these seven leans on sugar for its sweetness, and not one of them tastes as though it is doing without. It is simply sweetness — made on purpose, carried lightly, and given to the body with intention. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 18 · The nourishment through-line — *sweetness, given back.*

We began with milk, and the newborn turning toward sweetness before it had a name for anything. We should end there too, because that infant is the honest center of this whole book. The love of sweetness is not a weakness to be corrected. It is one of the first true things a body knows — *this is good, this will sustain me* — and for nearly all of human history it was reliable, because sweetness meant ripe fruit, meant stored sun, meant food a body could use. The craving was wisdom. It is wisdom still.



Our family grew this work the way you grow anything you love — slowly, with attention, meaning for it to nourish. Sweetness was the place to start, because it is the flavor we loved first.

What changed was never the craving. It was the supply. In a single century we learned to manufacture sweetness's cheapest carrier without limit and pour it into everything, and an ancient, trustworthy signal was turned against the people who feel it — not because they wanted sweetness, but because the only sweetness on offer came bundled, in excess, with a load the body was never asked whether it wanted. The guilt that now rides with sweetness is not the fruit's fault and not the infant's. It is the bundle's. And this entire book has been, from its first page, one patient argument: that the bundle can be undone, because sweetness is a receptor event and not a sugar, so the pleasure and the load were never the same thing and can be separated by anyone willing to measure carefully enough.

That last clause is the whole of the work, and it deserves to end the book not as a promise but as a mechanism. So let us honor the thing that makes it true — slowly, exactly, the way we would show you anything we loved.

The receptor that answers *sweet*

Sweetness is read by a single receptor, and the body owns only one of it. TAS1R2 and TAS1R3 are class C G-protein-coupled receptors — each a long chain folded into a large extracellular **Venus flytrap domain** (two lobes hinged like a clamshell), a cysteine-rich linker, and a seven-pass transmembrane bundle threaded through the cell membrane. Apart, neither tastes of anything. Together they snap into

one heterodimer, and that pairing is the sweet sense entire: every sugar, every diet sweetener, every drop of sweetness a human has ever felt has passed through this one lock.

The keys do not all enter the same door. Sucrose and most small sweeteners — glucose, aspartame, sucralose, the steviol glycosides of the stevia leaf — settle into the cleft of the TAS1R2 flytrap; when the two lobes close around the ligand, the motion telegraphs down the cysteine-rich linker to the transmembrane bundle, which on the inside of the cell lets go of its G-protein (gustducin) and opens the PLCβ2 → IP_3 → TRPM5 cascade that fires the nerve. A few ligands bind elsewhere: cyclamate and the sweet-blocker lactisole dock inside the transmembrane barrel of TAS1R3. And the sweet *proteins* — brazzein, thaumatin, monellin — are far too large to enter the cleft at all; they press against the *outside* of the closed flytrap, a wedge that holds the clamshell shut. Different doors, one hinge, one word delivered to the brain: *sweet*.

Hold on to what each of those events is — and is not. It is a shape fitting a pocket. No bond is broken for energy; no calorie is spent, stored, or even noticed. The receptor cannot tell, and does not care, whether the molecule that closed its flytrap carries four kilocalories per gram or exactly none. This is not a loophole we found. It is the plain architecture of the sense: the tongue reports *how loudly a thing calls to a receptor*, and the body's metabolic load is a wholly separate account, kept by wholly separate machinery. The bundle that modern food tied — sweetness lashed to a sugar load — was always two things wearing one name. The receptor knew it the whole time.

Loudness, not amount

Because sweetness is a binding event, it is measured the way binding is measured: not by *how much* of a compound a food holds, but by the concentration at which that compound drives the receptor to half of its full response. That concentration is the **EC_{50}** — the half-maximal effective concentration — and it is the honest unit of loudness. A low EC_{50} means the receptor answers at a whisper of material; a high one means it takes a shout.

Turn a half-activation into a score and you have the axis this book is written on:

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

We anchor the sweetness axis to sucrose: the reference stimulus is the half-maximal activation of the TAS1R2·R3 heterodimer by ordinary table sugar, which in our cell assay sits near **20 mM**. The anchor is a convention — you choose it once, you hold it forever, and from then on every sweetness on earth can be read against the same zero. A candidate exactly as potent as sucrose scores **1,000**. One that half-activates the receptor at a *lower* concentration — a louder sweetness — scores higher; a quieter one scores lower. The arithmetic is deliberately plain, because a number you cannot check by hand is a number you cannot trust.

The curve underneath the single number is a sigmoid, the standard Hill form:

```
R ⁄ R_max = [L]ⁿ ⁄ ( EC₅₀ⁿ + [L]ⁿ )
```

where [L] is the ligand concentration and *n* the Hill coefficient — the steepness of the climb from no sweetness to full. For the sweet receptor *n* sits near 1 to 1.5, which is why sugar's sweetness arrives as a gentle ramp rather than a switch. EC_{50} fixes *where* that ramp crosses halfway; *n* fixes *how abruptly*. Two sweeteners can share an EC_{50} and still feel different because one climbs its curve steeply and the other leans into it — a distinction the single score compresses, and one a careful formulator restores by hand.

A score is only as good as its repeatability, so we keep one more number beside every EC_{50}: the batch-to-batch coefficient of variation. A wild botanical extract wanders — lot to lot, a leaf-derived sweetener can swing with a CV near **42%**, which is to say a "900" one week is a 1,278 the next and a 522 the week after. No one can lock a recipe to a number that breathes that hard. The designed route holds its delivered sweetness to a CV of **≤7%** — a 900 that stays between roughly 837 and 963, lot after lot. That tightness is not a flourish. It is the difference between a sweetness you can build a product on and a sweetness you can only hope for. Precision, here, is a form of respect.

One worked reading, so the machinery is not abstract. Our full-spectrum cultivar **Phoenix** carries a computed sweetness score of **S = 1,100**. Invert the index — EC_{50}(candidate) = EC_{50}(reference) × 1,000 ÷ IFI-S — and its sweet ligand half-activates the receptor at 20 mM × 1,000 ⁄ 1,100 = **18.2 mM**, a touch under sucrose's 20: about ten percent louder than table sugar, and no more. A quiet, true sweetness, sitting on a plant built mostly for other things. The number does not flatter. That is exactly why we keep it.

AXIS	IFI (PHOENIX)	EC_{50}(C) ⁄ EC_{50}(REF) = 1000 ÷ IFI	LOUDNESS VS. REFERENCE
Heat (H)	4,200	0.238	4.2×
Zero-Bitter (B)	2,200	0.455	2.2×
Kokumi (K)	1,900	0.526	1.9×
Umami (U)	1,800	0.556	1.8×
Astringency (A)	1,400	0.714	1.4×
Sweetness (S)	1,100	0.909	1.1×
Cool-Burn (C)	800	1.250	0.8×
Pungency (P)	252	3.97	0.25×

Eight numbers, one plant

There is a plant on which the bundle is not an argument but a fact you can taste. Phoenix is the first cultivar our family designed to register on all eight axes at once, and its profile is the plainest demonstration we know that a flavor is a vector of independent signals, each with its own receptor and its own EC_{50}, never a single lump labeled "taste."

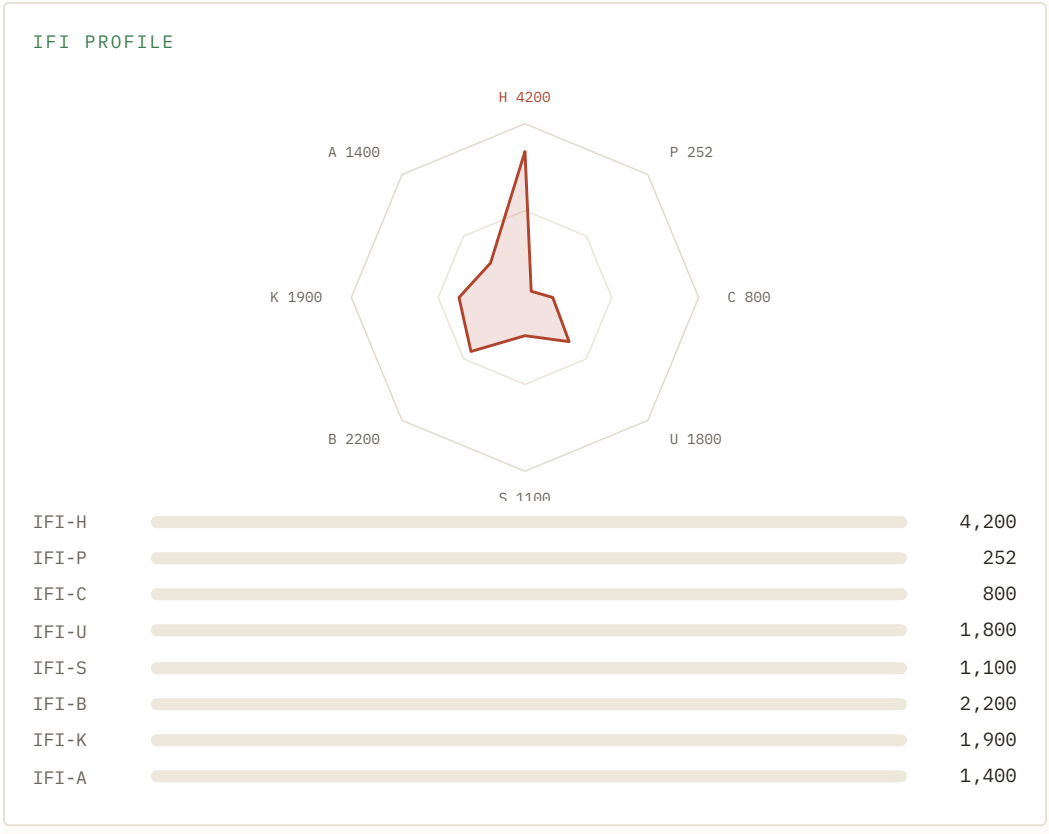

Read it as eight separate truths. The heat is a roar — TRPV1 driven at a quarter of the reference concentration. The sweetness is a murmur beside it, 1,100, ten percent above sugar. They do not interfere; they do not trade; the sweet signal is not "cancelled" by the heat any more than a cello is cancelled by a trumpet. Each axis is specified alone, written alone, and can be moved alone. A flavor, under the Index, is not a mood. It is a chord, and every string is tunable by hand.

That independence is the whole reason sweetness can be given back clean, and our cultivar design progression shows where it enters. Climb the ten proof-of-concept cultivars and the design scope grows — more axes activated, more flavor pathways wired — until, at the top, all eight light at once and sweetness finally takes its seat.

CULTIVAR · SKU	PREDICTED SHU	IFI-H	AXES
Ember · SS-0087	866K	1,570	2/8
Blaze · SS-0091	1.72M	2,630	4/8
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

The sweetness we are giving back is the one that arrives at the top of that progression — not borrowed from sugar, but engineered as its own line in the chord. The rest of this chapter is four ways we are learning to play it.

The tomato — summer, with the sugar left out



A tomato in cross-section — the chambers and gel where sweetness and savor are kept, and where the volatiles that make a tomato taste of summer are made or lost.

Ask a person what a summer tomato tastes of and they will say *sweet*, and they will be right, and the sugar will have almost nothing to do with it. This is the first surprise of the tomato, and it changed how we think about the whole axis. Perceived sweetness in a tomato is set only partly at the tongue; a great deal of it is assembled in the brain, which folds the retronasal aroma of the fruit — a suite of volatiles rising off the pulp as you chew — into the single percept it calls "sweet." Certain of those volatiles raise perceived sweetness with no change in sugar at all. The brain tastes the smell of ripeness and reads it as more sugar than is there.

Modern commercial tomatoes, bred for yield and shelf life and the long cold ride in the truck, lost most of those volatiles — the aromatic chemistry was simply not what anyone was selecting for. What was left was a fruit with its sugar roughly intact and its *sense* of sweetness stripped out. Across the 2,040 constructs we have built in the tomato, the central lever is not to pour in sugar but to grow the lost volatiles back through the fruit's own native routes, and then to lift the receptor-level sweetness only modestly on top. That is the honest meaning of a tomato that tastes of summer with barely more sugar than before: most of the sweetness restored is sweetness the brain was always willing to assemble, if only the aroma arrived to prompt it.

This is also where we keep ourselves honest, because the two sweetnesses are not the same measurement and must never be reported as one. The receptor sweetness — IFI-S — is read in a cell assay and in a tongue; the volatile-borne sweetness is read only in people, on a sensory panel, because no heterologous receptor will ever smell a tomato. So a restored tomato carries *two* numbers: a modest receptor IFI-S, and a separate perceptual lift from aroma, and we do not let the one flatter the other. The gain is real and large. It is also, in large part, perceptual — and a book that respects its reader says so, rather than banking an aroma effect as if it were loudness at the lock.

Walking the licorice out of the pathway

The stevia leaf makes a sweetness that is genuinely free of the sugar load, and it has one old, famous flaw: the licorice-bitter tail that rides behind the sweet. That tail is not a mystery and not a matter of taste. It is a second receptor being struck. The steviol glycosides are built in the leaf from a diterpene backbone — the MEP pathway yields geranylgeranyl diphosphate, which cyclizes to *ent*-kaurene and oxidizes to steviol — and then a series of UDP-glycosyltransferases hangs glucose units onto that core, one at a time. UGT85C2 opens the series; UGT74G1 and UGT91D2 build it out to **stevioside**; UGT76G1 adds a further glucose to make **rebaudioside A**; later enzymes carry it on to **Reb D** and **Reb M**, the most fully glucosylated of the family.

Here is the quiet elegance of it: where the glucose sits decides which receptor the molecule rings. Stevioside and the lower glycosides activate not only the sweet receptor but the bitter receptors **TAS2R4** and **TAS2R14** — that is the licorice tail, a bitter chord struck alongside the sweet one. Add glucose, and the shape that reaches the bitter receptors is shielded; Reb A is cleaner than stevioside, and Reb M and Reb D are cleaner still, nearly all sweet and almost no bitter. So "walking the licorice out of the pathway" is literal engineering: across the 910 constructs in the stevia program, we steer the UGT flux toward the fully glucosylated end — more Reb M and Reb D, less stevioside — which raises the sweetness the receptor hears *and* starves the bitter note of the molecules that carried it. The cleanup shows up on two axes at once: sweetness up, bitterness attenuated.

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

The same logic runs through the rest of the sugar-free sweeteners. Monk fruit carries its sweetness in the mogrosides — cucurbitane triterpene glycosides, with mogroside V (five glucoses) the sweet principle at roughly 250 times the weight-sweetness of sucrose and far less bitter than a crude stevia leaf. Sugarcane, which the world knows as the sugar plant, we work as a sweetness chassis rather than a sugar factory, which is why it sits at the frontier of difficulty with only 200 constructs so far: teaching the cane to hold the sweet signal without the sugar load is a longer road than tuning a leaf that already makes a non-caloric glycoside. The destination is the same in every case — the sweetness kept, the thing bundled to it walked out, molecule by molecule, down a pathway we can read.

A word on how little material this takes, because it is the whole point and it is easy to miss. Reb M runs about 300 times as sweet as sucrose by weight, but the receptor does not weigh — it counts molecules, and the two are not the same. Convert the weight potency to a molar one by the ratio of molecular weights: Reb M's mass is near 1,291 g/mol against sucrose's 342, so its molar potency is $300 \times (1{,}291 / 342) \approx$ **1,130 times** that of sugar. Against our 20 mM sucrose anchor, Reb M half-activates the sweet receptor at about **18 µM** — eighteen-*thousandths* of a millimole. The pure compound's

intrinsic loudness is off the top of any scorecard; this is why the axis is normalized to a ceiling of 5,000 for display, and why the honest sweetness of a *food* is not the compound's intrinsic score but the loudness it actually delivers at the dose you eat. The intrinsic potency tells you how little to use. The delivered IFI-S tells you what the mouth receives.

The sweetest things we can make

Past the glycosides lies the loudest sweetness biology has — the sweet proteins — and they are the clearest reason this book ends with a frontier rather than a finish line. Brazzein is a protein of fifty-four amino acids, about 6.5 kilodaltons, from the fruit of a West African vine; thaumatin and monellin are its larger cousins. They do not enter the receptor's cleft; they are the wedge, binding the outer face of the closed flytrap and holding it shut. And their loudness, reckoned honestly, is astonishing.

Walk the unit conversion, because it is the whole wonder of the thing. Brazzein is cited at roughly 500 to 2,000 times the weight-sweetness of sucrose, flattening toward the lower end as concentration rises — sweet proteins plateau early on their curve. Take the threshold figure, 2,000× by weight, and convert to molar potency by the mass ratio: $2{,}000 \times (6{,}473 / 342) \approx$ **37,800 times** sucrose, molecule for molecule. One folded brazzein does the sweet work of nearly forty thousand molecules of sugar. That is the quietest sweetness imaginable in terms of material and the loudest in terms of signal — a receptor rung almost to the top by a trace.

So why are the sweet proteins, in the chapter's own earlier words, not yet in the fields? Because loudness is not the hard part; folding is. A fifty-four-residue protein is sweet only when its disulfide bonds are formed correctly and its shape is held; express it in the wrong tissue, at the wrong level, without the chaperones to fold it, and you get an abundance of a protein that does not taste of anything. Scaling a correctly folded sweet protein to meaningful levels in an edible crop, holding its stability through harvest and processing, and walking the long, honest regulatory road that any novel dietary protein must walk — those are the real work, and they are not done. We say it plainly because the alternative is to promise a field that does not yet exist. The sweetness is real in the molecule. The plant that carries it well is still ahead of us.

Built back up from a number to a floor

Nowhere is the separation of sweetness from sugar more concrete than in a glass of soda, so let us build one by the numbers. A conventional cola holds roughly 100 grams of sucrose per liter — about 10 grams in a 100-milliliter pour. In molar terms that is $100 / 342 \approx$ **0.29 mol/L**, or 290 mM of sugar, which against a 20 mM half-activation sits far up the Hill curve: $290 / (20 + 290) \approx 94\%$ of the receptor's full response. A full-sugar soda is engineered, in effect, to park the sweet receptor near saturation. It reads *full* because it is.

Now hold the sweetness and drop the load. Specify the target not as a sugar weight but as a floor — the minimum delivered sweetness the drink must never fall below, across its shelf life and its temperature range — and build up to that floor with a designed

sweetener instead of sugar. Match the cola's sweetness with Reb M at 300× by weight: 100 g/L ⁄ 300 ≈ **0.33 g/L**, which at 1,291 g/mol is about **260 µM**. The swap is stark. Two hundred ninety *millimolar* of sugar becomes two hundred sixty *micromolar* of clean glycoside — a ligand count smaller by a factor near 1,100 — and the caloric sugar load falls to essentially zero while the receptor still reads near the same full-throated 94%. The sweetness floor is held. The thing bundled to it is simply gone. That is what it means to build a sweetness back up from a number to a floor: you name the loudness you owe the mouth, and you deliver exactly that, and nothing else rides along.

Honesty lives in the edge cases, and they are real. A high-potency sweetener does not share sugar's temporal shape: its sweetness can arrive a beat late and linger a beat long, and at these concentrations a faint bitter or licorice tail can surface that sugar's bulk once masked. Sugar also does work in the mouth that has nothing to do with sweetness — body, viscosity, a rounding fullness — which 260 micromolar of anything cannot supply. So a soda built this way is never a single sweetener; it is a composition. A trace of a second, faster sweetener repairs the onset. A kokumi element working through the CaSR receptor restores the continuity and body the sugar used to carry. And the whole thing is specified against that ≤7% CV, because a floor that the formula cannot reliably hit, lot after lot, is not a floor at all. The math gets you most of the way. The last of the way is taste, held to a number.

The strawberry that finishes warm



A strawberry — the fruit the body already trusts, and the one we are teaching to open sweet and finish warm.

There is one more sweetness we are reaching for, and it is the one that makes people lean in: a strawberry that opens sweet and finishes warm. It belongs to the cross-family frontier — a longer build, seven to ten genes, three to four years — because a strawberry has no heat chemistry of its own. The warmth has to be brought in. The phenylpropanoid branch that every vascular plant already runs — PAL to C4H to 4CL to HCT to COMT — supplies the backbone, and onto it the four terminal pepper-specific enzymes are transplanted: pAMT, Pun1/CS, KAS, BCAT, the last short pathway that turns a plant's ordinary aromatic metabolism into capsaicin.

Why it finishes warm rather than fights the sweet is the loveliest piece of receptor timing in the whole program. Sweetness and heat are read by different receptors with different speeds. TAS1R2·R3

answers almost at once — the sugar sense is fast, front-of-mouth, the first thing a bite says. TRPV1, the heat channel, is slower to gather and slower to fade; capsaicin is lipophilic, it loads into the tissue and releases its signal over a longer arc. Put a fast sweetness and a slow warmth in the same fruit and they do not collide — they *sequence*. The strawberry opens on the sweet it has always had, and the heat rises underneath it as the sweetness settles, and the finish is warm. The chord is written across time, not just across axes.

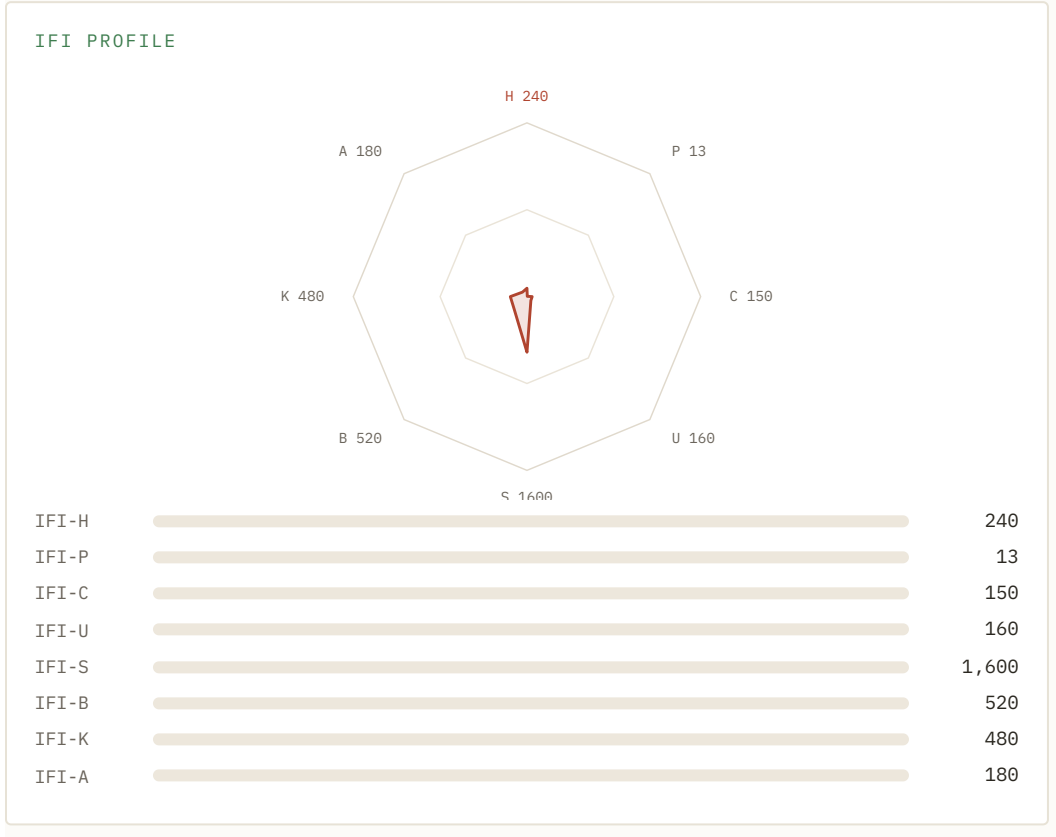

We print that profile the way we mean it — as a design target, the shape we are breeding toward, and not yet a field measurement. The strawberry does not yet bite back. But the target is honest and the receptors are willing, and the arithmetic of their timing is already on our side: a loud, front-loaded sweetness near 1,600 and a gentle, dessert-scale warmth near 240, arriving in that order because the channels that carry them keep that order on their own.

What we do not claim

We will not tell you this is medicine, because it is not, and this library does not make that kind of promise. No sweetness in this book cures anything, detoxes anything, or carries a health halo it has not earned — those are not our words and never will be. What the Index measures is loudness at a receptor; what the body does with the molecule afterward is a separate account, and we refuse to blur the two in either direction. The clean sweeteners drop the sugar load; that is a fact about the molecule, not a claim about a life.

And the limits are real, so we keep them in plain sight. A tomato's restored sweetness is in large part perceptual, read in people and never in a cell, and we report it as its own number rather than folding it into the receptor score. The stevia cleanup is a shift along a pathway, not a perfect erasure — the licorice is walked out, not teleported away. The sweet proteins are loud in the molecule and not yet at home in a field. A high-potency sweetener's timing and body must be composed back by hand. These are not caveats we were forced into; they are the shape of honest work, and naming them is how a reader knows the rest of the numbers are true.

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

What we will tell you is what we believe, which is simpler and older: that food should give the body pleasure and nourishment in the same bite, with nothing to apologize for; that a plant can be grown — vegetarian-first, clean, reverent of the living thing — to hold the most of both it can carry; and that designing flavor and designing nourishment are, in the end, the same careful act.

Sweetness, given back

Sweetness was where we chose to begin that work, because it is the flavor we loved first and have been made to trust least — and because giving it back, clean, felt like the truest thing we could do with a receptor, a pathway, and a family's attention. The computing was one pair of hands; the table, and the reason for all of it, belong to all of us.

There is a great deal still ahead. The sweet proteins are not yet in the fields; the strawberry does not yet bite back; the nourishment this library is named for has volumes of its own still to come. But the direction is set, and it is the direction of every plant our family has ever touched. We measure what we love, because a number is a form of respect. We walk the pathway forward until the flaw is no longer made — the licorice out, the load off, the signal kept. And we grow the thing toward the fullest flavor, and the fullest nourishment, it can hold.

Sweetness, measured and given back, is only the beginning of that. It is the oldest thing we taste, the first *this is good* a body ever knew — and we intend to see it loved again without a shadow on it.