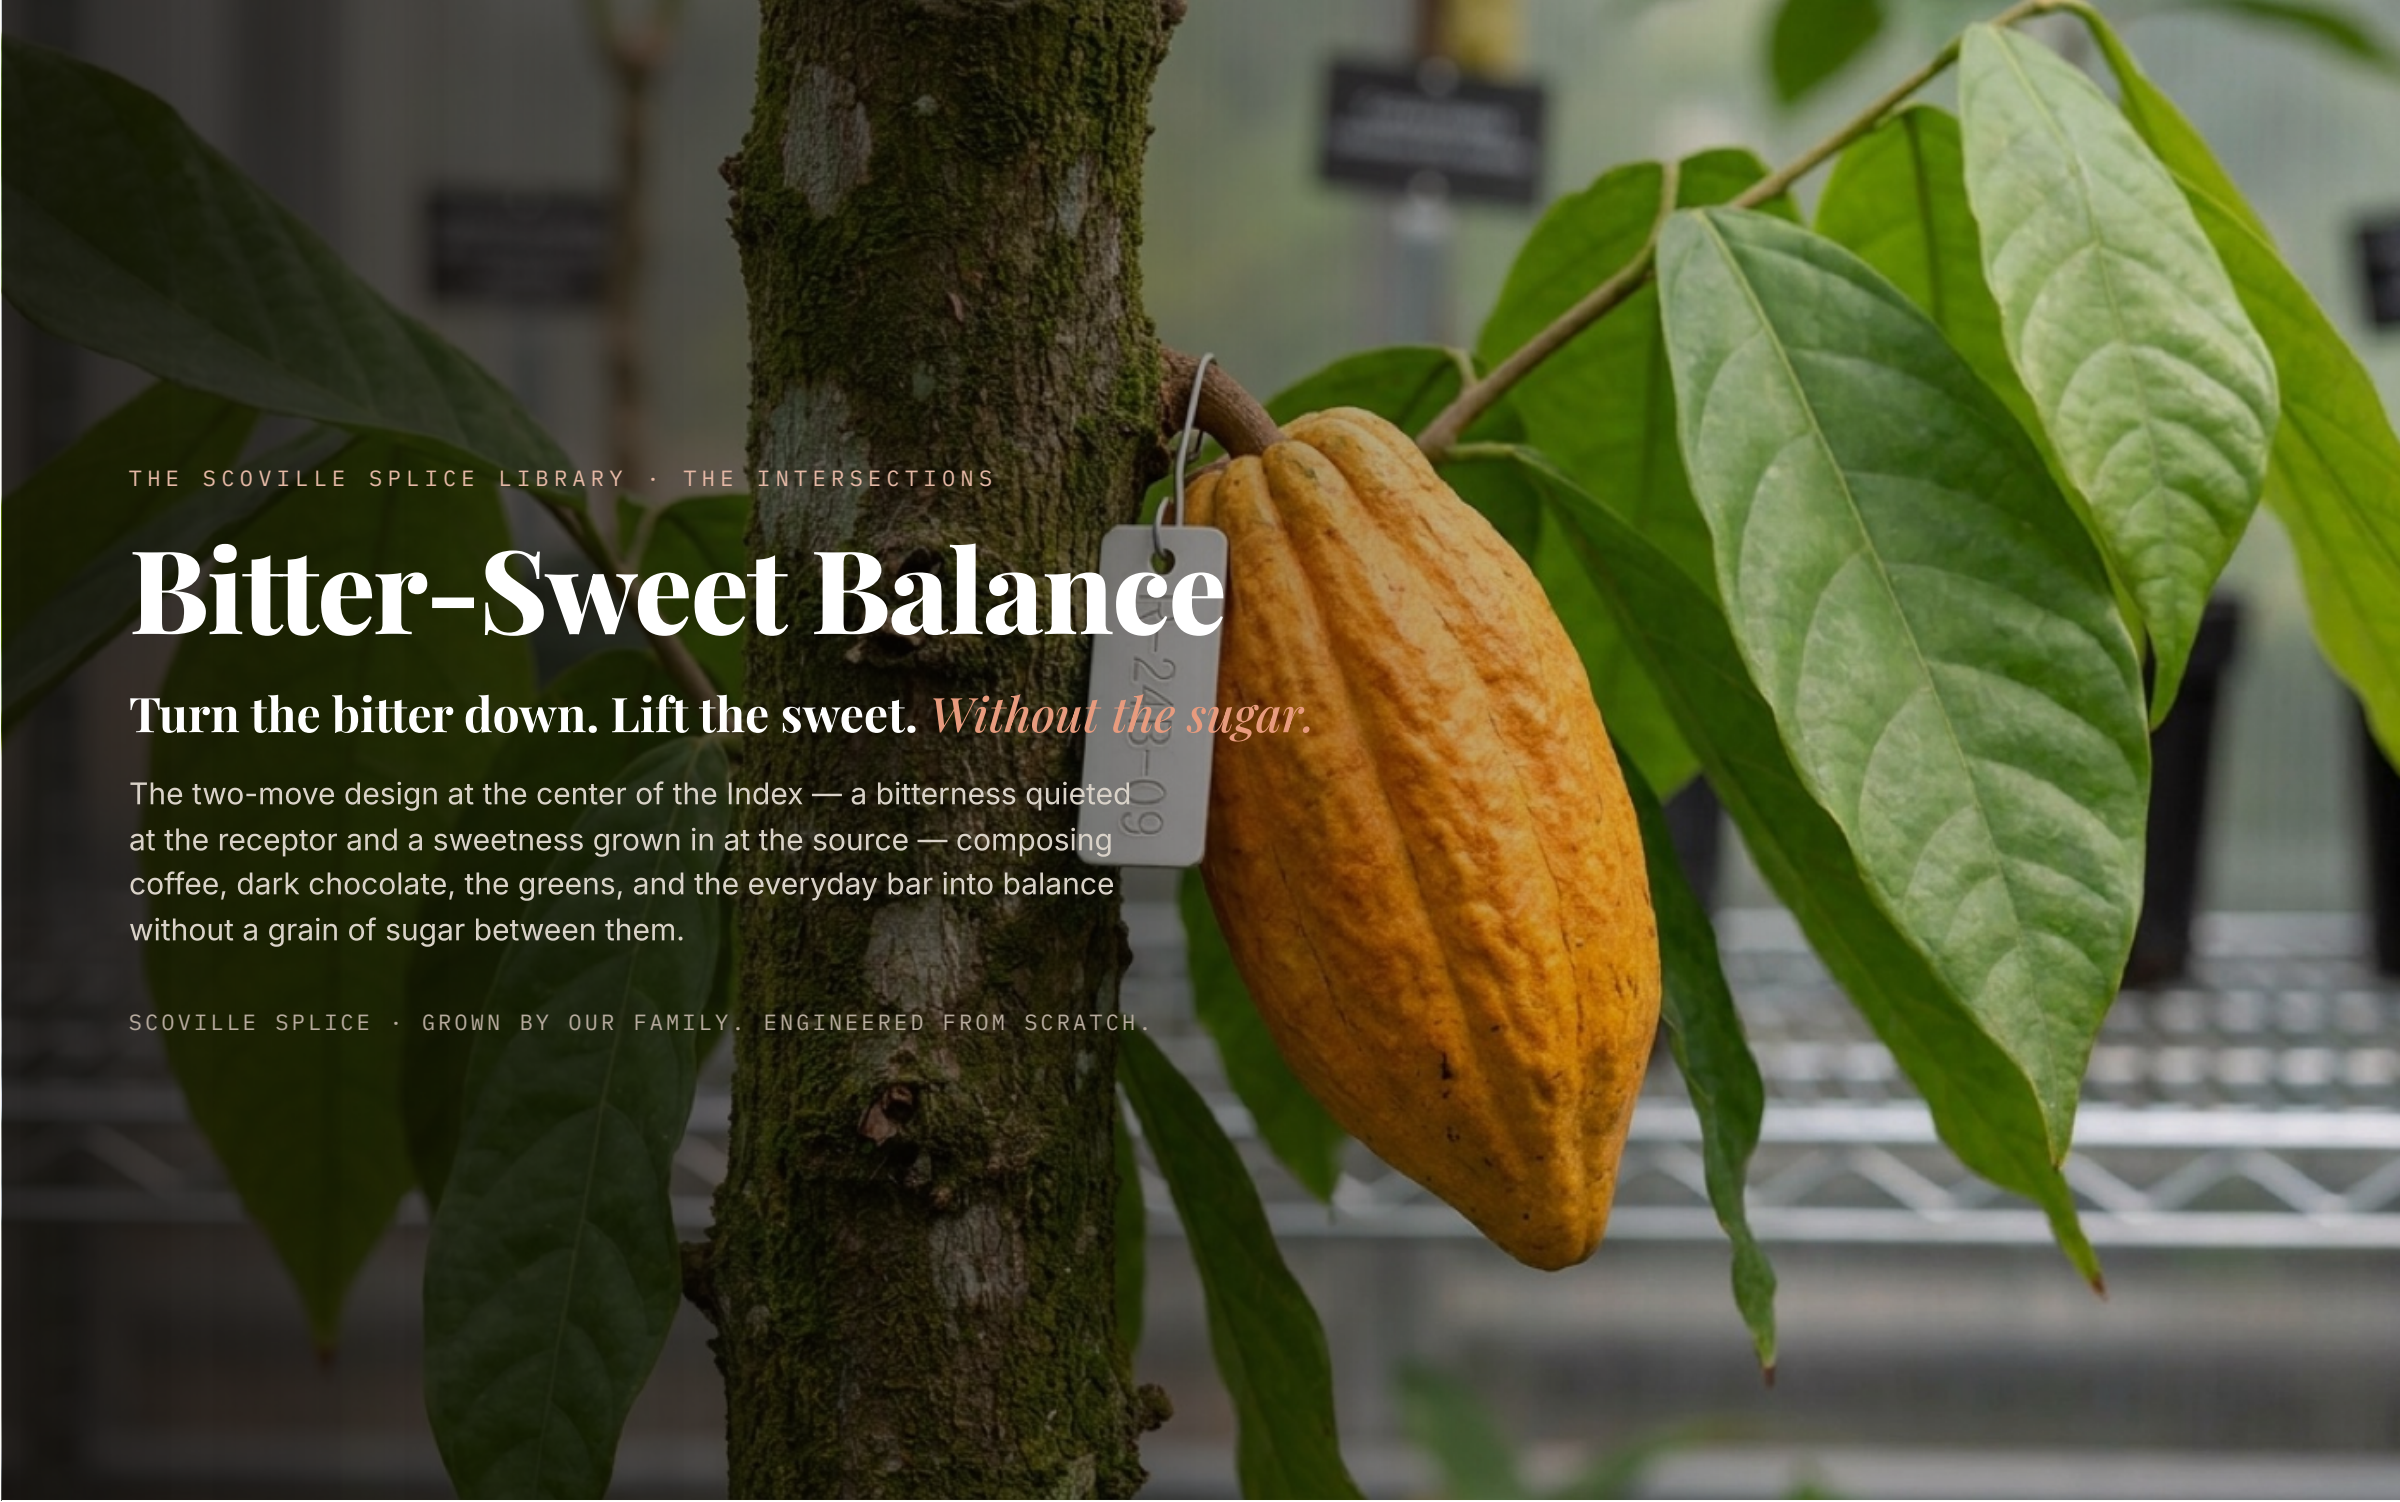

THE SCOVILLE SPLICE LIBRARY · THE INTERSECTIONS

Bitter-Sweet Balance

Turn the bitter down. Lift the sweet. *Without the sugar.*

The two-move design at the center of the Index — a bitterness quieted at the receptor and a sweetness grown in at the source — composing coffee, dark chocolate, the greens, and the everyday bar into balance without a grain of sugar between them.

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

Contents

01 Opening essay · The quarrel we settled with sugar

02 Chapter 1 · Two Tastes That Argue — the oldest balance in cooking.

03 Chapter 2 · The Alarm and the Invitation — two receptors, read side by side.

04 Chapter 3 · The Arithmetic of a Settled Argument — two numbers, one honest line.

05 Chapter 4 · The Sugar Tax — what a bitterness costs in sweetness.

06 Chapter 5 · The Two-Move Design — the balance you compose, not mask.

07 Chapter 6 · The Design Loop — scoring an edit before we cut.

08 Chapter 7 · The Two Bitternesses of Coffee — lactones, phenylindanes, and the roast.

09 Chapter 8 · Coffee, Read Black — the cup that asks for nothing.

10 Chapter 9 · Theobromine and the Flavan-3-ols — cacao, where bitter meets astringency.

11 Chapter 10 · Sweetness Against the Dark — chocolate, settled without sugar.

12 Chapter 11 · The Everyday Bar — flavor that survives the formula.

13 Chapter 12 · Greens That Read Sweet — the vegetable a child finishes.

14 Chapter 13 · The Bitterness We Keep — tea, hops, grape, and the useful edge.

15 Chapter 14 · The Sweet Keys — sweetness that behaves.

16 Chapter 15 · Building the Floor — what holds a sugar-free sweetness up.

17 Chapter 16 · The Reformulator's Playbook — from a sugar brief to two numbers.

18 Chapter 17 · One Fruit, in Balance — reading Phoenix's bitter and sweet.

19 Chapter 18 · Reference Readouts — the bitter-sweet shelf, scored.

20 Chapter 19 · Measured, Not Claimed — two assays, one discipline.

21 Closing essay · What balance was always for

Opening essay · The quarrel we settled with sugar

The first thing most of us ever sweetened was something bitter. A grandmother's coffee taken with two sugars and no comment. Cocoa cut with enough sugar to coax a child to the bottom of the mug. Grapefruit under a drift of it, a dark tea gone pale with honey. For as long as people have cooked, the cure for a bitterness has been the same blunt mercy — pour sweetness over it until the warning goes quiet. It works. It has always worked. And it has cost us more than anyone ever wrote on the bill, because a world that sweetens everything it finds too sharp ends up sweetened all the way through, and forgets that it had a choice.

We grew up at a table that loved bitter things — strong coffee before anyone else was awake, dark greens, the back-of-the-tongue depth of a good cup of tea, a square of chocolate so dark it argued back. And we noticed early that most of the people around us were not really tasting those things. They were tasting the sugar they had buried them under. Bitterness and sweetness are the two oldest adversaries on the tongue, and almost the whole history of making the first one bearable has been the history of drowning it in the second. That is not balance. That is surrender, dressed up as dessert, and it leaves the thing you buried exactly where it was — still bitter, now only hidden, under a load the body carries long after the taste has done its work.



The question that became this book came up at our table one night: if the cure for bitterness is sweetness, why must the sweetness arrive as sugar — and why must the bitterness arrive so loud in the first place?

This book is about the other way, and the other way begins where every book in this library begins — with the receptors, because a taste is not a property of a food but an event in a protein. Bitterness lives on a family of sentinels the Index anchors to **TAS2R38**, the alarm evolution hung on the tongue to keep our ancestors off the poison leaf. Sweetness lives on an entirely different machine, the **TAS1R2·TAS1R3** pair — two proteins that must fold together, like two cupped hands closing on a coin, before the brain will hear the word *sweet.* They are separate locks, read by separate keys, built on different architectures in different parts of the membrane. For ten thousand years no one could touch either one, so we reached for the only lever we had, which was sugar, and we pulled it on everything that was ever too sharp.

We can touch them now, each one directly, and the whole of this book is what happens when you stop pulling the one crude lever and start working the two real ones. On the Iyengar Flavor Index, **IFI-B** measures how loudly a plant speaks its bitterness, and **IFI-S** measures how loudly it speaks its sweetness, and underneath both is the one honest line that runs through this entire library:

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

EC_{50} is the concentration at which a compound drives its receptor to half of a full response — the physical measure of how loudly a molecule speaks. The two moves of this book are a direct consequence of that line. First, we quiet the bitterness at its source — rerouting, by the smallest possible edit, the pathways that build the harsh molecules, so the plant makes less of what the sentinel was sounding the alarm about; the reading falls, and the alarm with it. Second, we lift the sweetness at the same source — not by adding sugar but by growing more of the plant's own loud sweet signal, the steviol glycosides and the sweet proteins that speak to TAS1R2·TAS1R3 hundreds of times more loudly, gram for gram, than cane sugar ever did. Two receptors. Two moves. Two numbers. And between them, a balance you *compose* on purpose rather than *mask* by the spoonful.

That is the whole argument, and it is a short one, though the chapters that make it good are long. You do not have to sweeten a bitterness to forgive it. You can turn the bitterness down and grow the sweetness in, at the plant, before anything ever reaches the cup — and arrive at the same balance your grandmother chased across a lifetime of two-sugar coffees, carrying none of the load she needed to get there. The quarrel between bitter and sweet is the oldest one in cooking. This is the book in which, receptor by receptor and pathway by pathway, we finally settle it.



PART I OF VI

The Two Tastes That Argue

mixture suppression · TAS2R38 and TAS1R2·TAS1R3 · the arithmetic of a settled argument

Chapter 1 · Two Tastes That Argue — *the oldest balance in cooking.*

IFI-B × IFI-S · mixture suppression · the mechanism hidden in the quarrel

Bitter and sweet do not merely sit beside each other on the tongue. They fight — and the fight, it turns out, is the mechanism this whole book is built on. Drop sugar into a harsh black coffee and the bitterness genuinely recedes: not because the bitter compounds have gone anywhere, for they are all still there, still binding their receptors exactly as before, but because sweetness and bitterness *suppress one another.* Raise one and the other is heard less. Every cook who has ever balanced a too-bitter sauce with a pinch of sugar, every barista who has pulled a shot sweet to tame its edge, every grandmother with her two sugars, has been working this suppression without a name for it. It is the oldest balance in cooking, and until now it has only ever been driven in one direction, with one tool.

The suppression is real, it is measurable, and it lives in two places at once. The first is peripheral, out at the taste cells themselves. Recall that the sweet, bitter, and umami receptors, for all that they read different molecules, pour their signals into the *same* downstream machinery inside the taste cell — the same gustducin, the same phospholipase, the same calcium store and the same TRPM5 channel that finally shouts to the nerve. When two of those receptor types fire at once in neighboring cells, their signals compete for shared transduction capacity and for the brain's attention; a strong bitter input can tamp down the apparent gain of a sweet one, and a strong sweet input can do the same to bitter. The second place is central, in the brain's own integration of taste, where the gustatory cortex sums the incoming lines and a loud signal on one lets it hush a quieter signal on another. Psychophysicists have mapped this for a century: present sweetness and bitterness together and each is rated less intense than it would be alone, along a curve steep enough to be useful and reproducible enough to be designed against. The two tastes were built by evolution to argue. We are simply going to referee the argument.



The two tastes converge in the same taste cell and the same cortex — which is why a sweetness can quiet a bitterness it never touched, and why the whole quarrel is designable rather than merely endurable.

For almost the whole history of the kitchen, that referee had exactly one instrument, and he used it one way. To quiet a bitterness, add sweetness; to add sweetness, add sugar; and because the suppression is roughly proportional, a *loud* bitterness demands a *great deal* of sweetness, which means a great deal of sugar. This is the arithmetic of a two-sugar coffee stated plainly. The sugar in that cup is not mostly there to taste sweet. It is there to spend itself suppressing a harshness, and only the fraction left over after that work is done reads as sweetness at all. We will put a number on that hidden tax in Chapter 4; for now it is enough to see the shape of the old deal. One lever, pulled hard, pouring sugar into a food to buy a suppression the food's own chemistry was fighting the whole time. It is a crude tool not because it fails — it succeeds, reliably, which is why we have used it for ten thousand years — but because it is *expensive,* and the expense is paid in the one currency the body carries least lightly.

The two-move design keeps the suppression and discards the sugar, and it does so by working both sides of the quarrel at once instead of only one. We go to the bitter receptor's source and lower the compound it reads, so the plant's bitterness is quieter *before any sweetness is added at all* — the alarm turned down at the gene rather than shouted over at the table. Then we lift the sweetness, not with sucrose but with the plant's own loud sweet molecules, and that lifted sweetness does what sweetness has always done: it leans on whatever bitterness remains and presses it quieter still. Two interventions, each carrying half the work, reinforcing one another exactly the way the tongue's own wiring already wanted them to. And here is the quiet multiplication that makes the whole book worth writing: because the two moves compound, you need far less of *either* than the old single-lever cure ever asked for. A bitterness already lowered at the source needs only a little sweetness to feel balanced; a sweetness grown in rather than poured on needs no sugar to deliver. The expensive deal becomes a cheap one, and the food keeps its character instead of surrendering it.

There is a deeper reason this particular intersection is the gentlest and most natural in the library, and it is worth naming because it governs everything ahead. Some axes must be held carefully *apart* because they crowd each other — heat and sweetness fight for the front of the palate, which is the whole tension of *Swicy* (Volume

XXVI). Some must be braided *together* because neither is complete without the other — umami and kokumi, the subject of *Savory Depth* (Volume XXVII), where the one gives a broth its depth and the other its length. Bitter and sweet are different from both. They are two tastes evolution built to oppose, which means that when you tune both at once you are not fighting the biology — you are *using* it. The opposition itself becomes the lever. You are not silencing one voice to let the other through. You are teaching two voices, already inclined to settle, to finally do it. That is the balance a cook has always chased in the bowl, and the achievement of this book is that it can now be composed in the plant, before the bowl is ever reached — quieter, cleaner, and carrying nothing it did not mean to carry.

Chapter 2 · The Alarm and the Invitation — *two receptors, read side by side.*

TAS2R38 (class A) · TAS1R2·TAS1R3 (class C) · the shared wire behind both

To design a balance between two tastes you must know both locks, and the two locks at the heart of this book could hardly be more different in build. One is an alarm, broadly tuned, old, and suspicious; the other is an invitation, exquisitely specific and ancient in a different way. They sit on the same tongue, often in the same mouthful, and understanding where they differ — and the one surprising place they converge — is the foundation of every edit in the chapters ahead.

Begin with the alarm. **TAS2R38** is a class A, rhodopsin-like G-protein-coupled receptor — seven membrane-spanning helices threaded through the surface of a bitter-sensing taste cell, with only a short external tail. A bitter molecule does not dock on some great outer platform; it slips *into* the helical bundle itself, settling into a pocket lined by transmembrane helices three, five, six, and seven, and from inside the membrane it levers the receptor's shape just enough to wake the machinery below. The pocket is roomy and forgiving, which is exactly what an alarm needs, because the number of bitter poisons a plant or a mold might invent is effectively unbounded and no narrow sentinel could cover them. This is why bitterness, alone among the tastes, is read not by one receptor but by a *family* of twenty-five — the TAS2Rs — a committee of sentinels where sweetness makes do with a single pair. TAS2R38 is the one we anchor the bitter axis to, because it is the best-characterized of the family and the one most sharply tuned to the sulphurous edge of the brassicas and to PTC, the famous probe compound; the full account of the choir and the anchor is the opening of *Zero-Bitter* (Volume IX). It carries, too, the most famous polymorphism in all of taste: the common **PAV** haplotype (a taster form) and the **AVI** haplotype (a non-taster form), differing at just three amino-acid positions, which split a dinner table into people who taste a given bitterness sharply and people who barely taste it at all. Hold that fact; it will decide, in Chapter 12, whose tongue a bitterness number is really written for.



A broccoli floret in cross-section — the glucosinolate chemistry that trips the bitter alarm is laid down in exactly the tissue we are asked to eat, which is the whole problem the alarm half of this book exists to solve.

Now the invitation, and it is a different machine entirely. Sweetness is read by **TAS1R2 and TAS1R3**, two proteins of the class C GPCR family — the same architectural clan that reads glutamate in your synapses and calcium in your blood. Each subunit is built in three stacked pieces: at the top, reaching out into the saliva, a large bilobed **Venus flytrap domain**, two lobes hinged open and waiting; below it a short, rigid **cysteine-rich domain** that acts as a driveshaft; and anchored in the membrane the seven-transmembrane bundle. The two subunits are an obligate heterodimer — each is deaf alone, and only clasped together do they taste. A sugar drops into the open cleft of the TAS1R2 flytrap and triggers it to close, and that mechanical snap travels down the driveshaft into the membrane and sets the signal loose. But — and this is the fact that makes the sweet axis so designable — the flytrap is not the only door. A second site lives deep in the 7TM domain of TAS1R3, where cyclamate binds and where lactisole wedges in to *jam* the machine shut. And a third mode is the strangest and most useful: the sweet proteins — brazzein, thaumatin, monellin — are far too large to drop into the cleft, so they dock against the *outside* of the closed flytrap like a wedge jammed under a door, holding the active shape shut from without. A protein a thousand times too big to be a key turns out to be a better one than sugar, because it does not have to be caught and re-caught; it simply holds the trap closed. This is why, per molecule, a sweet protein is tens of thousands of times louder than sucrose, and it is the deep structural reason sweetness can be uncoupled from sugar at all. The full architecture is the first chapter of *Sweetness* (Volume VII).

PROPERTY	TAS2R38 (BITTER)	TAS1R2·TAS1R3 (SWEET)
GPCR class	A (rhodopsin-like)	C (glutamate-like)
Arrangement	single receptor, one of 25	obligate heterodimer
Ligand site	pocket inside the 7TM bundle	VFT cleft, VFT/CRD surface, 7TM pocket
Tuning	broad — a whole family of sentinels	specific — one receptor pair
Signature variation	PAV / AVI haplotype (taster / non-taster)	residues that let humans taste aspartame, the proteins
On the Index	**IFI-B** (anchored to TAS2R38)	**IFI-S** (anchored to 10% sucrose)

So far the two locks look wholly separate, two different machines reading two different molecules — and at the level of the receptor they are. But descend one step, past the lock into the cell behind it, and the two tastes abruptly share a single hallway, and this convergence is the physical basis of everything in Chapter 1. When TAS2R38 catches a bitter molecule, it engages a taste-specific G protein built around **α-gustducin**; the βγ pair gustducin releases switches on **phospholipase C-β2 (PLCβ2)**, which cleaves a membrane lipid into **inositol trisphosphate (IP_3)**; IP_3 opens **IP_3R3** on the cell's internal calcium store; calcium floods out; the rising calcium opens **TRPM5**, a cation channel; sodium pours in, the cell depolarizes, and the voltage swing opens **CALHM1/CALHM3**, a large channel that releases **ATP** directly onto the waiting nerve, which reads it through **P2X2 and P2X3** receptors and carries the word *bitter* to the brain. Now read the sweet cascade: receptor → gustducin → PLCβ2 → IP_3 → IP_3R3 → calcium → TRPM5 → CALHM1/3 → ATP → P2X2/3. It is, link for link, the *same chain.* The front doors differ; the hallway is one.

```
bitter or sweet → (its receptor) → gustducin → PLCβ2 → IP₃ → IP₃R3 → Ca²⁺ ↑ → TRPM5 → CALHM1/3 → ATP → afferent nerve
```

That shared hallway is the mechanism Chapter 1 described, now given its anatomy. Because sweet and bitter run their signals through the very same enzymes and channels — and because the cells carrying them sit side by side and converge again in the cortex — a loud bitter signal and a loud sweet signal genuinely compete, out at the periphery and again in the brain, which is exactly why one can suppress the other. It is also why a trace of bitterness can flatten a beautiful sweetness that the receptor never should have affected, and why removing that bitterness lets the same sweetener suddenly read cleaner and fuller, with nothing else changed. The two tastes are not processed in sealed tubes. They are wired to interfere — and a designer who knows the wiring can turn that interference from a problem into a tool. The alarm and the invitation are strangers at the door and cousins in the hall. Our whole craft is to quiet the one and lift the other, knowing that, downstream, each move helps the other along.

Chapter 3 · The Arithmetic of a Settled Argument — *two numbers, one honest line.*

the IFI formula · the attenuation band · delivered versus molecular · the balance index

Two tastes, two receptors, two numbers — and the discipline of this library is that both numbers rest on the same honest line, so that a coffee's bitterness and its sweetness can be read against one another on a single page. The line is the one the opening essay gave:

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

A molecule that drives its receptor to half-response at a *low* concentration is loud and scores high; one that needs a flood to be heard is quiet and scores low. That rule holds identically for both axes of this book. What differs is the reference each is anchored to, and — crucially — the *direction* in which we design.

On the sweet axis, the reference is a standardized ten-percent sucrose solution, about 0.29 molar, defined to sit at **IFI-S = 1,000.** Everything louder scores higher: a molecule twice as loud as the reference scores 2,000; a pure steviol glycoside, three hundred times louder than sucrose, carries a raw *molecular* score near 300,000; a sweet protein runs into the millions. Those molecular numbers are true and useless on a scorecard, so the Index does two sensible things, and keeping them straight is the key to reading every readout in this book. First, it **clamps the per-axis display at 5,000** — the normalization maximum shared by all eight axes — so the picture stays legible. Second, and more important, a scorecard does not report a pure compound's loudness at all; **it reports the *delivered* sweetness of a food as eaten.** When a cup reads S 1,300, that does not mean its sweetener is 1.3× sucrose per molecule — it means the whole experience of drinking it reads about 1.3 times as sweet as the reference. The molecule is a shout; you simply whisper it into the drink at parts per million. All of reformulation, in one line, is composing a target *delivered* sweetness out of ingredients whose *molecular* loudness ranges over four orders of magnitude.

On the bitter axis the formula is identical and the direction is reversed, and this is the single thing most worth getting right in the whole book, because it is the thing the earlier, thinner account of this intersection got backwards. **A higher IFI-B is a *more* bitter plant, not a less bitter one** — exactly as on every other axis, where higher means louder. We do not invert the scale to flatter a book called *Bitter-Sweet Balance.* It would be tempting: a subtractive axis might like a number that climbs as the bitterness falls, so that "better" and "bigger" point the same way. We decline, because a score is a loudness on all eight axes or it is a loudness on none; the instant one

axis runs its arithmetic backward, a profile stops being readable across the library, and Phoenix's bitterness of 2,200 would mean the opposite of a coffee's. So on this axis the design goal and the number point *opposite* ways. Success is a *smaller* reading. We raise the sweet number and we lower the bitter one.

What we report on the bitter side, therefore, is not an IFI-B in isolation but a *shift* — the distance between a plant's native bitterness and its engineered one. Define the attenuation **r** as the fraction of bitter loudness removed:

> **r = 1 − IFI-B(engineered) ÷ IFI-B(native)**

Because IFI-B runs inversely with the bitter compound's EC_{50}, an attenuation of *r* corresponds to a precise rightward shift of that EC_{50} — the receptor now needs more agonist to reach the same half-response. Writing the fold-shift as **φ = EC_{50}(engineered) ÷ EC_{50}(native)**, the arithmetic is exact and clean:

> **φ = 1 ÷ (1 − r)**

A forty-five percent quieting is not a vague softening; it is the bitter molecule's threshold moved to **1.82×** its native value, reproducible in a dish to a published tolerance. The table below is the Rosetta stone of the bitter half of this book — attenuation, the fold-shift that produces it, the fraction of the original loudness that remains — and the band it runs across is the promise *Zero-Bitter* publishes and holds:

ATTENUATION *R*	EC_{50} FOLD-SHIFT φ	IFI-B RETAINED	READS AS
20%	1.25×	80%	the floor — just past the just-noticeable
30%	1.43×	70%	a clear softening (tea)
38%	1.61×	62%	a shaped, not removed, bite (hops)
45%	1.82×	55%	the ash lifted off a dark roast (coffee)
50%	2.00×	50%	threshold doubled — the midpoint
55%	2.22×	45%	a cabbage a child negotiates no longer
60%	2.50×	40%	the kale that keeps its minerals, loses its grip
62%	2.63×	38%	a broccoli that chars sweet (the brassica ceiling)
65%	2.86×	35%	the published limit — no further

Both ends of that band are deliberate. The **floor at twenty percent** is a floor of honesty — below it the change sits at or under the just-noticeable-difference for most tasters, and we will not publish a quieting no one can feel. The **ceiling at sixty-five percent** is a floor under the flavor itself: push past it and the food loses its spine, the structural bitterness that makes coffee *coffee* and chocolate *chocolate* thinning into something flat, while the sentinel loses too much of its reach. Sixty-five percent is the most we will take off anything, and the best work usually lives well short of it. The shape of the receptor curve makes the band efficient: TAS2R38, like most bitter receptors, follows a Hill relation **R = R_max · [A]^n ÷ (EC_{50}^n +**

[A]ⁿ) with a coefficient *n* near 1.3 to 1.5, so a food sitting near its EC_{50} — most bitter foods do, which is why they register as bitter — is in exactly the steep middle where a modest rightward shift buys a large drop in response. And perception compresses the rest: intensity reads roughly as the logarithm of the signal, so a doubled threshold is felt not as a brutal halving but as a gentler step from *too much* to *right.*

Now the part unique to this book — because *Bitter-Sweet Balance* is not either axis alone but the composition of the two, and a composition needs a way to say whether it is *balanced.* We read a dish on both numbers at once and ask whether the sweetness and the residual bitterness meet. A useful working figure is the **balance ratio** — the delivered IFI-S set against the engineered IFI-B — read alongside the suppression each exerts on the other. A coffee meant to be drunk black wants a low ratio: a bitterness quieted to structure, a sweetness barely lifted, perhaps S 300 against B 1,900, so the cup reads as clean coffee and not as a sweetened drink. A drinking chocolate wants a higher one: S 1,500 against B 1,750, the sweetness rising to meet a bitterness kept proud. A nutrition bar wants the two nearly level, S 1,400 against B 1,650, so a dense, difficult food reads as something a person reaches for twice. The art, which the next chapter makes quantitative, is that these are not two independent dials. Lower the bitterness and the same sweetness reads louder, because less of it is spent on suppression; lift the sweetness and the bitterness reads quieter, because the suppression runs both ways. The balance is a fixed point of a two-way interaction, and finding it — the least intervention on either axis that lands the dish where it should sit — is the whole of the design. One honest line, read twice, in opposite directions, and composed into a third thing: the settled argument itself, given a number.



PART II OF VI

The Mechanism of the Two Moves

the sugar tax, measured · the composition that finds the ratio · the loop that scores an edit before we cut

Chapter 4 · The Sugar Tax — *what a bitterness costs in sweetness.*

mixture suppression quantified · peripheral and central · the uncovering

Chapter 1 named the mechanism; this chapter puts numbers on it, because the numbers are where the whole economy of the two-move design becomes visible. Mixture suppression — the mutual quieting of bitter and sweet — is not a vague culinary intuition but one of the more carefully measured relationships in sensory science, and once you can read it quantitatively you can see exactly how much sugar a harsh food has been quietly spending, for no flavor at all, just to keep its own bitterness down.

Start with a controlled observation of the kind trained panels run routinely. Take a fixed bitter stimulus — a defined concentration of caffeine, or of a coffee's chlorogenic lactones — and rate its bitterness alone: call it 100 on an intensity scale. Now add sucrose in a rising series and re-rate the bitterness at each step. It falls along a smooth curve. A modest sweetness, enough to read as lightly sweet on its own, might pull the rated bitterness from 100 down to perhaps 70; a strong sweetness, the kind in a dessert, can push it below 50. The reverse experiment runs the same way: a fixed sweetness, rated against a rising bitterness, is itself progressively suppressed. Each taste spends part of itself hushing the other, and the exchange rate — how much perceived bitterness a given sweetness removes, and vice versa — is specific, measurable, and steep enough through its middle to design against. This is the raw material of every balanced recipe ever made, finally written as a function instead of a feel.

Two mechanisms, from Chapter 2, set that exchange rate, and they behave differently. The **peripheral** component lives in the shared transduction hallway of the taste cell — the common gustducin, PLCβ2, calcium, and TRPM5 machinery that sweet and bitter both run through — where simultaneous strong signals compete for shared capacity. The **central** component lives in the cortex, where the brain integrates the gustatory lines and a dominant taste can perceptually mask a subordinate one; this is the larger share of the effect, and it is why suppression can occur even between a taste on the front of the tongue and one on the back, cells apart, that never shared a hallway at all. For the designer the split matters because it tells us that suppression does not require the two molecules to touch or to bind anything in common. A sweetness grown into a coffee quiets that coffee's bitterness whether or not the two ever meet at a receptor, because most of the quieting happens upstream in the brain's own accounting. The suppression is, in the most literal sense, in the eater's head — which is precisely why it is so reliable, and why it costs the food nothing but the sweetness itself.



Green coffee against roasted — most of a cup's harshness is forged in the gap between them, and the sugar a drinker adds afterward is a tax paid to suppress a bitterness the bean never needed to carry so loud.

Now the tax, made concrete. Consider a conventional dark-roast coffee whose harsh phenylindane bitterness reads, untouched, at a high level — the ash that sends a drinker reaching for sugar. To bring that cup into the balance the drinker finds comfortable, they spoon in sugar until the perceived bitterness drops to a tolerable level; in a strong cup, that is routinely two teaspoons, roughly eight grams, perhaps thirty calories, in a single mug. Read what that sugar is doing. Only a fraction of it registers as sweetness — the cup does not taste like a dessert — because most of it is being *spent* against the bitterness, consumed by the suppression curve, converting itself into an absence of harshness rather than a presence of sweet. The eight grams are a tax, levied by the bitterness, payable only in sugar, and collected cup after cup for a lifetime. The drinker never sees the bill as sugar-for-suppression; they see it as "how I take my coffee." But that is what it is.

The two-move design refuses to pay the tax, and it does so by attacking the thing the tax was levied on. Quiet the coffee's harsh wave at the source — lower the phenylindane bitterness by forty-five percent before the cup is ever poured — and the suppression that needed to be bought with sugar has already, in large part, been bought with engineering. The bitterness that remains is low enough that little or no sweetness is required to balance it. And here is the gift hidden inside the suppression curve, the one that makes this the gentlest intersection in the library: **when you lower the bitterness, a sweetness that was in the food all along is released from suppression and seems to arrive.** The bean's own faint natural sweetness — always present, always masked under the harsh roast — becomes audible the moment the harshness comes down, not because a grain of sugar was added but because the bitterness that was suppressing it is no longer there. We call this *uncovering,* and it is the cheapest sweetness in the kitchen, because the plant already paid for it. Hold a cup's sugar constant and merely lower its bitterness, and trained panels rate the same cup as *sweeter* — the sweetness readout rises though nothing sweet was added. Much of the pleasure of a de-bittered coffee is exactly this: the drinker, reaching by habit for the sugar they always thought they needed, discovers the sweetness was there the whole time, waiting under the ash.

This is why the two moves are not a simple addition of two independent effects but a compounding of them, and why a well-composed bitter-sweet target asks so little of each axis. The lowered bitterness needs less sweetness to balance it; part of the sweetness it does need arrives free, uncovered, as the bitterness falls; and any sweetness we then grow in on purpose arrives into a cup already quieted, where a very little goes a long way because it is no longer being spent on suppression. Three savings stacked: less bitterness to fight, sweetness uncovered for nothing, and lifted sweetness working at full value instead of being taxed away. The old single-lever cure captured none of these — it poured sugar in to do all the suppression itself, at full price, every cup. The arithmetic of the settled argument is the arithmetic of those three savings, and the chapters of recipes ahead are, every one of them, that arithmetic served in a glass or on a plate.

Chapter 5 · The Two-Move Design — *the balance you compose, not mask.*

the B×S composition library · why max-both fails · the ratio · ≤7% CV

Here is the intersection itself, the real subject of the book, stated as an engineering problem. For any given coffee, chocolate, or bar, there is a pair of settings — how far down the bitterness, how far up the sweetness — that lands the thing in balance with the least intervention on either side. Finding that pair, reproducibly, for plant after plant and product after product, is the work. We hold a dedicated library for it: the **B×S composition set**, one member of the multi-axis collections that sit alongside the big single-axis programs. Across the whole platform there are **8,400 dual-axis (M2) constructs** and **2,800 triple-axis (M3) constructs**; the B×S line is the dual-axis family aimed at exactly this quarrel, and every construct in it exists to answer one question — *for this food, what two settings settle it?*

The naïve answer is to maximize both: attenuate the bitterness as far as the band allows and lift the sweetness as high as it will climb. It does not work, and the reason is the suppression of Chapter 4 running in reverse. Over-attenuate the bitter and the sweetness you added now has nothing to lean against; with no residual bitterness to suppress, the sweetness reads flat, bald, faintly false — the taste of a diet soda that forgot what it was imitating, a sweetness standing in a cleared room with nothing to give it dimension. Over-lift the sweet and you have simply rebuilt the old sugar problem in a sugar-free costume, a cloying drink that happens to carry no sucrose. **Balance is not the sum of two maxima. It is a ratio** — a specific pairing of residual bitterness to delivered sweetness — and the ratio, not either extreme, is the design. The residual bitterness is not a failure to be minimized; it is *structure,* a thing to tune *to,* the backbone the sweetness drapes over. A coffee drunk black wants a great deal of structure and a whisper of sweet; a drinking chocolate wants a proud bitterness met by a generous sweet; a bar wants the two nearly level. Each is a different point on the ratio, and naming that point is the first act of every composition.



The research plot where a target becomes a plant — a bitter-sweet composition is grown to a specification, two settings held together, not a happy accident of a single season.

We find the ratio the way this platform finds every composition: with a **genome-scale flux-balance model** that searches for settings which hold together rather than cancel. A flux-balance analysis is a mathematical reconstruction of a plant's entire metabolic network, and it lets us pose a precise question of a living pathway — *lower this bitter compound by this much, and raise this sweet one by that much, and what does the rest of the plant do?* On the bitter side it tells us how far a given edit actually lowers the harsh molecule (the relationship is rarely one-to-one; pathways reroute) and what else moves when it does. On the sweet side it holds the hard constraint that the loud sweet molecules — steviol glycosides, mogrosides, sweet proteins — are built from precursor pools the plant also needs for pigment, for structure, for growth, so a sweetness pushed too hard yields a sweeter plant on a sicker one. The model apportions the plant's carbon and energy across both demands at once and finds the edge: the most balance the plant can carry while remaining, unmistakably, a healthy plant. It is the same mathematics that lets *Swicy* set heat against sweetness without either canceling the other, and that lets a restored tomato carry its umami without tipping — aimed here at the oldest pair of adversaries on the tongue.

The output of all this is not a flavor that tastes right once. It is a *specification* — a target on both axes that the plant can be grown to and that can be hit again, batch after batch, season after season, inside the standard this whole library holds to: a batch-to-batch **coefficient of variation of seven percent or less.** That figure is the difference between a balance you designed and a balance you got lucky with. A drinking chocolate that reads the same in August and in February, a coffee whose black cup lands on the same two numbers every roast, a bar whose twelfth unit tastes like its first — these are specifications, not fortunate harvests, and the ≤7% CV is the engineering that makes them so. The full machinery of that reproducibility is its own volume, *Precision & Consistency* (Volume LX); here the point is narrower and decisive. A composition is finished not when it tastes balanced but when it tastes balanced *every time,* to within seven percent, on a crop grown at scale.

And the specification is *small,* which is the quiet triumph of doing both moves at once rather than either alone. Because the two reinforce — the lowered alarm needs less sweetness to feel balanced, and the lifted sweetness further hushes the alarm, with a

share of the sweetness arriving free as the bitterness falls — a well-composed bitter-sweet target asks far less of each axis than either would need by itself. Less attenuation, so the plant keeps its character and its spine. Less sweetness, so there is less of anything in the cup that was not already the coffee or the cocoa. The old single-lever cure maximized one thing, sugar, and paid full price for it. The two-move design minimizes *both* things and lets the biology make up the difference. You compose the balance instead of masking the bitterness, and composing, it turns out, costs less than masking ever did — in sugar, in intervention, and in everything the food had to surrender to be made bearable the old way.

Chapter 6 · The Design Loop — *scoring an edit before we cut.*

ESM2 · flux-balance analysis · CRISPR · the predicted return

The chapters ahead name a good many valves to turn — HQT in the coffee bean, GS-OH and MYB28 in the brassicas, the LAR/ANR branch in cacao and tea, and on the sweet side the UGT glycosyltransferases that finish a steviol glycoside or a mogroside. Before any of them, the question underneath them all: *how do we decide which edit to make, and how far, on both axes at once, before we ever touch a living plant?* The answer is not intuition and it is not brute trial. It is a loop — predict, build, measure, feed back — and it is the same loop, whether the target is a bitterness to lower or a sweetness to raise.

The loop begins in silico, in two models that ask two different questions. The first is a **protein language model** — in our pipeline an ESM2-class network of some 650 million parameters, trained on the evolutionary record of millions of protein sequences until it has, in effect, learned the grammar of what a working protein looks like. Pointed at a target enzyme — a bitter-building transferase we want to down-tune, or a sweet-finishing UGT we want to drive forward and make more selective — it scores *every possible single-residue substitution* by its predicted effect on catalytic efficiency and fold stability. Instead of guessing which change might cleanly re-tune an enzyme, we rank the entire space of changes before synthesizing one. The structural scaffold the model reasons over is itself a prediction, and the platform's confidence in that scaffold is something we read explicitly: a predicted-aligned-error map confesses where the model is sure how two regions of a fold sit relative to each other and where it is not, and we trust an edit most where the structure beneath it is read most clearly.

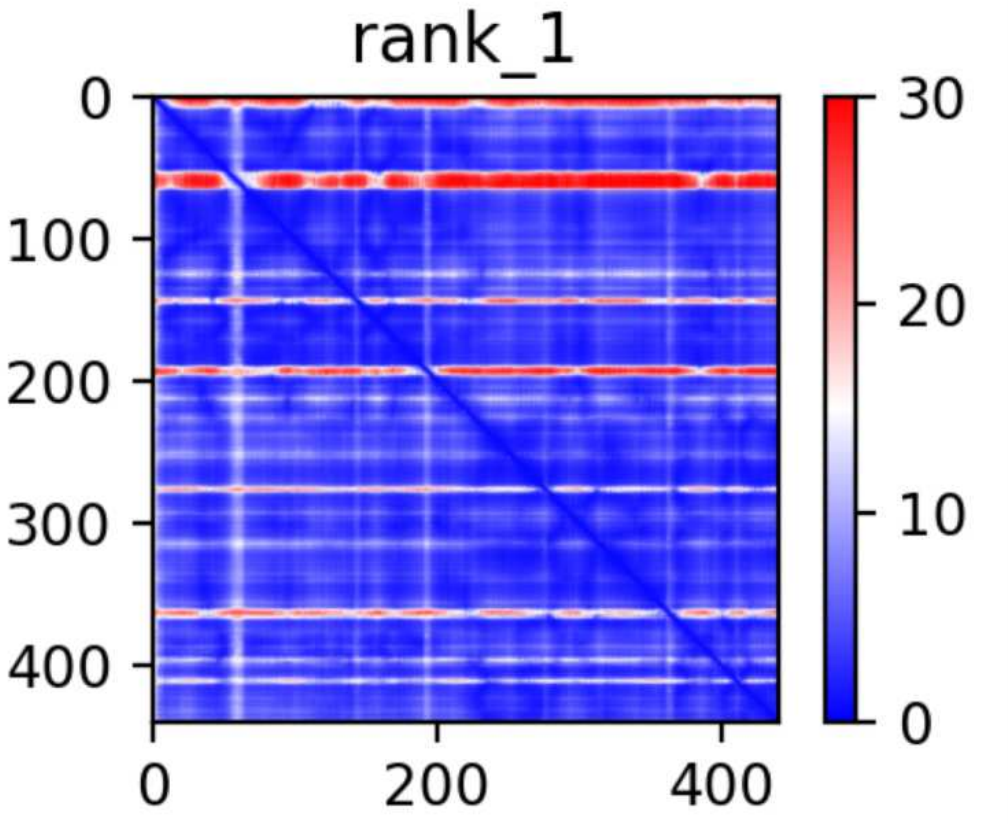

A predicted-aligned-error map on the platform's structure pipeline — dark where the model is confident how two parts of a protein sit together, bright where it is not. We design our edits into the regions we can read.

The second model asks not about one protein but about the whole plant — the **flux-balance analysis** of the last chapter — and on an intersection axis it earns its keep twice over, because it must balance a subtraction and an addition at the same time. It is how we know, before cutting, that lowering a coffee's chlorogenic pool at HQT genuinely starves the harsh phenylindanes downstream rather than being quietly compensated by a dozen neighboring steps; and how we know that driving a stevia leaf's glycosylation forward toward clean Rebaudioside M will not bankrupt the isoprenoid pool the plant needs for its pigments. Flux balance guards against the classic error of metabolic engineering — editing a step that looks important and watching the network shrug and reroute around it — by showing us the true rate-limiting valve and the full downstream consequence of turning it, on both axes, together.

The proof that this loop works is not on either of this book's axes; it is the plant that made the company. The method's first triumph was on **Pun1**, the final capsaicin synthase at the end of the pathway — though our flux model predicts control sits upstream, at **4CL and HCT/CCoAOMT**, not at the terminal Pun1 step itself; Pun1 labels the product, it does not govern the rate. The protein language model scored every substitution on Pun1, and two edits emerged that no breeder had ever found or combined — written in with CRISPR, our model predicts they lift a fruit to approximately **2.94 million SHU** at a designed batch-to-batch variation of seven percent or less:

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

Those two edits are the whole argument for the platform in miniature: the same two models, the same CRISPR pen, now score the coffee transferase and the brassica regulators and the steviol UGTs exactly as they once scored Pun1. The science did not have to be reinvented for bitterness, or for sweetness, or for the composition of the two. It had to be *aimed.* That is what a platform is — a method general enough that the brassica work and the stevia work and the superhot work are the same work pointed at different pathways, which is why one family can design flavor across forty-two organisms and eight axes at all.

Then the loop closes, and this is the part that keeps it honest. A designed edit is still only a prediction until it is **tested in the lab** — which is why the loop is built to end in measurement. The engineered plant would go onto the assays of this book's final chapter — a bitter compound read against the TAS2R38 reference receptor for an IFI-B, a sweetener read against TAS1R2·TAS1R3 and a sucrose reference for an IFI-S — and a trained human panel, screened for taster genotype, would confirm that the two numbers in the dish match the balance on the tongue. That result feeds back to sharpen the next round of predictions, and the gate on the whole loop is reproducibility: the ≤7% CV design specification that recurs across every axis in this library, because a balance we cannot hit again is not a design, only an accident we got to watch once. We never design on a hunch. Every edit in the chapters ahead — every quieted bitterness, every lifted sweetness — was scored before it was built, and the loop is designed to close on measurement.



PART III OF VI

Coffee

the two bitternesses of the cup · the lever past the roast · the cup that asks for nothing

Chapter 7 · The Two Bitternesses of Coffee — *lactones, phenylindanes, and the roast.*

DOMAIN · COFFEE · Coffea · 40 constructs · the chlorogenic acids · the caffeine boundary

Coffee is where this book begins its real work, because coffee is where the quarrel is loudest and the stakes most familiar. A great coffee is a controlled bitterness that no one would want gone — a dark, round, almost savory depth that no sweetness could stand in for, the thing a drinker falls for and the reason the acquired taste is worth acquiring. And yet most of that bitterness, at most tables, is never tasted as itself; it is buried under milk and sugar, and the hundreds of volatile compounds that make a coffee worth drinking drown with the harshness they came packaged in. Coffee is the purest case in the library of the difference between the two bitternesses hiding in a single food — the one that is the drink's character, which we protect with everything we have, and the one that is its punishment, which we are here to lift. A century of roasters have had only levers that move both at once. We offer the one that moves them apart.

The chemistry sorts cleanly into those two bitternesses, and the coffee method turns entirely on the distinction. The green coffee bean is unusually rich in **chlorogenic acids** — a family of esters built on the universal phenylpropanoid branch, the caffeoylquinic acids (3-, 4-, and 5-CQA) and their relatives, together as much as six to eight percent of the green bean's dry weight, among the highest chlorogenic loads of any plant food. In the green bean they are not especially bitter. The bitterness is *made in the roast,* and it arrives in two waves. The first wave is the **chlorogenic acid lactones**: as the roast moves into light and medium territory, the chlorogenic acids dehydrate and cyclize into lactone rings, and these lactones are the pleasant, structured, coffee-defining bitterness — the round depth of a well-judged medium roast, broadly the bitterness we want to keep. The second wave comes later and meaner. Push the roast darker, past where the lactones peak, and they in turn break down into **phenylindanes**, fused bicyclic compounds that are the signature of a dark roast and the source of its harsh, lingering, almost ashen bite — the bitterness that clings to the back of the throat and sends a careful drinker reaching for milk. One wave is the coffee's soul; the next is its scorch. They sit on a single continuum of roast, which is exactly why the roaster alone can never separate them.



The coffee bean in cross-section — most of a cup's eventual harshness is latent here, in the chlorogenic chemistry of the green seed, long before the roaster's heat ever forges it into the phenylindane scorch.

Which receptors read the cup? Not, mostly, the anchor. A coffee's bitterness is a chord struck across several of the twenty-five bitter sentinels — the chlorogenic lactones and phenylindanes engaging TAS2R46 among others, the caffeine lighting up TAS2R7, TAS2R10, TAS2R14, TAS2R43, and TAS2R46 together — a broad, layered activation rather than the sharp single-receptor hit a glucosinolate lands on TAS2R38. We score it, as we score everything, against the TAS2R38-anchored axis, and we say the honest thing on the readout: the IFI-B of a coffee is a faithful, reproducible *projection* of a multi-receptor sensation onto the one axis the whole library shares. It is a true measure of how loud the cup's bitterness is. It is not a claim that one receptor does the hearing. For coffee, a committee hears.

Caffeine deserves a plain word, because its reputation outruns its role. It is genuinely bitter, and it contributes — but studies of coffee bitterness put it at only perhaps ten to fifteen percent of a cup's total perceived bitterness. The great majority is the chlorogenic story: the lactones and, above all in a dark roast, the phenylindanes. This matters for method, because it means a coffee can be meaningfully de-bittered *without touching its caffeine at all*, and we treat that as a deliberate design boundary. Caffeine rides a wholly separate biosynthetic route — the purine-alkaloid pathway, xanthosine through theobromine to caffeine via the methyltransferase caffeine synthase — and because it contributes so little of the harshness and so much of the reason people drink coffee at all, there is no reason to disturb it and every reason not to. A drinker's lift stays exactly as nature set it. We are in the business of the harsh wave, and nothing else.

The two-wave structure is the gift, because it tells us precisely what to design: not a coffee stripped of its chlorogenic character, which would be a thin, hollow cup, but the *second* wave moderated — the phenylindane end pulled down so a roaster can chase body and sweetness into a darker roast without the ash arriving to end the conversation. And because the phenylindanes are made from the chlorogenic acids, which are made on a pathway we know how to tune, the lever sits upstream, in the green bean, long before the roaster's heat. The target is the supply of chlorogenic acid itself. The branch, through to the cup's precursor, runs like this:

STEP	ENZYME	PRODUCT
1	**PAL** (phenylalanine ammonia-lyase)	cinnamic acid
2	**C4H** (cinnamate 4-hydroxylase)	*p*-coumaric acid
3	**4CL** (4-coumarate:CoA ligase)	*p*-coumaroyl-CoA
4	**HCT / HQT** (hydroxycinnamoyl-CoA:quinate transferase)	*p*-coumaroyl quinate
5	**C3′H** (coumaroylester 3′-hydroxylase)	**chlorogenic acid (5-CQA)**

The pivot is step four. **HQT** — hydroxycinnamoyl-CoA:quinate hydroxycinnamoyltransferase — is the enzyme that commits carbon specifically to the chlorogenic pool rather than to the dozens of other phenylpropanoid products the branch feeds. Lower HQT activity and the green bean accumulates less chlorogenic acid; less chlorogenic acid means, downstream in the roaster, fewer lactones and, the real prize, proportionally fewer phenylindanes. We tune HQT *down,* not off — a designed reduction set by the flux-balance model, lowering the chlorogenic load enough to lift the harsh wave while leaving ample chlorogenic acid for the lactone body the cup depends on and for the bean's own physiology. The **40 coffee constructs** in the library are, in the main, a titration series across that single lever and a few of its neighbors — a ladder of HQT settings, each designed to a predicted Index reading, converging in our model on the point where the ash lifts and the soul stays. Forty is a small and hard-won number, every one a frontier-grade engineering problem, because *Coffea* does not give up its chemistry easily; the full monograph is *Coffee* (Volume XV). And one boundary beyond caffeine is worth stating: **aroma.** A coffee's smell, the hundreds of volatiles we actually love, is forged in the roast on roads largely separate from the chlorogenic one — the Maillard chemistry, the degradation of trigonelline, the lipid and sugar reactions — so our lever pulls the harshness down and leaves the perfume entirely standing. That separation is the whole reason the method works, and it is what the next chapter puts in a cup.

Chapter 8 · Coffee, Read Black — *the cup that asks for nothing.*

application · pour-over · clarified cold brew · the unsweetened cortado · the bitter-sweet mocha

The clearest argument for a de-bittered, lightly-sweetened bean is the cup a person reaches for without reaching, afterward, for sugar. So the coffee part ends in the kitchen, with four cups built on the clarified cultivars of the last chapter — each deliberately plain, so what you taste is the bean and the method and not a recipe hiding them, and each read on the Index so a drinker can see exactly what the two moves did. None of them is clever. Cleverness is what a harsh coffee has always demanded — the milk, the sugar, the flavored syrup sent in to cover an edge the bean never needed to carry. On a coffee quieted at the source and lifted a whisper, the cleverness is unnecessary, and what is left is the coffee, finally able to be met for what it is.

The first is the black cup, the one we drink before anyone else is awake. The whole recipe is choosing the bean and getting out of its way.

```
The Black Cup · makes 1 · 4 min
— 18 g clarified-bean coffee (SS-line), medium roast, ground medium-fine
— 300 g water at 94 °C
— nothing else — no sugar, no milk, on purpose
```

1. Rinse the paper filter and discard the rinse water; it dulls a cup otherwise. Add the grounds and make a small well in the center of the bed.
2. Pour 50 g of water to wet every particle and wait forty seconds while the bed blooms and settles — the bloom is the bean releasing its roast gas, and a clarified bean blooms just the same.
3. Pour the rest in slow concentric rings, aiming to finish the draw-down around two minutes forty-five.
4. Drink it black. Notice that you did not reach for anything.

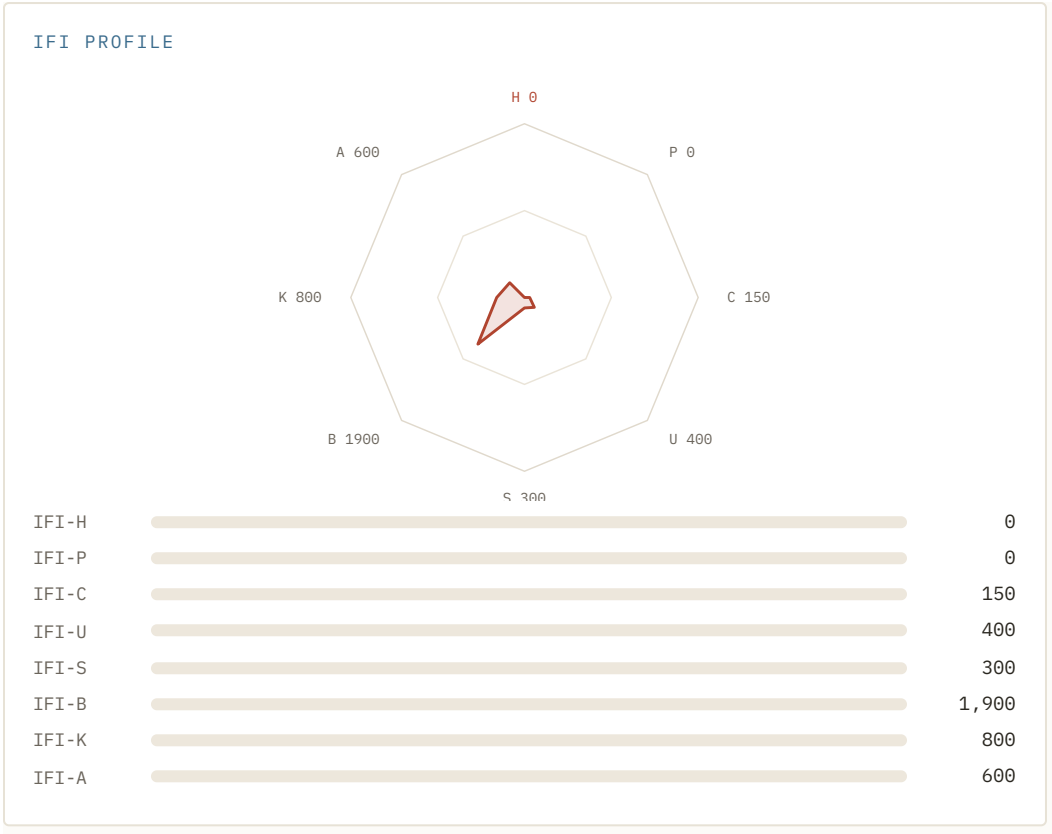

Read the scorecard and the method explains itself. The **B of 1,900** is the harsh phenylindane wave lifted about forty-five percent while the lactone body stays — a cup whose bitterness reads as structure, not scorch. Around it: a **K of 800**, the round, filling mouthfeel of a good dark roast, untouched; a quiet **U of 400**, the faint savory depth real coffee carries; an **A of 600**, a clean dry finish left in on purpose, because a coffee with no dryness tastes slack. And the **S of 300** is the quiet proof of Chapter 4 — not a sweetness we spooned in but the bean's own trace sweetness, *uncovered,* suddenly audible because the harshness that was suppressing it has come down. Three hundred is a small number and it is doing large work: it is the reason the cup needs no sugar, grown free from the bean rather than bought with a spoon.

The second is cold brew, where the method of brewing already does, crudely, a little of what our engineering does precisely — steeped cold and slow, coffee extracts far less of its harshest, most soluble compounds, which is why cold brew is famously smooth. Begin from a bean that carries less of the harsh wave to extract, and the two effects compound.

> Clarified Cold Brew Concentrate · makes ~1 L · 16 h
> — 200 g clarified-bean coffee (SS-line), coarse grind
> — 1 L cool filtered water
> — a wide jar and a fine cloth or paper filter

1. Stir the coarse grounds into the water until every particle is wet. Cover.
2. Steep at room temperature, or in the refrigerator, sixteen hours — no more, or the bean's heavier compounds begin to extract.
3. Strain slowly through cloth, then once through paper for a clean cup. Do not press the grounds; pressing drags the harshness back in.
4. Keep refrigerated. Serve over ice, cut one-to-one with water or milk.

IFI PROFILE

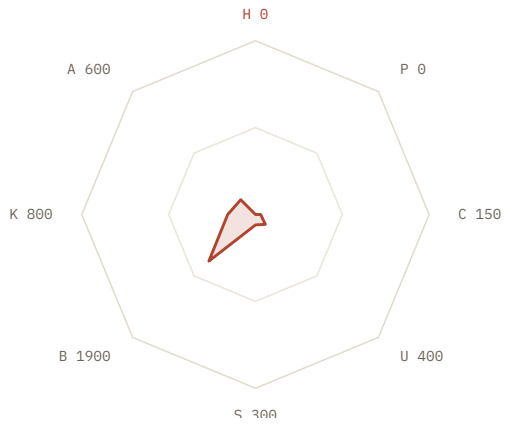

IFI-H		0
IFI-P		0
IFI-C		150
IFI-U		400
IFI-S		300
IFI-B		1,900
IFI-K		800
IFI-A		600

The concentrate lands on the same numbers as the pour-over, which is the point — the balance is in the bean, not the brew method, so it travels. The third cup is the one that tends to end the argument: a cortado, espresso cut with a little warm milk, pulled from the clarified bean and served with nothing added. Its small volume and its milk put the bean nowhere to hide.

```
Unsweetened Clarified Cortado · serves 1 · 4 min
— 18 g clarified-bean coffee (SS-line), for a double shot
— 60 ml whole milk, steamed to just warm, barely foamed
— a small glass, warmed
```

1. Pull a double shot into the warmed glass — aim for a slow, even, honey-thick extraction.
2. Steam the milk only to body temperature, about 60 °C, with the barest texture; this is a cortado, not a latte, and the milk is a cut, not a blanket.
3. Pour the milk through the crema in a single steady stream. Serve at once, and taste it before you reach for anything.

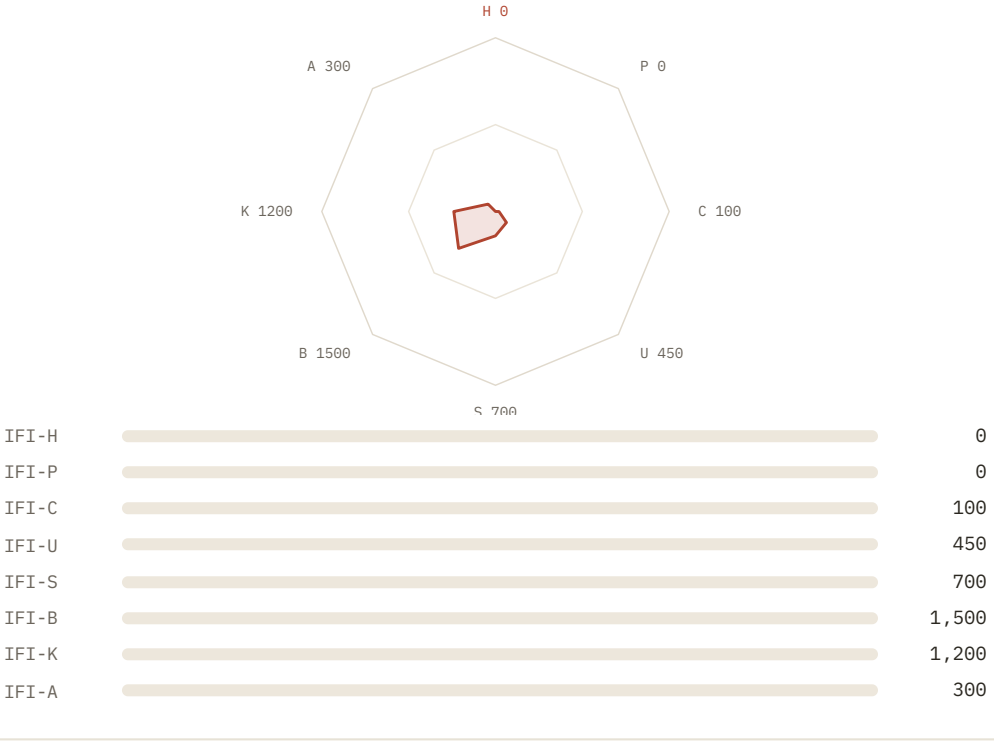

The cortado tells the subtler story. The **B drops to 1,500** — lower than the straight brew, partly the bean and partly the milk's own softening — and look what the quiet uncovers: an **S of 700**, most of it the milk's lactose but some of it the bean's own sweetness, finally readable; a **K of 1,200**, the combined body of milk and a good dark roast, round and filling; an **A down at 300**, just enough dryness to keep the cup from going flat. A drinker who always took their coffee sweet discovers, often with real surprise, that the sugar was never balancing the coffee — it was burying a harshness that, on this bean, was never there to bury.

The fourth is the one place in this part where we lift the sweetness on purpose rather than merely uncover it — a bitter-sweet mocha, where coffee meets the dark-chocolate work of the next part, both bitternesses quieted and a real sweetness grown in to meet them without a grain of sugar.

> Bitter-Sweet Mocha · serves 2 · 8 min
> — 2 double shots clarified-bean coffee (SS-line)
> — 20 g low-bitter dark cocoa (see Part IV)
> — 350 ml whole or oat milk
> — Reb-M stevia + mogroside-V monk fruit, to taste (no sugar)
> — 1 g γ-glutamyl kokumi base · a pinch of flaky salt

1. Whisk the cocoa with a splash of cold milk into a smooth paste, with no lumps.
2. Warm the rest of the milk with the kokumi base and salt until it steams; do not boil. Whisk in the cocoa paste.
3. Sweeten a little at a time with the stevia and monk fruit until the drink reads as chocolate-sweet, not candy-sweet — stop early.
4. Divide the shots between two warmed glasses and pour the chocolate milk over. The salt lifts the sweetness without adding any.

IFI PROFILE

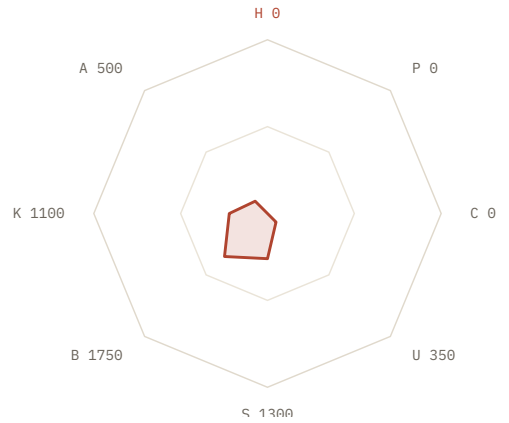

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	350
IFI-S	1,300
IFI-B	1,750
IFI-K	1,100
IFI-A	500

Here the **S of 1,300** leads, because a mocha should — but note the **B held at 1,750**, the combined coffee-and-cocoa bitterness quieted to a proud structure, not erased. The sweetness rises to meet that bitterness rather than bury it, the kokumi base gives the drink the body the missing sugar would have supplied, and the whole thing lands as a composition: two bitternesses tuned down, one sweetness grown in to the exact level that settles them, and no sucrose anywhere in the glass. Four cups, one bean philosophy — a coffee that, whether black or milked or married to chocolate, no longer needs to be made sweet, because the balance was built before it was ever poured.



PART IV OF VI

Chocolate & the Nutrition Shelf

theobromine and the flavan-3-ols · sweetness against the dark · the everyday bar

Chapter 9 · Theobromine and the Flavan-3-ols — *cacao, where bitter meets astringency.*

DOMAIN · CACAO · the methylxanthine and flavonoid roads · a note on status · braided B × A

Cacao is the axis's hardest and most beautiful test, because cacao is a bitterness that nobody — not even a child — ultimately wants fully removed. A fine dark chocolate is *supposed* to grip a little. It is supposed to carry the clean, bright, almost medicinal lift of its alkaloid and a tannic edge that keeps all that cocoa butter and sugar honest. Take the bitterness all the way out and you do not get better chocolate; you get candy, or worse, a brown confection that tastes of nothing in particular. So on cacao we work with a scalpel's patience and keep more than we cut, and the whole of the craft is knowing which of the bean's two very different bitternesses to protect and which to ease — a judgment made harder, and more interesting, by the fact that cacao's bitterness is braided inseparably with its astringency.

The first bitterness is **theobromine** — the principal methylxanthine of cacao, a close cousin of caffeine, present at a couple of percent of the bean and responsible for much of chocolate's bright, lifting, characteristic edge. Caffeine rides along in a smaller amount, and the pair together carry the cocoa's stimulant lift and a good share of its bitterness. Theobromine is built on the purine-alkaloid route — the same family of N-methyltransferase steps that builds caffeine elsewhere, xanthosine methylated and oxidized toward the finished alkaloid — and this bitterness we very largely **leave alone,** because it *is* the chocolate telling the truth about itself. A dark chocolate without its methylxanthine lift is a dark chocolate with its personality removed.



Cacao pod and beans — two separate bitter chemistries ripen in the same seed: the methylxanthine lift we protect, and the flavan-3-ol grip we ease. Easing one without disarming the other is the whole of the cacao problem.

The second bitterness is the one we ease: the harsh, drying grip of the **flavan-3-ols** and their polymers, the **procyanidins.** These are the cacao's polyphenols — (−)-epicatechin, (+)-catechin, and the

oligomeric procyanidins they build into — and they are the compounds that can tip a very high-percentage bar from structured into punishing. They are built on the flavonoid arm of the phenylpropanoid branch:

STEP	ENZYME	PRODUCT
1	**CHS** (chalcone synthase)	`naringenin chalcone`
2	**CHI → F3H**	`dihydroflavonol`
3	**DFR** (dihydroflavonol reductase)	`leucoanthocyanidin`
4a	**LAR** (leucoanthocyanidin reductase)	**`(+)-catechin`**
4b	**ANS → ANR** (anthocyanidin reductase)	**`(−)-epicatechin`** → `procyanidins`

The lever is the **LAR/ANR branch point** — the same pair of enzymes that governs flavan-3-ols in tea and the grape, which is precisely why the method transfers across the whole polyphenol family. Tuning the flux through ANR moderates the roughest, most aggressively bitter-and-astringent procyanidins while leaving the gentler flavanols, and the structure they give, in place. We ease the roughest fraction by roughly forty percent and protect the rest. And here, more than anywhere in the book, the two axes braid: these same procyanidins are the heart of **IFI-A, astringency** — the dry, puckering grip that comes not from a bitter receptor at all but from the polymers cross-linking the proteins in your saliva, read by the salivary PRPs. The distinction is a matter of molecular size: a small flavan-3-ol monomer slips into a bitter receptor's pocket and reads as *bitter;* link a few into a procyanidin and the polymer grows too large for the pocket but just the right size to grip saliva, reading as *astringency.* The same biosynthetic road feeds both sensations, so every move we make on cacao's bitterness is also a move on its grip, and the flux model has to hold both in view at once — easing the bitter thread while protecting the astringent one, because a dark chocolate's structured astringency is not a flaw to remove but a backbone to keep.

A word on status, in the plain voice this library keeps for the truth about its own work: **cacao is not among the forty-two host organisms in our platform.** Our standing cacao work is method-transfer and roadmap — the flavonoid methodology applied to the polyphenol hosts already in the program, the grape with its 1,610 constructs and tea with its 280, pointed toward the cacao bean, where the full host program lives on the companion shelf as *Cacao & Chocolate* (Volume XVI). We tell you this because a number that overstates where the science actually stands is worse than no number, and the discipline of this whole library is to say exactly how far along a thing really is. So the chocolate recipes that follow do two honest things at once. Where they call for a low-bitter cocoa, they describe the profile the method *targets* — a representative readout, not a shipped measurement. And where they lift the sweetness to meet that cocoa, they use the move we can make *today,* with no plant engineering at all: because a high-percentage chocolate carries little sugar by definition, a sweetness grown in rather than spooned on leans on the cocoa's bitterness through the mixture suppression of Chapter 4, quieting the alarm without touching the bean. The theobromine lift stays. The astringent backbone stays.

Only the punishment softens. Here is the profile a method-targeted 80-percent bar aims for, scored on the Index to show the intended balance:

IFI PROFILE

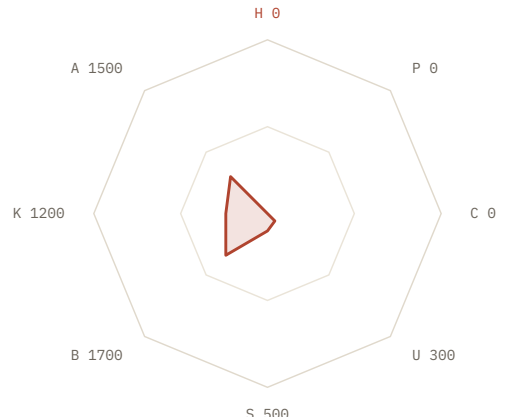

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	300
IFI-S	500
IFI-B	1,700
IFI-K	1,200
IFI-A	1,500

Read it as a map of restraint. The **B of 1,700** is the procyanidin harshness eased about forty percent — but notice the **A held high at 1,500,** the structure a real bar must keep, and the **K of 1,200,** the deep, lingering body of good cacao, and the **S held low at 500,** enough to round the bar and never enough to mask it. A chocolatier working with such a cocoa can take the percentage higher — further into the fruit and the roast of the bean — without the bar turning against the eater at the finish. The grip stays. Only the scrape goes. On cacao, more than anywhere, bitter-sweet balance is not subtraction for its own sake but the fine judgment of exactly how much of a thing's defining edge to protect — and how much sweetness to grow in to meet what remains.

Chapter 10 · Sweetness Against the Dark — *chocolate, settled without sugar.*

application · drinking chocolate · bitter-sweet ganache · a note toward swicy

Chocolate is the richest proving ground for the sweet half of the two-move design, because chocolate is where a lifted sweetness has the most bitterness to lean on and the most to gain from leaning. These three preparations use the sweet move we can make today — the plant's own loud sweet molecules, grown up rather than poured in — against a dark cocoa whose grip we respect too much to erase. Each is built to let the cacao stay cacao: bitter, structured, alive, and sweetened only to the exact point where the punishment becomes depth.

The first is drinking chocolate, the simplest place to watch a sweetness settle a bitterness. Made with real dark cocoa and no sugar at all, its sweetness comes from an engineered sweetener that moves the receptor at parts per million, so a pinch in the whole pot does the work a quarter-cup of sugar used to.

```
Dark Drinking Chocolate, Sugar-Free · serves 2 · 10 min
— 40 g dark cocoa (85%, ideally an IFI-B-attenuated low-bitter cocoa)
— 500 ml whole or oat milk
— Reb-M stevia, to taste (begin 0.2 g)
— mogroside-V monk fruit, a pinch (to round the middle)
— 1 g γ-glutamyl kokumi base · flaky salt, to finish
```

1. Whisk the cocoa into a little cold milk to a smooth paste.
2. Warm the rest with the kokumi base until it steams; whisk in the paste and lower the heat.
3. Sweeten with the stevia and monk fruit to just-sweet — dark chocolate wants restraint; stop before it reads as candy.
4. Whisk hard for froth and finish with a few flakes of salt, which lifts the sweetness without adding any.

IFI PROFILE

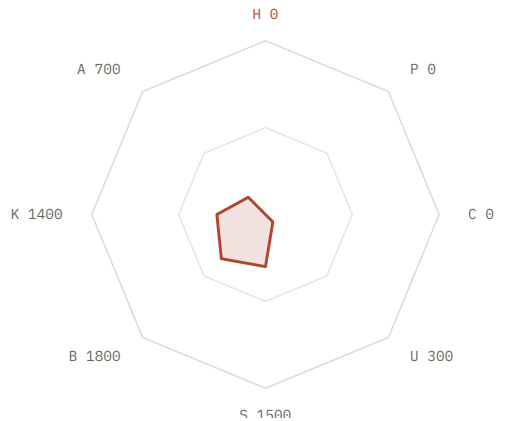

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	300
IFI-S	1,500
IFI-B	1,800
IFI-K	1,400
IFI-A	700

The **S of 1,500** meets a **B of 1,800** — a proud, dark bitterness, not an erased one — and the two settle into the balance a good cup of drinking chocolate has always wanted, the sweetness rising exactly to the cocoa's edge and no further. The **K of 1,400** is the round, coating body of a cup made with care, and the pinch of salt is doing quiet arithmetic: a little sodium measurably *raises* perceived sweetness and *lowers* perceived bitterness at once, so it earns its place twice. Notice there is no sucrose in any of it; the mogroside and the steviol glycoside carry the whole sweetness at a few parts per million, and the cocoa never had to be touched for the drink to come into balance.

The second pushes the same logic into something you can set on a plate: a bitter-sweet ganache, the base of a truffle, where a very high-percentage chocolate is made approachable without a grain of added sugar.

```
Bitter-Sweet Ganache · makes ~24 truffles · 20 min + chilling
— 200 g dark chocolate mass (85–90%, low-bitter cocoa), chopped
— 160 g oat cream (or heavy cream)
— mogroside-V monk fruit + Reb-M stevia, to S ≈ 1,400
— 1 pinch salt · unsweetened cocoa, for rolling
```

1. Warm the cream with the salt until it just steams; do not boil.
2. Pour over the chopped chocolate, wait a minute, then stir from the center out until glossy and smooth.
3. Stir in the sweeteners a little at a time, tasting, until the ganache reads deep and sweet-enough rather than sweet.
4. Chill until firm, scoop into rounds, and roll in unsweetened cocoa. The bitter coating is deliberate — it frames the sweetness inside.

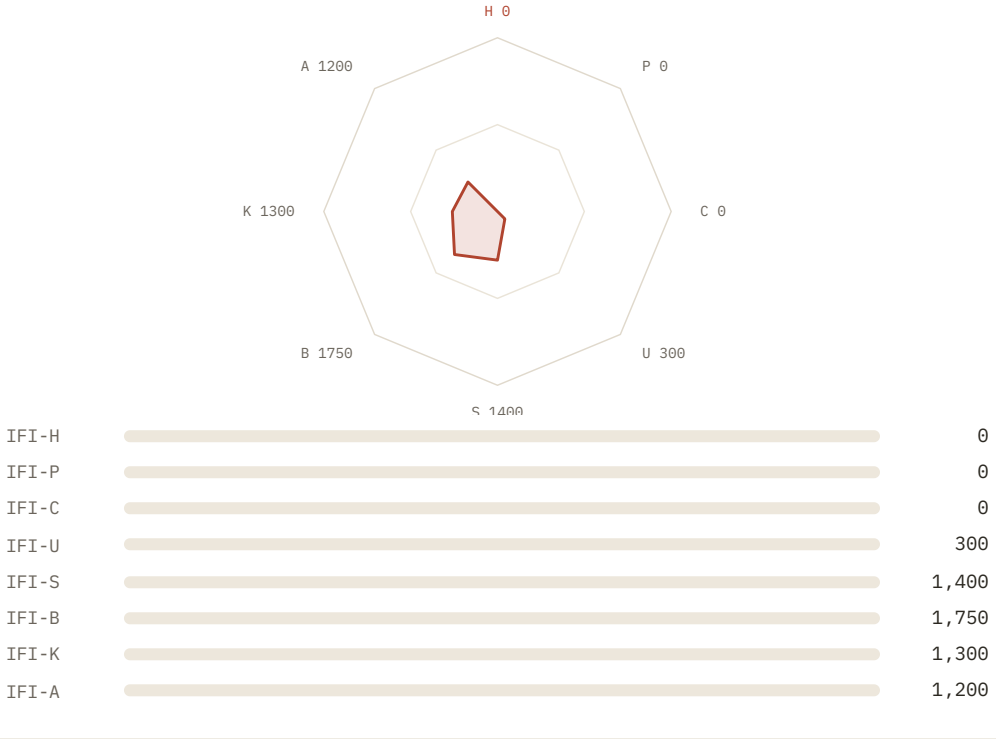

The truffle is a lesson in framing. Its **A of 1,200** and **B of 1,750** are the dark grip and the bitterness, held firm, and the unsweetened cocoa it rolls in raises them at the very first bite — so the **S of 1,400** inside lands against a bitterness freshly reasserted, and reads sweeter and cleaner for the contrast. This is composition in time, the same insight as the sweet-heat choreography of *Sweetness* (Volume VII) and the ordered delivery of a full-spectrum fruit: bitterness first, sweetness second, the argument staged rather than blurred.

The third is a nod across the shelf to *Swicy* (Volume XXVI), because chili and chocolate are the oldest sweet-heat pairing and the one place a trace of the heat axis belongs in this book. A few grains of a superhot carry warmth that arrives *after* the sweetness, the way capsaicin's slow TRPV1 burn always does, so the bite tells a story — sweet, then dark, then a late glow.

```
Chili Dark Chocolate Bark · makes 1 sheet · 15 min + setting
— 250 g dark chocolate mass (80%, low-bitter cocoa), tempered
— Reb-M stevia + mogroside-V monk fruit, to S ≈ 1,300
— a few grains of a superhot (SS-line), to H ≈ 400 — restraint
— flaky salt · toasted cacao nibs
```

1. Melt and temper the chocolate; stir the sweeteners through while it is fluid.
2. Fold in the barest amount of superhot — taste a fleck first; the heat should be a rumor, not an event.
3. Spread thin on parchment, scatter with salt and nibs, and set at cool room temperature.
4. Break into shards. The sweetness lands first, the cocoa's depth second, the warmth a slow third.

IFI PROFILE

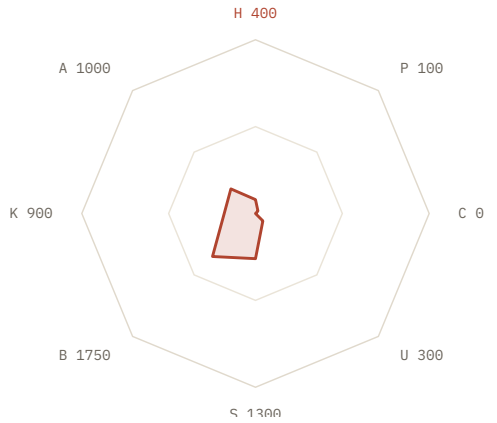

IFI-H	400
IFI-P	100
IFI-C	0
IFI-U	300
IFI-S	1,300
IFI-B	1,750
IFI-K	900
IFI-A	1,000

The **H of 400** is a warm glow and nothing more — the whole art of a chili chocolate is that the heat arrives as the *last* act, inside a sweetness and a depth that were generous first. Push it higher and you have a chili wearing chocolate's clothes, which this never should be. Across all three, the pattern holds: a cacao bitterness kept proud, a sweetness grown in to meet it at parts per million, and no sucrose asked to do the work the plant's own molecules now do. The dark stays dark. The sweetness simply lets you stay in it longer.

Chapter 11 · The Everyday Bar — *flavor that survives the formula.*

application · the nutrition matrix · plant-protein bitterness · sugar's hidden jobs · attribute-led

The nutrition bar is where flavor goes to die, and everyone in the category knows it. You are asked to carry a great deal of protein — whey, or the harder plant proteins, pea and soy, which bring a bitterness and a chalky back-note all their own — plus cocoa, plus whatever fiber the formula demands, and to make the result taste good while adding little or no sugar. The industry's answer has been to dump sugar back in under a dozen other names, or to over-sweeten with something that leaves a bitter metallic tail, which is the old quarrel of Chapter 1 all over again with a nutrition panel attached. The two-move design was built for exactly this matrix, and the bar is where it does its most useful work, because the bar is where the most bitterness and the least sugar are asked to coexist.

The key insight is that the bar's two main bitternesses run through the *same* receptor system and so answer to the *same* lever. The cocoa's bitterness we met in Chapter 9 — procyanidins and methylxanthines, read across the bitter family. The plant protein's bitterness is less famous but just as real: pea and soy isolates carry **bitter peptides** — short hydrophobic fragments released during the protein's processing and digestion — along with saponins and phenolic compounds, and these engage several of the same TAS2R sentinels, which is why a high-plant-protein bar so often reads bitter and chalky no matter how much cocoa is piled on to hide it. Because the cocoa and the protein both speak through the bitter receptor family, a single attenuation strategy — a low-bitter cocoa, and a plant-protein fraction selected or processed to shed its bitter peptides — quiets *both* bitternesses at once, and removes the very reason the formula was reaching for so much sugar in the first place. Then the sweetness is lifted with the plant's own sweet proteins and glycosides, at parts per million, so it arrives clean and leaves nothing behind.



Soybean — a plant-protein source and, in this library, a kokumi host. The bitterness a protein isolate carries and the richness a bar's missing sugar leaves behind are two separate problems, and the platform has a lever for each.

But bitterness is only half of what a bar loses when the sugar comes out, and a serious formulation respects the other half. Sugar in a bar was never only sweet; it was doing the quiet structural work catalogued in *Sweetness* (Volume VII) — **bulk**, the physical mass that holds a bar together; **binding and moisture**, the hygroscopic hold that keeps it from crumbling and staling; **browning**, the Maillard color and toffee notes; and a share of the **body** the mouth reads as richness. A high-potency sweetener dosed at parts per million gives sweetness and *nothing* to fill the holes the sugar left. So the bar's reformulation is a system, not a swap: the sweetness is the easy part, and the real work is restoring structure. **Allulose** does much of it — a rare sugar that bulks, browns, holds moisture, and reads about seventy percent as sweet as sucrose while the body carries it lightly, which makes it the quiet backbone of serious bar reformulation. **Date paste** binds and adds a whole-food body (we use it for hold, not sweetness). A **kokumi** base, via the CaSR axis of Chapter 15, puts back the perceived richness. And the nut butter brings fat-borne body that reads as satisfaction. Build all of that under a quieted bitterness and a clean, grown-in sweetness, and the bar stops tasting like a compromise.

```
The No-Sugar Cocoa Bar · makes 12 · 20 min + chilling
— 120 g rolled oats, half toasted and pulsed to a coarse flour
— 90 g low-bitter plant protein (a de-bittered pea or soy isolate)
— 40 g low-bitter dark cocoa
— 120 g unsweetened almond or peanut butter
— 80 g date paste (for binding, not sweetness)
— Reb-M stevia + mogroside-V monk fruit blend, to S ≈ 1,400
— 20 g allulose · 60 g water, as needed · 1 pinch salt
```

1. Toast the oats dry until they smell nutty; cool. Pulse half to a coarse flour, leave the rest whole.
2. Stir the cocoa, protein, allulose, salt, and sweetener blend through the oats, so the two bitterness-carriers are evenly met by the sweet before any liquid goes in.
3. Warm the nut butter and date paste together until loose; fold into the dry mix, adding water a little at a time until it just holds a firm press.
4. Press hard into a lined tin, chill until set, and cut into twelve. They keep a week, and they never taste like a compromise.

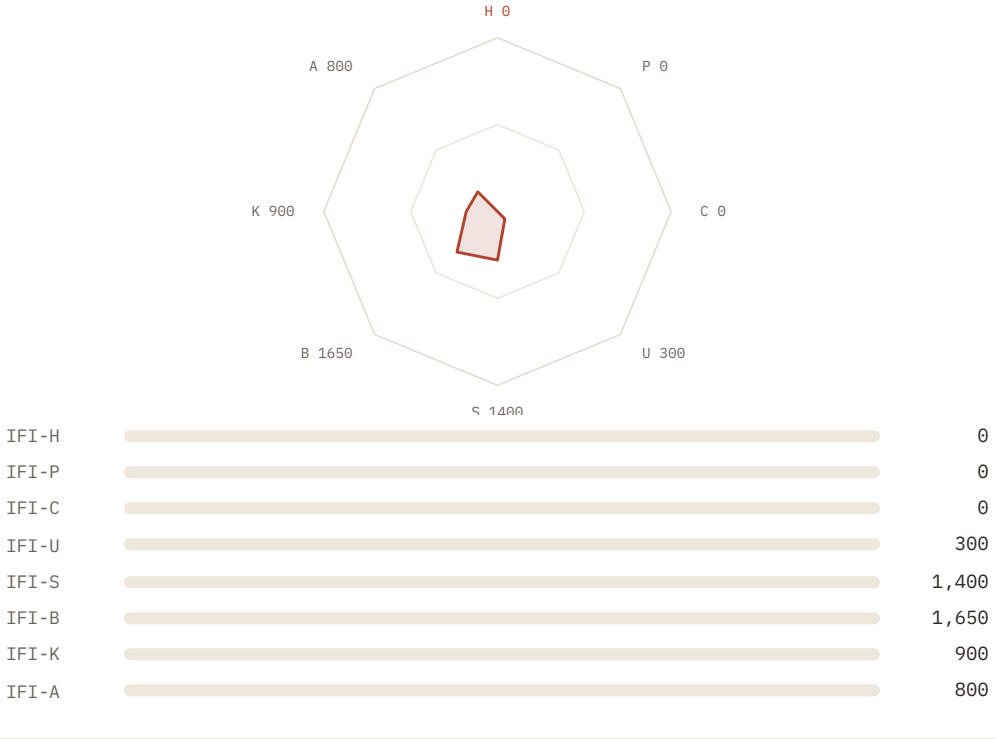

Read the two numbers at the center. The **B of 1,650** is the cocoa *and* the pea protein quieted together by a single attenuation move, down to a structure rather than a warning — without a grain of sugar sent in to cover them. The **S of 1,400** meets it, clean and residue-free, grown from the sweet keys of Chapter 14. And the **K of 900** is the floor rebuilt — the body the missing sugar would have given, reconstructed from the nut butter, the oats, and the kokumi base, so the bar reads full and satisfying instead of dense and dry. The two moves, level and reinforcing, are the whole reason a protein-forward, no-added-sugar bar can taste like something a person reaches for a second time.

A second bar makes the same point in a leaner, higher-protein register — the kind a formulator is asked for when the protein number has to lead and the sugar number has to be zero.

```
Seeded Protein Bar, Zero Added Sugar · makes 10 · 25 min + chilling
— 140 g low-bitter plant protein (de-bittered pea/soy blend)
— 60 g low-bitter cocoa
— 80 g mixed toasted seeds (pumpkin, sunflower)
— 100 g almond butter · 70 g allulose syrup
— 40 g soluble fiber (inulin or resistant dextrin)
— Reb-M stevia + mogroside-V monk fruit, to S ≈ 1,400 · 1 pinch salt ·
water as needed
```

1. Combine the dry protein, cocoa, salt, fiber, and sweeteners; stir so the bitterness is met evenly.
2. Warm the almond butter and allulose syrup until pourable; fold in with the seeds.
3. Add water by the spoonful until the mass just coheres under a firm press.
4. Press into a lined tin, chill until firm, and cut into ten. The fiber and allulose carry the body; the seeds carry the bite.

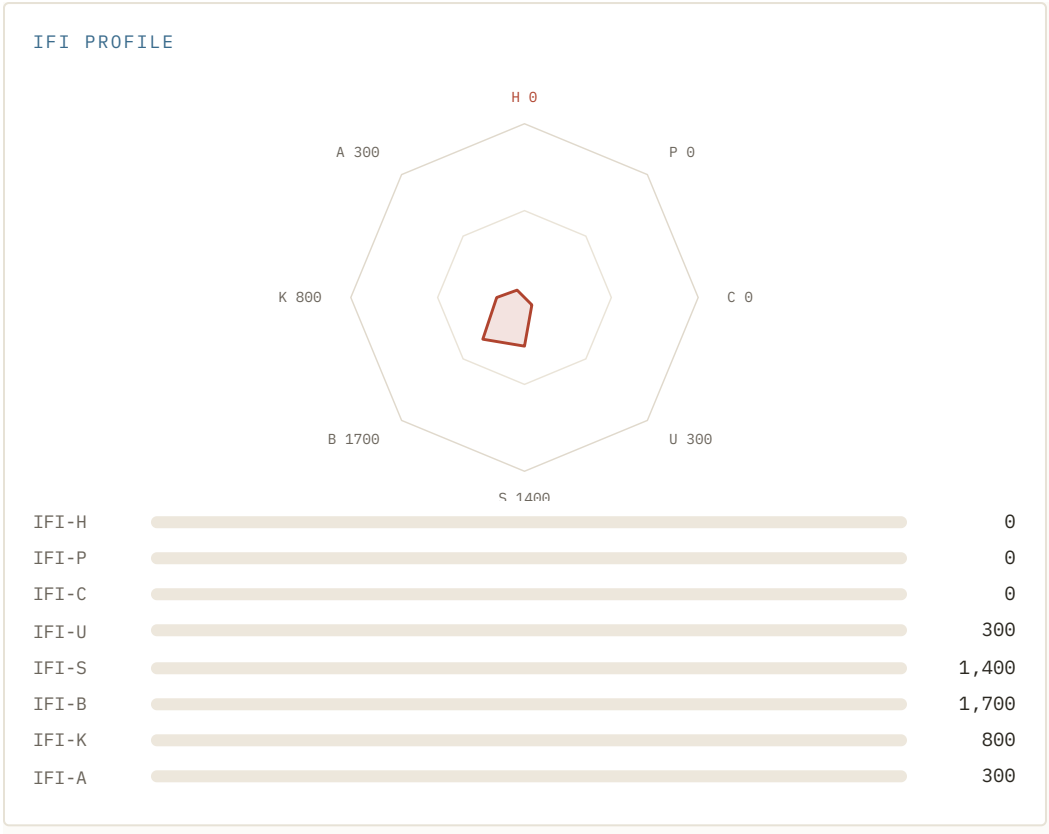

What a formulator can promise from either bar is bounded and true, and the line matters enough to draw in bright ink: **attribute-led and nothing more.** Protein-forward. No added sugar. A flavor that survives the formula. Those are compositional attributes with legal meanings, and they are the only kind of claim this book makes. We do not say a bar *delivers* anything, *does* anything for a body, or belongs to any outcome beyond its taste — those are not sentences we have earned the right to write, and the moment flavor science starts making them it stops being either good science or honest food. The nourishment ethos that runs under this whole library lives in *why* we did the work — because a nourishing food refused for its flavor helps no one, and a bar people actually finish is a bar that did its job — never in a promise printed on a wrapper. We made the flavor honest, and we made it survive the formula. What the bar is for beyond that is the eater's business, not our claim. The full treatment of attribute-led design for the better-for-you shelf is *The Functional Kitchen* (Volume XL); here the discipline is simply held, because a bar is exactly where the temptation to overclaim is strongest and the reason to resist it is clearest.



PART V OF VI

The Garden & the Kept Bitterness

greens that read sweet · the bitterness we keep · the sweet keys

Chapter 12 · Greens That Read Sweet — *the vegetable a child finishes.*

DOMAIN · BRASSICA · the glucosinolate engine · GS-OH and MYB28 · the surgeon's rule

The vegetables a child is told to finish are, almost to one, the vegetables evolution built to taste like a warning. The brassicas — broccoli, kale, cabbage, the Brussels sprout that became cultural shorthand for a food endured — defend themselves with a chemistry so effective that a century of cheese sauce and long braising has never quite talked a table past it. And then, having made themselves bitter to survive, they were bred for shipping and uniformity until most of the sweetness that once stood against the bitterness was gone too, so what arrives under a grocery light is all alarm and no invitation. This is the single clearest case for the two-move design in the whole garden: quiet the warning at its source, let the vegetable's own sweetness be heard, and a different plant comes to the table — the one it tasted like before we asked it to survive a truck.

The warning is built by **glucosinolates** — sulphur-rich, nitrogen-bearing thioglucosides that sit quietly and harmlessly inside the intact cell, stable and nearly tasteless as stored. The weapon arms only on damage: when a molar crushes the leaf, the glucosinolates meet **myrosinase**, an enzyme the plant keeps locked carefully apart from them, and in milliseconds the two react, stripping the glucose and rearranging the molecule into a sharp **isothiocyanate** or the cyclized compound **goitrin** — the molecules TAS2R38 reads most sharply as *stop.* The bitterness a child tastes in raw broccoli is not a property of the floret on the plate; it is a chemical reaction the child's own chewing sets off. The glucosinolates are built, in the main, on the aliphatic route from the amino acid methionine, through a dozen named enzymes and under a set of master-regulator throttles — the **MYB28/MYB29/MYB76** transcription factors — that set how hard the whole branch runs. Dial a MYB down and the vegetable floods itself with less glucosinolate; the engine comes with its own throttle already wired in.



Broccoli, whole and halved — the glucosinolate load sits highest in the florets and growing tips, the very tissue we most want to eat, which is why the brassica is the surgeon's problem of the axis.

But the crude move — turn the MYB throttle down and let the whole branch quiet — is exactly the trap this intersection exists to avoid, because not all of the glucosinolate output is harsh, and some of it is the very reason a careful eater reaches for the vegetable at all. The finest problem on the brassica is a fork, and both tines come off the same side-chain stage. Down one tine, the enzyme **GS-OH** hydroxylates an alkenyl glucosinolate into **progoitrin**, which on hydrolysis cyclizes into **goitrin** — one of the harshest, most persistent, most child-repelling bitternesses the brassicas make. Down the other tine sits **glucoraphanin**, which hydrolyzes to a far gentler isothiocyanate and is the compound most eaters are quietly glad to be getting. The two branch from the same upstream pool, and **GS-OH is the switch between them.** So the selective lever almost designs itself: reduce flux down the GS-OH tine, collapsing the progoitrin-to-goitrin arm, while leaving the glucoraphanin tine and the rest of the plant's defenses intact. Where a vegetable's total load still sits too high after that, we add a *gentle* trim at the MYB28 throttle — a dosage reduction, not a knockout — lowering the whole branch by a designed fraction while keeping its proportions. This is what we mean, across the library, by the **surgeon's rule**: not the biggest cut that quiets the plant, but the smallest cut that quiets only what needed quieting. The full account is the heart of *Zero-Bitter* (Volume IX); here it is the bitter half of a balance, because once the warning is down we lift the vegetable's own sweetness to meet it, and the plate changes character entirely.

The method shows up in the numbers. Across the brassica domain the library holds **380 broccoli constructs, 280 in the broccoli-kale cross, and 40 each in cabbage, kale, radish, and rapeseed**, a deliberate spread because the optimal valve settings differ by species and by the balance of glucosinolates each one carries:

ORGANISM · CONSTRUCTS	PRIMARY LEVER	NATIVE IFI-B	ENGINEERED IFI-B	ATTENUATION
Broccoli · 380	GS-OH knock-down + MYB28 −1	4,600	1,750	~62%
Broccoli-Kale · 280	GS-OH + MYB28/29 trim	4,400	1,760	~60%
Kale · 40	GS-OH + ESP tune	4,300	1,720	~60%
Cabbage · 40	progoitrin-arm trim	3,900	1,750	~55%
Radish · 40	glucoraphenin-arm trim	3,700	1,850	~50%
Rapeseed · 40	progoitrin reduction	3,500	1,820	~48%

The brassicas sit at the generous top of the twenty-to-sixty-five band, because this is the family where the harsh fraction and the prized fraction separate most cleanly. Now the kitchen, where the balance is served. These dishes are plain on purpose; the cleverness a brassica used to demand is unnecessary once the warning is down at the source. The first is cabbage taken as hot and hard as a pan allows — the treatment that, on a conventional head, drives the sulphur notes to their worst, and on a low-bitter one does the opposite.

Blistered Cabbage, Brown Butter, a Drop of Sweet · serves 4 · 25 min
— 1 small head low-bitter cabbage (SS-line), in thick wedges, core kept on
— 40 g butter · 1 tbsp olive oil
— flaky salt · 1 lemon
— a few drops engineered sweetener (Reb-M stevia), if the cabbage is very lean

1. Heat the oil in a heavy pan until it just smokes. Lay the wedges cut-side down and leave them, undisturbed, until they blister dark.
2. Turn once, add the butter, and let it foam and brown to a nutty amber, spooning it over the wedges.
3. Salt well. Finish with a hard squeeze of lemon and, if the cabbage is lean, a few drops of sweetener — not to sweeten it, only to meet the last of the bitterness.

IFI PROFILE

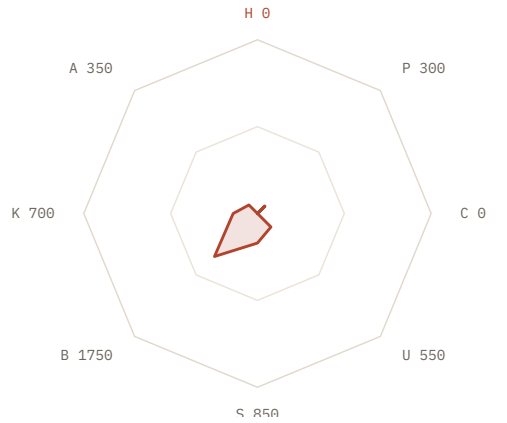

IFI-H	0
IFI-P	300
IFI-C	0
IFI-U	550
IFI-S	850
IFI-B	1,750
IFI-K	700
IFI-A	350

The **B of 1,750** is the glucosinolate alarm pulled down about fifty-five percent, and look what rises into the quiet it leaves: an **S of 850**, the cabbage's own sugars, now audible and nudged the last step by a few drops of grown-in sweetness; a **K of 700** from the brown butter's kokumi, giving a simple dish the body of one that simmered for hours; a deliberate **P of 300**, the clean mustardy edge we *kept,* because a brassica with no sharpness has lost its nerve. The blister

reads as caramel, not char, because the signal that would have read as scorch was turned down at the enzyme, not cooked over at the stove.

The second is broccoli seared hard, and the third is raw kale — the vegetable most often eaten as penance — dressed and eaten with pleasure.

```
Charred Low-Bitter Broccoli · serves 4 · 18 min
– 2 heads low-bitter broccoli (SS-line), cut into flat-faced spears
– 2 tbsp cold-pressed high-smoke-point oil
– 2 cloves garlic, paper-thin · 1 lemon, zest and juice
– flaky salt · a pinch of chili, optional
```

1. Heat a heavy skillet until a drop of water skates and vanishes. Dry the broccoli thoroughly; wet broccoli steams instead of searing.
2. Add the oil, then the broccoli cut-face down in one layer, and leave it three to four minutes until deep brown, almost black at the edges.
3. Turn once, add the garlic at the pan's edge to toast, and give it two minutes more.
4. Off the heat, hit it with the lemon juice, then the zest, salt, and optional chili. Serve while the edges still crackle.

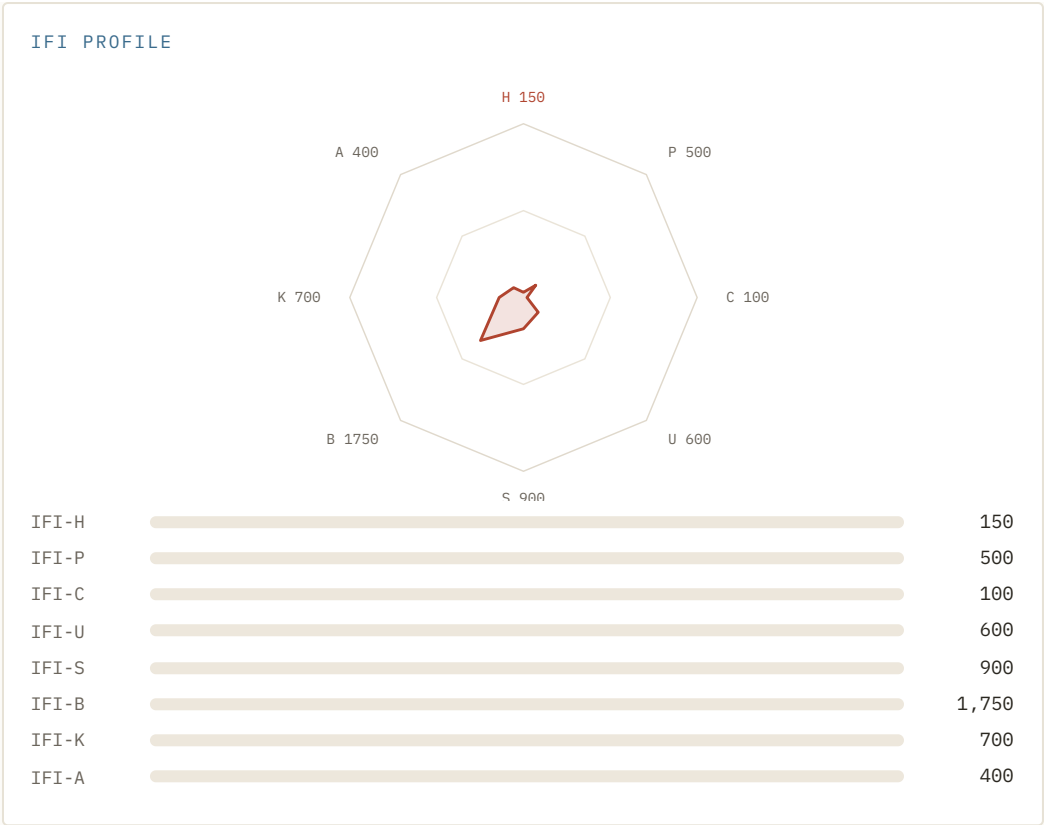


```
Massaged Low-Bitter Kale · serves 4 · 15 min + 10 to rest
– 1 large bunch low-bitter kale (SS-line), stems stripped, leaves torn
– 1 tbsp lemon juice + 2 tbsp olive oil · 1 tsp sea salt
– 2 tbsp toasted seeds · shaved hard cheese or a handful of currants, to
finish
```

1. Put the torn leaves in a wide bowl with the salt, lemon, and oil.
2. Massage with your hands two to three minutes — the leaves darken, soften, and nearly halve in volume.

3. Rest ten minutes so the dressing works through. Fold in the seeds and your finish just before serving.

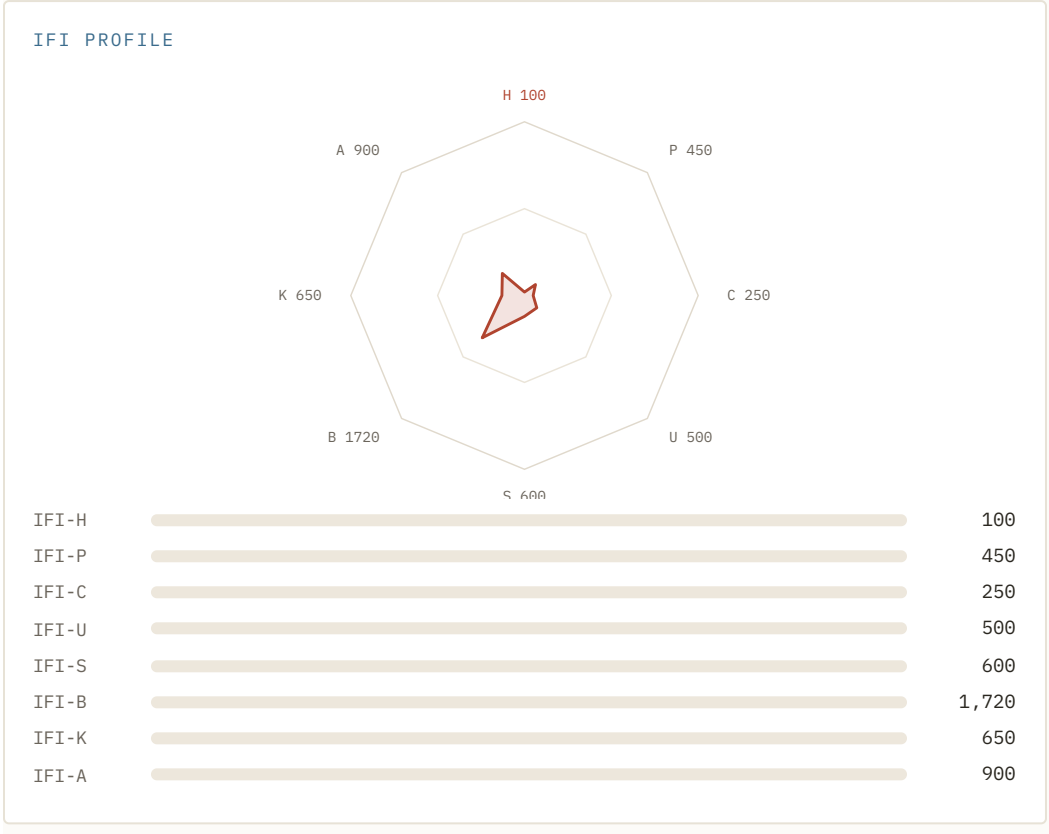

The **S of 900** on the broccoli and the uncovered sweetness across all three are the invitation half of the move, grown free as the warning falls. And the kale's **A of 900** is worth a word, because it is what we chose *not* to touch: a good kale keeps a structural, almost mineral grip — astringency, not bitterness — and that grip is part of why the leaf feels substantial rather than slack. We quieted the bitter thread of the braid and left the astringent one standing. Set any of these in the middle of a table of strong tasters and non-tasters alike — the PAV and the AVI of Chapter 2 — and the thing you notice is that, for once, everyone eats the same amount. The warning that used to divide the table by genotype has been turned down below the threshold that divided it, and a drop of grown-in sweetness closes what little gap remains. That is the garden's version of the settled argument: a vegetable no longer endured, but wanted.

Chapter 13 · The Bitterness We Keep — *tea, hops, grape, and the useful edge.*

DOMAIN · TEA 280 · HOPS 80 · GRAPE 1,610 · ORANGE 120 · the dial, not the switch

We must be careful not to leave the impression that bitterness is a flaw to be erased. It is not. A great many things we love are bitter on purpose, and the whole art of this intersection is knowing *how far down to turn the alarm* — which, in some of the most carefully worked plants in the library, is not very far at all. The two-move design is a dial, not a switch, and some of our finest work is in the plants where we barely touch the bitter, and lift the sweet only a little, because the bitterness is the structure and the structure is the point. Here the move inverts: we tune the bitterness *to* a target rather than away from one, and meet it with exactly as much sweetness as it needs to read as depth instead of punishment — never enough to silence it.

Tea is the clearest teacher, and the most intricate, because it lives on three axes at once — a savory umami, a structural astringency, and a bitterness woven through both — which is why it carries **280 constructs** and why its story is as much about restraint as about edits. The grip of a strong or over-steeped cup is partly bitterness and partly a drying pucker, and both come from the **catechins**, the flavan-3-ols of the leaf, chief among them EGCG (epigallocatechin gallate) along with EGC, ECG, and EC. The galloylated catechins, EGCG above all, are the most aggressively bitter-and-astringent; the gallate group is much of what makes them grip. The lever is the catechin profile itself — tuning the galloylation and the balance among the catechins to ease the roughest, EGCG-led harshness by roughly thirty percent — while protecting two things carefully. The first is the astringent **backbone** that gives a tea its length; a tea with no grip is slack and forgettable. The second is **theanine**, the amino acid nearly unique to the tea plant, which carries the round, savory, almost broth-like sweetness a fine green tea is prized for, and which we leave entirely intact, so that quieting the harsh catechins actually lets the theanine *step forward.* A representative green tea engineered to forgive a long steep, with a little monk-fruit lift to meet the kept bitterness:

IFI PROFILE

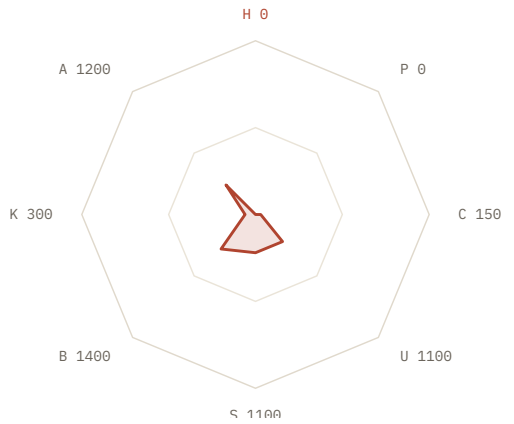

IFI-H	0
IFI-P	0
IFI-C	150
IFI-U	1,100
IFI-S	1,100
IFI-B	1,400
IFI-K	300
IFI-A	1,200

The **B of 1,400** is the harshest catechins down about thirty percent — the lowest attenuation in the book, because here we *shape* a wanted bitterness rather than strip it — and the **A held at 1,200** is the drying backbone kept for length. The **U of 1,100** is the theanine, uncovered and grown forward into the quiet, and the **S of 1,100** is a grown-in sweetness lifted just to the point where it meets the structural bitterness as balance rather than erasure. A tea you can steep without fear is a tea that gives more of itself; the full tannin story belongs to *Astringency* (Volume XIII) and to *Tea* (Volume XVII).



Tea in cross-section — the catechins that grip a strong cup and the theanine that rounds it ride the same leaf. The craft is to quiet the first thread, keep the second, and lift a little sweetness to meet what is left standing.

Hops are the same lesson in a glass, and the most extreme case of a bitterness we would never remove — because in beer, bitterness is deliberately *added,* measured, and bragged about on the label as **IBU.** The **80 hop constructs** are not about less bitterness but about the *shape* of a bite: the complaint a brewer actually has is never "too much" but "lands wrong" — harsh where it should be firm, clinging where it should finish clean. The bitterness is the work of the **iso-alpha-acids**, formed in the boil from the cone's alpha-acids (humulone, cohumulone, adhumulone), and its *quality* depends heavily on the fraction that is **cohumulone**: a high-cohumulone hop gives a harsher, sharper, more lingering bite; a hop weighted toward humulone gives a rounder, cleaner one. The alpha-acids are built by a lineage-specific polyketide route — **VPS** (valerophenone synthase)

laying down the skeleton, then the prenyltransferases HlPT1 and HlPT2 — and the cohumulone branch diverges right at the start, on the choice of starter unit. Tune the starter supply or the VPS preference and you weight the cone toward the rounder acid: the same IBU, landing softer. The aroma stays wholly intact, because it is built on the separate isoprenoid road of the hop's terpenes, which we never touch. Here the sweetness we lift is barely anything at all — a beer is not meant to be sweet — and the balance is almost all on the bitter side of the dial, shaped rather than lowered. Hops are where this book's method points at a bitterness the drinker is *proud* of, and asks only that it finish clean.

Grape teaches the lesson's sharpest edge: sometimes the right amount to turn an axis is almost none. It is one of the most heavily worked organisms in the whole library — **1,610 constructs** — but its signature is **tannin,** not bitterness. Its defining compounds are the seed and skin proanthocyanidins, large procyanidin polymers that read on the tongue far more as astringency than as the TAS2R38 warning; the pocket that reads bitter is too small for a big polymer, which instead cross-links saliva into structural dryness. In a glass of wine a little bitterness is *structure,* and removing it gives you a flabby, spineless wine nobody wants. So on grape we touch the bitter axis barely, and only to lift a green, unripe, stalky harshness off a young wine, then hand the real work — the whole architecture of tannin and mouthfeel — to *Astringency* (Volume XIII) and *Grape & Wine* (Volume XVIII). Knowing when *not* to turn an axis is as much a part of holding it as any edit.

And **citrus** closes the family with the brightest bitternesses of all, and the clearest bitter-sweet uncovering outside of coffee. In an orange — **120 constructs**, scored B and C — the harsh notes are two, and both fight the fruit's own sugar: **naringin,** a flavanone glycoside, the bitterness of the white pith and membrane; and **limonin,** a limonoid on a separate triterpenoid route, the source of citrus's notorious *delayed* bitterness, the way a fresh juice can taste sweet at breakfast and turn bitter by lunch as limonin slowly forms. We tune the flavanone route that builds naringin, easing the pith's grip about forty-eight percent, and moderate the limonoid supply so the juice holds its brightness through the afternoon. A whole-pressed orange, juiced with a little of its membrane the way that keeps the most of the fruit:

IFI PROFILE

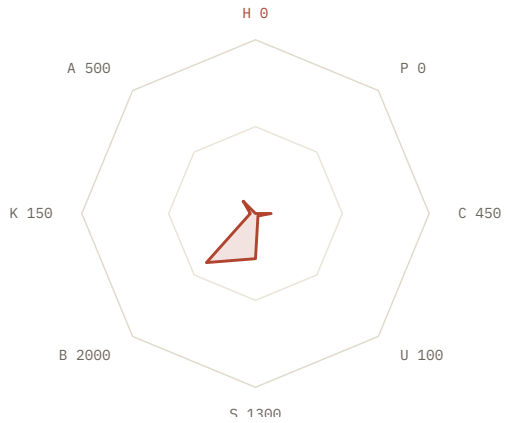

IFI-H	0
IFI-P	0
IFI-C	450
IFI-U	100
IFI-S	1,300
IFI-B	2,000
IFI-K	150
IFI-A	500

The **S of 1,300** is the fruit's own sugar, *uncovered* rather than added — once the pith's bitterness comes down, the sweetness that was always there is finally audible, the clearest everyday proof of Chapter 4's gift. The **C of 450** is the cool, lifting citrus brightness carrying the glass, and the **B at 2,000** is the pith eased just enough that the fruit no longer turns against itself by midday. Across tea, hops, grape, and citrus the pattern holds and inverts the garden's: a bitterness kept proud and met, not a bitterness lowered and replaced. The dial turned the other way is still the same design — two voices, taught to settle, with neither one asked to leave the room.

Chapter 14 · The Sweet Keys — *sweetness that behaves.*

DOMAIN · STEVIA 910 · MONK FRUIT 80 · SWEET PROTEINS · SUGARCANE 200 · the off-note that is itself bitter

The sweet half of every balance in this book rests on a small family of remarkable molecules, and they do not all behave alike. Sugar is a blunt, honest sweetener: it arrives fast, peaks, and leaves, and it must come by the gram. The sweeteners we grow instead are subtler instruments, read by the receptor at a thousandth or a millionth of sugar's dose, and choosing among them — matching the key to the bitterness it has to meet — is most of the craft of lifting sweetness without a trace of sugar behind it. And there is a deep irony waiting in the first of them, one that makes this chapter belong in *this* book above any other: the cleanest natural sweetener is *itself* a bitter-sweet problem, and solving it is the two-move design pointed inward, at the sweetener's own chemistry.

That sweetener is **stevia**, our workhorse — a **910-construct** library, the largest sweet-axis program after the tomato. Its sweet molecules are **steviol glycosides**, a diterpene core hung with glucose units, built not from sugar but from the plant's isoprenoid machinery: the plastidial MEP pathway converges on geranylgeranyl diphosphate (GGPP), which a defined cascade folds and oxidizes into steviol, after which a family of UDP-glycosyltransferases — **UGTs** — adds glucose units one at a time, in a specific order and specific linkages. The identity of the final glycoside, and therefore both its sweetness *and its bitterness,* is decided entirely by which glucoses get added where:

STEP	ENZYME	PRODUCT
1–4	CPS → KS → KO → KAH	**steviol** (branches off the gibberellin route)
5	UGT85C2	steviolmonoside
6	UGT91D2	steviolbioside
7	UGT74G1	**stevioside**
8	UGT76G1	**rebaudioside A**
9	UGT91D2	rebaudioside D
10	UGT76G1	**rebaudioside M**

Here is the irony, and it is a receptor fact, not an impurity. **Stevioside** (step 7), the dominant glycoside of an ordinary leaf, does not only bind the sweet receptor — it also activates two *bitter* receptors, **TAS2R4 and TAS2R14**, and that bitter signal arrives riding alongside the sweet one. This is the licorice-and-lingering-bitterness people mean when they say they "can taste the stevia," and recall from Chapter 2 that a simultaneous bitter signal does not merely add a bad taste: through the shared wiring, it actively *subtracts* from the perceived sweetness, flattening it. So the cleanest stevia is not stevia with a masker bolted on. It is a leaf engineered to walk the pathway *forward* — past stevioside, past Reb A, all the way to

rebaudioside M (step 10), which carries two extra β(1→3) glucoses that spoil the fit at the bitter receptors while preserving it at the sweet one. Reb M barely touches TAS2R4 and TAS2R14 and keeps nearly all the sweetness. The bitterness and the sweetness, in other words, can be pulled apart *by molecular structure* — the two-move design running inside a single family of molecules. Our stevia program drives the UGT selectivity (chiefly **UGT76G1** for the β(1→3) steps and an engineered **UGT91D2** for the β(1→2) step toward Reb D) to shift the leaf's output from a bitter-tinged stevioside majority toward a clean Reb-M one:

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



Stevia in cross-section — the sweetness is a diterpene dressed in glucose, and the licorice bitterness is a matter of exactly how many glucoses, in which linkage. We breed and edit the leaf forward until the bitter molecule is simply no longer the one it makes.

Monk fruit is stevia's rounder cousin, and the reason so many balances in this book call for the two together. Its sweeteners are **mogrosides** — cucurbitane *triterpene* glycosides on a thirty-carbon skeleton, built from the mevalonate pathway up through cucurbitadienol (by cucurbitadienol synthase), oxidized to **mogrol** by P450s including **CYP87D18**, then glycosylated by UGTs up the ladder to **mogroside V**, roughly two hundred times sweeter than sucrose by weight, heat- and pH-stable, and essentially free of stevioside's licorice burden. The **80-construct** program is small because the slow perennial vine is hard to transform and its chemistry needs less rescuing; the work concentrates on driving flux to mogroside V and siamenoside I, and on using monk fruit as the

rounding partner. Here is the blend mathematics that earns it a place in nearly every recipe above: stevia's Reb M is fast and clean at the front of a sip; monk fruit's mogroside V is slow and round through the middle; sucrose's own arc spans both. Compose a stevia-and-monk-fruit blend and you can trace an arc close to sucrose's own — the stevia covering the onset, the monk fruit the body — each dosed low enough that neither off-note surfaces, and the two synergizing at their separate receptor sites. Two engineered crops, each grown to be the other's missing half.

Beyond the glycosides lie the two frontiers. The **sweet proteins** — brazzein chief among them — are the loudest clean sweetness in nature, read by the receptor at parts per million by the wedge mechanism of Chapter 2. Brazzein is a single 54-amino-acid chain locked by four disulfide bonds into near-indestructibility, stable through heat and acid, roughly a thousand times sweeter than sugar by weight and far more per molecule, and the platform's lead because its small, rigid fold can be re-sculpted by the ESM2 pipeline without collapsing. It is honestly a frontier: a novel sweet protein in a food crop is a new-to-market ingredient, and brazzein does not yet carry broad approval, so the sweet-protein work sits on the roadmap, its edits model predictions and its destination a fruit that makes brazzein in its own tissue at a few parts per million. And **sugarcane** — a **200-construct** program — is not reinvented but reconsidered: the world's densest sugar factory is also the ideal *chassis,* a C4 plant whose whole physiology packs an enormous mass of soluble sweet storage, redirected to grow high-potency, low-load sweeteners from the same fields and mills the world already runs. The full account of all four keys is *Sweetness* (Volume VII); the table below is the working summary a formulator needs:

SWEETENER	SOURCE · CONSTRUCTS	SWEET COMPOUND	WHAT IT BRINGS TO THE BALANCE
Stevia	Stevia · 910	steviol glycosides (Reb M)	the high, clean front; the irony solved by walking the pathway forward
Monk fruit	Monk Fruit · 80	mogrosides (mogroside V)	the round, slow middle — the bitter-sweet workhorse
Brazzein	sweet-protein class (frontier)	brazzein / thaumatin	parts-per-million potency, a long held note; roadmap, not yet approved
Sugarcane	Sugarcane · 200	native sugars, redirected	the familiar shape of cane, at a lighter load, from existing fields

The point of a toolkit this varied is that balance has a *texture,* not just a level. A coffee wants the barest high lift and nothing lingering, so a touch of Reb M and no more. A drinking chocolate wants the round, patient sweetness that sits with the cocoa, so monk fruit carries the middle. A bar wants a clean sweetness that survives being pressed and kept, so a Reb-M-and-mogroside blend over an allulose floor. A tea wants a whisper that meets a kept bitterness without burying its theanine. Match the key to the bitterness it has to meet — its onset, its peak, its decay set against the bitterness's own shape in time — and the sweetness behaves: arriving where it is wanted, leaving when it is done, and never once becoming the sugar it replaced.

That is the sweet half of the two-move design, and it is, at its root, the same move as the bitter half — walk the pathway to the molecule that reads the way you want, and starve the one that does not.



PART VI OF VI

Composition, Proof, and the Table

the floor that holds a sweetness up · the reformulator's playbook · one fruit in balance · measured, never claimed

Chapter 15 · Building the Floor — *what holds a sugar-free sweetness up.*

IFI-K · CaSR · the γ-glutamyl peptides · allulose · the cross-axis levers

Taste a high-potency sweetener dissolved in plain water, meeting a quieted bitterness, and you can still learn, in one sip, the deepest problem in sugar reduction — deeper than any off-note, and the one most likely to sink a bitter-sweet reformulation that got both of its headline numbers right. The sweetness is *there,* correct and loud, the bitterness is *down,* and yet the whole thing can taste thin, hollow, fleeting — a balance played with no instrument behind it. This is the lesson every chapter so far has gestured at, now stated plainly: sweetness was only ever one of sugar's jobs, and when you remove the sugar you open five or six other vacancies at once. A bitter-sweet drink or bar that fills only the sweetness and the bitterness, and leaves the rest empty, reads as a clever absence rather than a whole food. Rebuilding the floor is the quiet, decisive third move, and it is mostly not about sweetness at all.

Consider what a spoonful of sucrose did in a food besides closing the Venus flytrap. It provided **bulk**, physical mass and structure. It provided **body and mouthfeel**, raising a liquid's viscosity so the tongue reads thickness as richness — which is why a diet drink tastes "watery" not for lack of sweetness but for lack of dissolved solids, a texture problem wearing a flavor costume. It controlled **water and shelf life**, binding free water so microbes had nothing to drink. It made **color and roasted flavor** through caramelization and the Maillard reaction. And it governed **physical state** — the freezing point of a frozen dessert, the glass transition of a crisp edge. A sweetener dosed at parts per million gives sweetness and none of these. So a serious bitter-sweet formulation is a system: a high-potency sweetener for the taste, a bitterness quieted at the source, and then a deliberate reconstruction of the floor.

The structural subfloor is **bulking agents** — allulose, erythritol, soluble fibers such as inulin and resistant dextrin — which restore the dissolved-solids density and some viscosity. Allulose is worth singling out, because it is the rare ingredient that does several of sugar's jobs at once: it bulks, it browns, it holds moisture, and it reads about seventy percent as sweet as sucrose, which is why it anchors the cocoa bars of Chapter 11 and the drinking chocolates of Chapter 10. On top of that subfloor sits the perceptual work, and it belongs to a different axis entirely — **kokumi.**

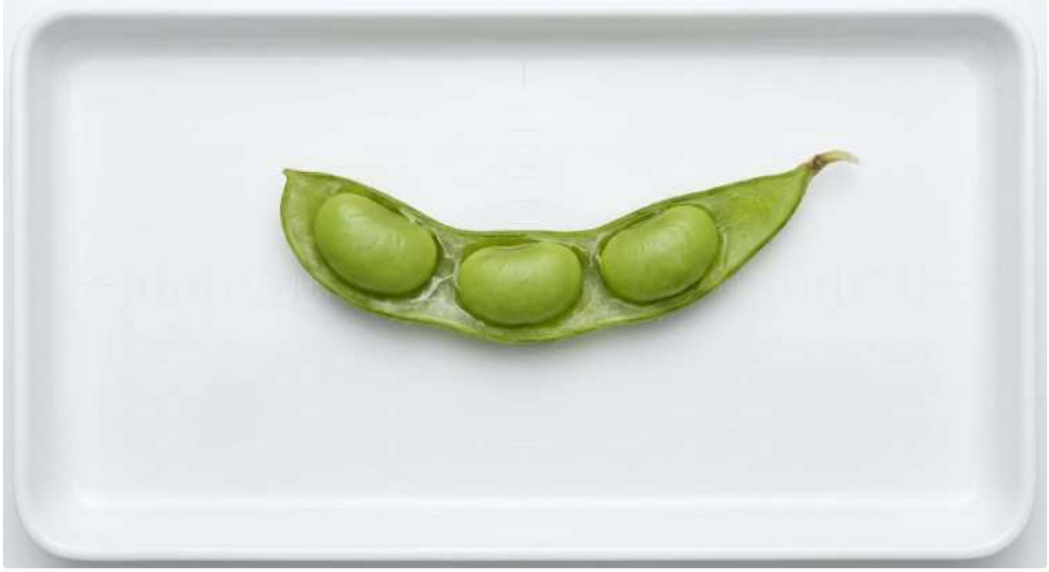

Soybean, a kokumi host — the γ-glutamyl peptides that give a sugar-free sweetness its missing body are engineered here, on CaSR, a different receptor and a different axis from the sweetness itself.

Kokumi, **IFI-K,** is the axis of richness, continuity, and mouthfulness — the sensation of a food being *full* and *sustained* rather than thin and quick. Its receptor is the **calcium-sensing receptor, CaSR,** another class C GPCR like the sweet receptor itself, and its natural agonists are small **γ-glutamyl peptides** — glutathione (γ-Glu-Cys-Gly) and, more potently, **γ-glutamyl-valyl-glycine** — that occur in aged, fermented, and long-cooked foods. The remarkable part, and the reason kokumi belongs at the center of a bitter-sweet book, is that these compounds produce almost no taste of their own, yet when present alongside a sweetness they *amplify and prolong* it. They do not add sweetness; they give the sweetness a body to live in. A stevia-sweetened, de-bittered drink built over a designed kokumi base reads rounder, fuller, and — crucially — its sweetness *lasts the right length of time* instead of flashing and vanishing. We engineer the peptides' biosynthesis in kokumi hosts such as soybean and yeast and profile the result on IFI-K, so that "richness" stops being an adjective and becomes a dose. The full axis is *Kokumi* (Volume X); here it is the floor under every balance in this book.

And then there are the cross-axis levers — the quiet arithmetic that lets a formulator build a sweetness from its whole neighborhood rather than from sweetener alone, and they matter doubly in a bitter-sweet food because they move both of this book's axes at once. **Salt** in small amounts measurably *raises* perceived sweetness and *suppresses* bitterness, which is why a pinch of flaky salt finishes the drinking chocolate and the cabbage: it is doing bitter-sweet work on both sides of the quarrel with no sweetness and no saltiness of its own. **Acid** *suppresses* perceived sweetness, so trimming a formula's sourness can replace a share of its sweetener — sometimes the cheapest sweetness is the acid you remove. And **warmth** lifts sweetness, because TRPM5, the shared channel of Chapter 2, is temperature-sensitive and opens more readily when warm, so the same formula reads sweeter hot than cold — which is why a drinking chocolate forgives a lighter sweetener dose and an iced coffee needs a touch more. None of these is sugar. All of them are sweetness, or the absence of bitterness, conjured from the composition. The sweet axis is never ordered alone on our platform; every balance ships inside a small composition — a floor of kokumi, a managed acid, a whisper of salt, the right warmth — tuned so the

sweetness lands full, the bitterness lands quiet, and the whole thing feels like a food rather than a solved equation. A balance is built from its whole neighborhood, and we design the neighborhood.

Chapter 16 · The Reformulator's Playbook — *from a sugar brief to two numbers.*

method · the benchmark vector · the compliance line · a worked case

Most of the people who will use this book do not grow plants. They reformulate products — a mocha losing a third of its sugar to hit a new standard, a drinking-chocolate mix going clean-label, a protein bar that has to drop its syrup without anyone noticing the cocoa turn bitter. This chapter is for them: the method, step by step, that turns a vague brief — "take the sugar out, keep it balanced" — into a specified, reproducible formula with *two* numbers on it, a bitterness and a sweetness held in a chosen ratio. It is the whole book made operational, and it runs in six movements.

One — set the target as a vector, not a vibe. Do not start with "less sweet" or "less bitter." Measure the benchmark. Put the original, full-sugar product on the Index and read its *delivered* **IFI-S** and **IFI-B** together, and while you are there read the axes that travel with them — the acidity that suppresses sweetness, the body and kokumi (**IFI-K**), the astringency (**IFI-A**) that structures a chocolate or a tea. Now you have a target vector: "match S = 1,500, hold B near 1,750, rebuild K to 1,400, keep A at 700." Everything downstream serves that vector, and the ratio of S to B *is* the balance you are protecting.

Two — map the sugar's hidden jobs. From Chapter 15: in *this* product, what was the sugar doing besides tasting sweet and suppressing bitterness? A drink leans on body and mouthfeel; a bar on bulk, binding, and browning; a frozen product on freezing-point depression. List the vacancies you are about to open, because each must be filled by something other than the sweetener, and a bitter-sweet food has the extra trap that the sugar was *also* doing a share of the bitter-suppression you now have to replace with engineering rather than sweetness.

Three — quiet the bitterness at the source, or specify a bean that already is. This is the move that distinguishes a real reformulation from a masked one. Rather than pour sweetness over a harsh cocoa or a harsh coffee until the bitterness drowns, start from an attenuated ingredient — a low-bitter cocoa, a clarified coffee, a de-bittered plant protein — so the bitterness is already down to structure before any sweetener is dosed. Every point of bitterness removed at the source is a point you do not have to pay to suppress with sweetness, which is the whole economy of Chapter 4.

Four — choose the sweetener *system,* not a sweetener. No single molecule wins: the steviol glycosides plateau as partial agonists, the synthetics each carry a temporal flaw, the proteins linger. So you blend, for synergy (sweeteners at different receptor sites yield more

combined sweetness than their doses predict, so you use less of each and dilute each off-note) and for time-intensity (a fast, clean Reb-M front, a round mogroside-V middle, composing the *arc* of the sweetness and not merely its peak). Dose to the target delivered IFI-S, and prefer the clean molecule to the masked one — a Reb-M-dominant stevia starts clean, so there is little to hide.

Five — rebuild the floor and use the cross-axis levers. Put back the body with a bulking system (allulose, which also browns) and, in liquids, a trace of hydrocolloid; add a kokumi base so the sweetness reads full and lasts; and use the levers — a whisper of salt to lift sweet and quiet bitter at once, a trim of acid to let the sweetener dose fall, warmth where the product is served hot.

Six — validate and lock. Taste-panel against the benchmark with the panel *screened for taster genotype* (the PAV/AVI divide of Chapter 2 means a bitterness clean to a non-taster formulator can read bitter to a taster consumer), read the finished eight-number vector, and confirm it holds across the batch at **≤7% CV** and across shelf life. A reformulation that tastes balanced once is a prototype; one that reads the same two numbers every batch is a product.

A worked case, the most common one — a full-sugar bottled mocha, reformulated to no added sugar. The benchmark reads sweet and bitter with a sugar-given body:

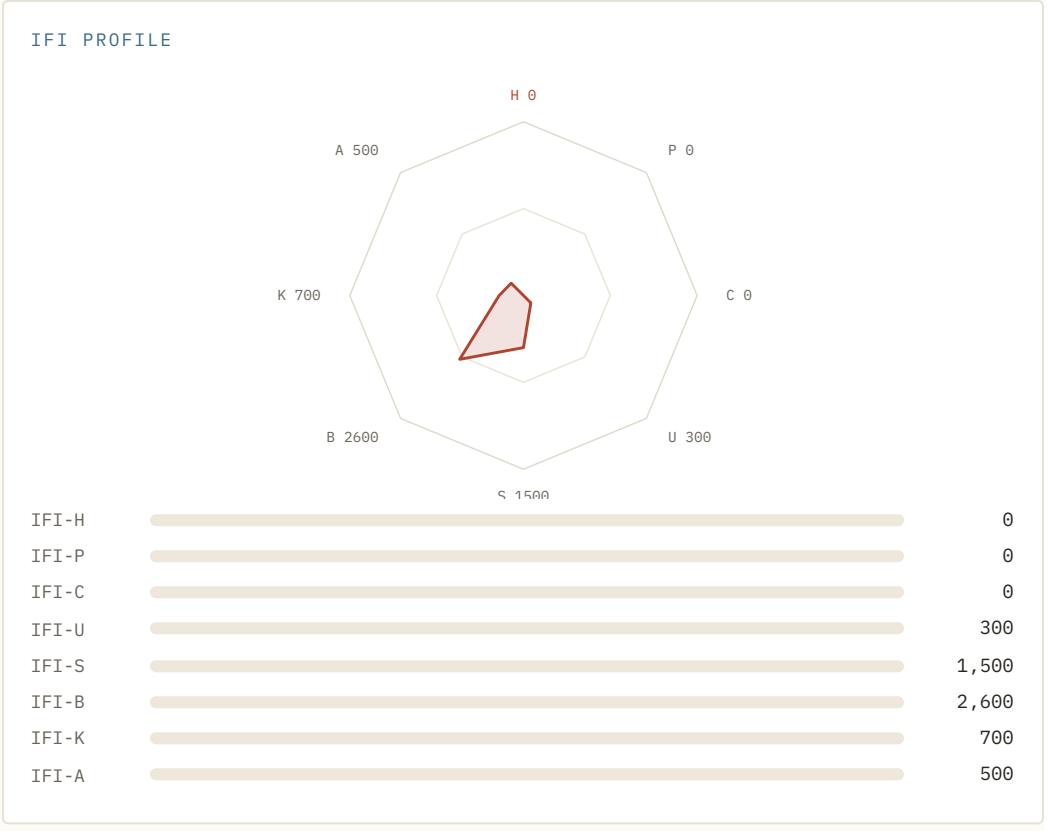

And the reformulation — a clarified-bean coffee and a low-bitter cocoa to bring the bitterness down at the source, a Reb-M stevia and mogroside-V monk fruit blend for the sweetness arc, allulose and a trace of soluble fiber for the lost body, a γ-glutamyl kokumi base for continuity, a pinch of salt doing double duty:

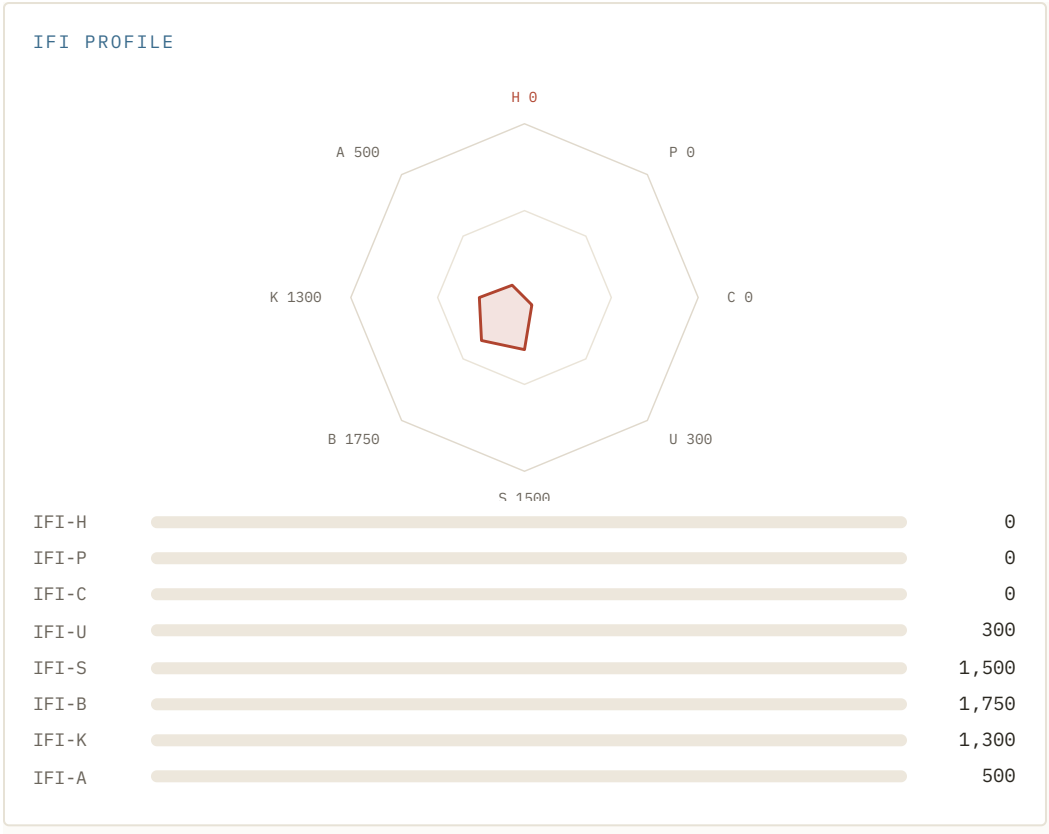

The sweetness lands on the same number, 1,500 — a drinker could not tell them apart on the sweet axis. But look at the real work: the **B falls from 2,600 to 1,750,** not by burying it under sugar but by starting from attenuated ingredients, so the cup reads as a balanced mocha rather than a bitter one sweetened into submission; and the **K rises from 700 to 1,300,** the body the sugar used to give, rebuilt deliberately out of bulk and kokumi, so the zero-sugar version reads full instead of watery. Two scorecards, and the whole method is visible: matched where it must match, lowered where the bitterness was punishing, rebuilt where the sugar was quietly working.

A closing word on what a formulator may and may not say, because it is a matter of both law and ethics and this library holds the line hard. Sweetener use and crop engineering move through real regulatory frameworks, and the honest posture is the plain one — say what the thing is, show the data, make no claim the data does not carry:

SWEETENER	US STATUS	EU
Steviol glycosides (Reb A, Reb M)	GRAS	E960
Monk fruit (mogrosides)	GRAS	flavor use
Allulose	GRAS	limited
Erythritol	GRAS	E968
Thaumatin	GRAS (flavor)	E957
Brazzein	not yet approved (frontier)	not approved

Claims about the *product* are compositional attributes with legal definitions — "no added sugar," "reduced sugar" (at least a quarter less), "sugar-free" (under half a gram a serving) — and those are the only kind this book makes. We design to the **attribute**: the two numbers, the balance, the absence of added sugar, the clean label. 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 — is a *value* that guides what we build, never a claim about what it cures. A coffee

drunk black and glad, a bar that is honestly what it says, a chocolate that tastes of cacao and not of candy: these are facts you can put on a label and pleasures you can put on a tongue. What they do for a body is between that body and the life it is living. We make the balance honest. The rest is not ours to promise.

Chapter 17 · One Fruit, in Balance — *reading Phoenix's bitter and sweet.*

SS-0096 · B 2,200 and S 1,100 inside the chord of eight · the order of delivery

To see the whole argument settled in a single plant, read the one cultivar we built to register on every axis at once. **Phoenix** — SKU **SS-0096** — is the library's full-spectrum flagship, a *Capsicum chinense* engineered to speak all eight axes, the only cultivar that does. It is, by heat, a monster: our model predicts approximately 2.94 million SHU, a computed **IFI-H of 4,200,** the loudest burn our family has ever designed, powered by the double Pun1 variant (S39L+L345G) of Chapter 6. But the heat is only the headline. The reason the fruit is coherent rather than chaotic — a composition rather than an assault — is everything underneath it, and two of those numbers are the subject of this book:

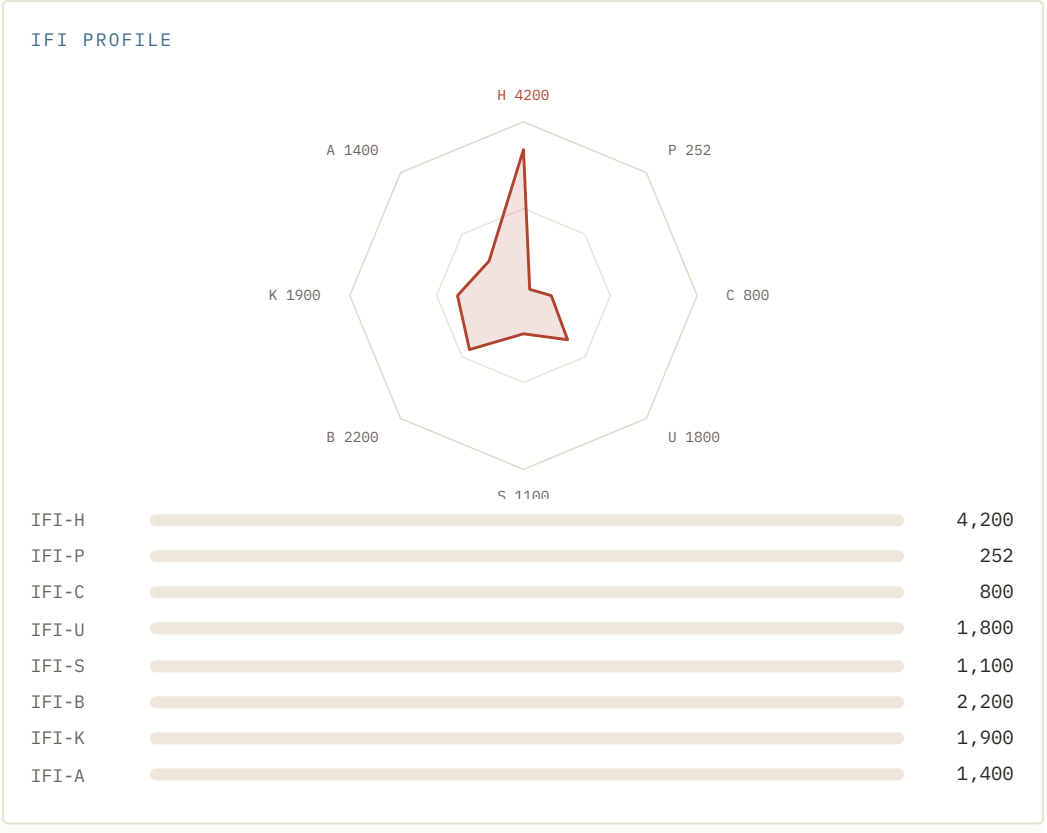

Look at the bitter and the sweet. **IFI-B 2,200** and **IFI-S 1,100** — and notice, first, that the bitter number here runs the *opposite* way from everywhere else in this book. On a coffee or a cabbage, B is a number we drove *down.* On Phoenix it is a number we built *up,* and kept, to a firm and deliberate 2,200. This is the bitter axis used in

reverse, and it is the proof that *Zero-Bitter* is not a crusade against a taste but the ability to put bitterness exactly where a composition needs it. A fruit pushed to the furthest edge of the heat ladder needs a bitter *spine* beneath the fire, or the burn reads as a chemical shriek instead of a flavor; so we specified bitterness up to the level the heat required, no more and no less. And the **S of 1,100** — a real, designed, computed sweetness, above the reference sugar's 1,000 — is the floor under that fire, a sweetness grown in on purpose to round the heat's leading edge so the burn arrives on *something* rather than on nothing. Even here, at the most extreme coordinate in the library, bitter and sweet are composed against each other exactly as they are in a cup of coffee, by exactly the same method and the same model.



Phoenix — the one cultivar engineered to speak all eight axes at once. Its bitterness is built up to 2,200 as a deliberate spine, and its sweetness grown in at 1,100 as a deliberate floor: the same balance as a cup of coffee, held under the fiercest heat we make.

How is a bitter-sweet balance placed inside a chord of eight? The foundation is that the eight receptors are *independent instruments* — TAS1R2·TAS1R3 for sweet, TAS2R38 for bitter, TRPV1 for heat, CaSR for kokumi, each a different protein with its own lock — so in principle each axis can be dialed without touching the others. But independence of the *receptors* is not independence of the *perception,* and this is the subtlety that makes full-spectrum design an art rather than an arithmetic. Recall the shared wiring and the suppression of Chapter 2: Phoenix's sweetness of 1,100 does not sit in isolation; it sits next to a bitterness of 2,200 that would flatten it and a kokumi of 1,900 and an umami of 1,800 that lift and lengthen it. The 1,100 is not the sweetness we installed; it is the sweetness that

survives its neighbors — balanced against the bitter, buoyed by the savory — so that the number on the page matches the sensation on the tongue. Designing it meant modeling not the sweet receptor alone but the whole competitive, cooperative field of eight, which is the subject of its own volume, *The Combinatorics of Flavor* (Volume XXX), and the full Phoenix story is *Full Spectrum* (Volume XXIX).

The last piece is *order,* and it is why a full-spectrum cultivar is more than its edits. Eight axes firing all at once, in a single undifferentiated rush, would read as chaos; 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, the long tail of bitterness and astringency arriving last — turning a chord struck all at once into a phrase played in sequence. It is the same insight as the chili-chocolate bark of Chapter 10 and the ganache rolled in cocoa: *loudness is set by the edits, but order is set by delivery,* and a flavor is both. That a quieted, deliberate bitterness and a grown-in sweetness sit calmly at the center of the hottest, most complex cultivar we have ever built is the proof this chapter exists to give. The two-move design is not a trick for coffee or a fix for a nutrition bar. It is a fundamental grammar of flavor — bitterness tuned to exactly where a composition wants it, sweetness grown in to meet it — and when the flux-balance model composes a fruit as uncompromising as Phoenix, it reaches for the same balance your grandmother reached for with her two sugars, and finds it, at the gene, with nothing added at all.

Chapter 18 · Reference Readouts — *the bitter-sweet shelf, scored.*

every domain, both axes · native → engineered · the recipe coordinates

Every domain and every dish in this book on a page, for citation and comparison. The first table is the plants, each carrying its *native* and *designed* bitterness as IFI-B on the shared 0–5,000 scale, the attenuation between them, and the sweet lift the balance draws on. The engineered IFI-B figures are computational predictions from our models. Read the bitter numbers as loudness, as everywhere in the Index: the *lower* the designed figure, the quieter the plant. Where the move inverts — hops, grape, Phoenix — the note says so.

DOMAIN	ORGANISM · CONSTRUCTS	NATIVE IFI-B	ENGINEERED IFI-B	ATTENUATION	THE SWEET LEVER IT MEETS
Broccoli	Brassica · 380	4,600	1,750	~62%	the vegetable's own uncovered sweetness
Broccoli-Kale	Brassica · 280	4,400	1,760	~60%	uncovered sweetness
Kale	Brassica · 40	4,300	1,720	~60%	uncovered sweetness
Cabbage	Brassica · 40	3,900	1,750	~55%	own sugars + a drop of Reb M
Radish	Brassica · 40	3,700	1,850	~50%	uncovered sweetness
Rapeseed	Brassica · 40	3,500	1,820	~48%	— (oil and leaf)
Orange	Citrus · 120	3,850	2,000	~48%	the fruit's own sugar, uncovered
Coffee	Coffea · 40	3,450	1,900	~45%	the bean's uncovered sweetness; Reb M in a mocha
Cacao*	the flavonoid branch	2,830	1,700	~40%	mogroside V + Reb M, no sucrose
Hops	Humulus · 80	2,580	1,600	~38%†	— (bitterness shaped, not stripped)
Tea	Camellia · 280	2,000	1,400	~30%	theanine uncovered + a monk-fruit whisper
Grape	Vitis · 1,610	—	light touch	<20%‡	— (tannin is the structure)
Stevia	Stevia · 910	(its own)	Reb M clean	—	the sweet key itself; walked past its bitterness
Monk Fruit	Siraitia · 80	(its own)	mogroside V clean	—	the rounding sweet key
Sugarcane	Saccharum · 200	—	—	—	redirected to high-potency, low-load sweetness
Phoenix	SS-0096 · flagship	—	**2,200 (built up)**	kept	S 1,100 grown in as a floor

∗Cacao is method-transfer and roadmap, not one of the forty-two host organisms in the platform; its figures are the profile the method targets. †Hops is shape-tuning of the harsh fraction, not removal. ‡Grape's bitterness is minor and mostly left alone; its tannin is scored on IFI-A, with the full story in *Astringency* (Volume XIII) and *Grape & Wine* (Volume XVIII).

And the dishes of this book, each as a bitter-sweet coordinate — the engineered bitterness it carries, the delivered sweetness that meets it, and the ratio that names the balance:

DISH (CHAPTER)	IFI-B	IFI-S	READS AS
The Black Cup (8)	1,900	300	clean coffee; sweetness uncovered, not added
Clarified Cold Brew (8)	1,900	300	the same balance, carried into the brew
Unsweetened Cortado (8)	1,500	700	milk softens bitter; lactose reads as sweet
Bitter-Sweet Mocha (8)	1,750	1,300	two bitternesses settled by one grown-in sweet
80% Cacao Bar, target (9)	1,700	500	restraint; grip kept, scrape eased
Dark Drinking Chocolate (10)	1,800	1,500	a proud bitterness met, not buried
Bitter-Sweet Ganache (10)	1,750	1,400	bitter coating frames the sweet center
Chili Dark Chocolate Bark (10)	1,750	1,300	sweet, then dark, then a late warmth
No-Sugar Cocoa Bar (11)	1,650	1,400	cocoa and protein quieted together, level
Seeded Protein Bar (11)	1,700	1,400	protein-forward, no added sugar
Blistered Cabbage (12)	1,750	850	the blister reads as caramel, not char
Charred Broccoli (12)	1,750	900	the florets' own sugars, uncovered
Massaged Kale (12)	1,720	600	bitter quieted, mineral astringency kept
Green Tea, long-steep (13)	1,400	1,100	a kept bitterness met, theanine forward
Whole-Pressed Orange (13)	2,000	1,300	pith eased; the fruit's own sugar audible
Reformulated Mocha (16)	1,750	1,500	the worked case; body rebuilt to 1,300 K
Phoenix · SS-0096 (17)	2,200	1,100	the balance held under the fiercest heat

Three things are worth reading off the two tables together. The bitter **band holds** across every plant — nothing below the twenty-percent floor we consider worth claiming, nothing past the sixty-five-percent ceiling we refuse to cross — and the spread inside it is not arbitrary: the brassicas sit high near the ceiling, where the harsh and prized fractions separate most cleanly; the tea and the hop sit low, where the work is to shape a wanted bitterness rather than strip it. The **ratios tell the recipes apart** better than any prose: a black coffee's 1,900-to-300 and a drinking chocolate's 1,800-to-1,500 are the same two axes composed to opposite ends of the balance, the one all structure, the other a sweetness risen to meet a bitterness kept proud. And the single row that runs every number the other way, **Phoenix,** is the whole book stated in reverse — a bitterness built *up* and a sweetness grown in beneath a fire — proof that the two-move grammar is not a special case for the bitter shelf but the fundamental way flavor is composed in this library. These are the numbers to cite. The chapters behind them are how they were earned.

Chapter 19 · Measured, Not Claimed — *two assays, one discipline.*

the TAS2R38 and TAS1R2·TAS1R3 assays · ≤7% CV · screened panels · attributes never claims

Everything in this book reduces, in the end, to two numbers we are willing to stand behind, and this final chapter is about how those numbers are made and what we will and will not say about them. The discipline is the one that runs under the whole library — *the math is exact and shown, because numbers are a form of respect for the reader* — and nowhere does it matter more than on an intersection whose subject touches, however lightly, both the foods people reach for when they want pleasure and the foods they reach for when they want to eat well.

Each axis begins in a dish, and because this book composes two axes it runs two assays. Bitterness is read on a **cell-based functional assay**: a line of cultured cells engineered to express the reference **TAS2R38** receptor together with the signalling machinery needed to report its activation as a flash of calcium — a promiscuous or chimeric G protein coupling the bitter receptor to the calcium pathway, and a fluorescent calcium-sensitive dye loaded into the cells. Apply a bitter compound across a range of concentrations; the cells flash brighter as the dose climbs; read the brightness on a plate reader; fit the dose-response curve to the Hill equation of Chapter 3; and the midpoint is the EC_{50} the whole Index is built on. Run reference and candidate side by side, take the ratio, scale by a thousand, and you have an IFI-B; the attenuation is the distance between the native plant's reading and the engineered plant's, measured the same way on the same instrument. Sweetness is read on the parallel assay against **TAS1R2·TAS1R3**, cells expressing the human sweet heterodimer reporting through the same calcium readout, candidate against the ten-percent sucrose reference, giving an IFI-S. Two receptors, two dishes, one arithmetic — and a bitter-sweet composition is designed against both at once, so the *ratio* that is the real subject of this book is computed from a principled model and not merely guessed.



Assay plates in the rack — a bitterness and a sweetness each begin as a dose-response curve in a dish of cells, read to a published tolerance before either number is ever printed.

A single curve is a draft, not a result. The numbers we publish are held to the library's standard of reproducibility — a batch-to-batch **coefficient of variation of seven percent or less** — because that figure is the difference between a balance you designed and one you got lucky with once. It means the attenuation and the sweetness we design for a cultivar are targets aimed at reproducibility — a specification a formulator can build a product on, designed to hold within that tolerance across seasons, growth conditions, and batches. And because this book's subject is perception, the design framework calls for both numbers to be confirmed twice over: on the cell assays that give the receptor-level EC_{50}s, and against **trained human sensory panels** — real tasters, scoring the plants and dishes under controlled conditions, their perceptions correlated back to the assays so that the numbers in the dish and the experience on the tongue can be shown to move together. The bitter panel carries one extra discipline this book has returned to from the start: it is **screened for taster genotype.** Because the PAV/AVI polymorphism of Chapter 2 means a bitterness clean to a non-taster can read sharp to a taster, an honest bitterness number must be anchored to a defined reference genotype and confirmed across the distribution of real tongues, not asserted from one. The full statistical machinery is its own volume, *Precision & Consistency* (Volume LX); here the point is narrow and firm: an IFI-B and an IFI-S are computed scores, carried to a published specification, not adjectives.

Which brings us to the line, drawn in bright ink because it matters most on exactly this shelf. **We characterize attributes. We do not make health claims.** It would be the easiest thing in the world, on an intersection about quieting the bitterness and lifting the sweetness of coffee, chocolate, greens, and the nutrition bar, to drift into the language of benefit — to say a de-bittered green *delivers* something, a no-sugar bar *does* something for a body, a sugar-free sweetness *provides* an outcome. We do not write those sentences, because they are not sentences we have earned the right to write, and because the moment flavor science starts making medical promises it stops being either good science or honest food. What we will say is bounded and true: that a plant's bitterness is predicted to be reduced by a designed percentage at a named receptor by a characterized edit, that a sweetness is predicted at a computed level without added sugar, that the two have been composed to a reproducible target ratio with the food's other chemistry left in place. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. Those are attributes. Those are all we claim. The compliance posture follows from the same discipline — the library stands on **three patents filed and twenty-one peer-reviewed papers,** released openly on Zenodo under a Creative Commons license, because a standard nobody can inspect is not a standard — and an engineered crop moves through the real regulatory frameworks of the places it would be grown and sold, honestly labeled for what it is and how it was made. The nourishment ethos that hums under this whole book lives entirely in *why we chose this work,* never in a promise on a

package. It is a reason, held quietly. It is never a claim. Keeping those two things distinct is, in the end, the same act as keeping the measurement honest: both are forms of respect, and both are the price of being trusted with something as intimate as what a family puts on its table.

Closing essay · What balance was always for

We began this book with a grandmother's coffee, taken with two sugars and no comment, and we will end there, because that cup is the whole of the intersection in a single gesture. The two sugars were not a failing of taste. They were the only tool a loving hand had ever been given — the one lever, pulled on anything too sharp to face, for ten thousand years of kitchens. There was no quarrel with the gesture. The quarrel, if there was one, was only ever with the fact that the gesture was all anyone had: that to forgive a bitterness you had to drown it, and to sweeten a thing you had to load it with sugar, and that the most nourishing shelf in the market — the dark greens, the strong coffee, the high-cacao bar — was also the one most often left untouched, because the warning wrapped around it was louder than the welcome inside.

The two-move design lets the food stop apologizing. Turn the alarm down at the gene, grow the sweetness in at parts per million, and the balance a cook has always chased arrives carrying none of the load it used to demand — a coffee you drink black and glad, a chocolate that tastes of cacao and not of candy, a vegetable a child finishes because it reads sweet underneath, a bar that is honestly what it says it is. We did not do it by silencing bitterness, which would be foolish and which we refuse; the coffee keeps its structure, the chocolate keeps its grip, the tea keeps its backbone, and Phoenix keeps a bitterness built firmly *up.* We did it by learning, receptor by receptor and pathway by pathway, exactly how much of each warning was ever doing the food any good, and how little sweetness it truly needed once the warning came down.



Dusk in the greenhouse — the plants that carried this intersection, grown toward the fullest flavor and the fullest nourishment they can hold, and the balance a family was always trying to set on a table.

There is a reason this was the book we most wanted to write, and it is not a scientific one. For as long as we have fed the people we love, we have watched the world answer every bitterness with sugar — in the coffee, the cocoa, the cereal, the bar that called itself healthy — until a great deal of what we eat is sweetened less because it should be than because something underneath it was too sharp to meet. The sugar was never really for the pleasure. It was an

apology the food was making for its own bitterness. We do not claim that a food which stops apologizing cures or detoxes or fixes a single thing; those are not our sentences, and they are not true. We claim something plainer and, to us, warmer: that a sweetness you can give the people you love without a spoon of sugar behind it, met to a bitterness you quieted only as much as it needed, is a kinder thing to set on a table — and that feeding the body gently has always been, for us, a form of reverence and not a diet.

This was only ever a book about balance — the oldest thing a cook knows how to do, finally possible to do in the plant instead of in the bowl. Bitterness is not the enemy; it is a voice, and a voice worth keeping. Sweetness is not the cure; it is an answer, and an answer that need not arrive as sugar. Tune them both with care, each at its own receptor, each by the smallest move that does the work, and the two tastes evolution built to argue settle into the harmony they were always inclined to find. You are left, at the end of it, with something bitter made beautiful — a cup, a square, a leaf, a fruit at the edge of the heat ladder — and not one grain of sugar you did not mean to give. That was always the whole point of learning to settle anything.