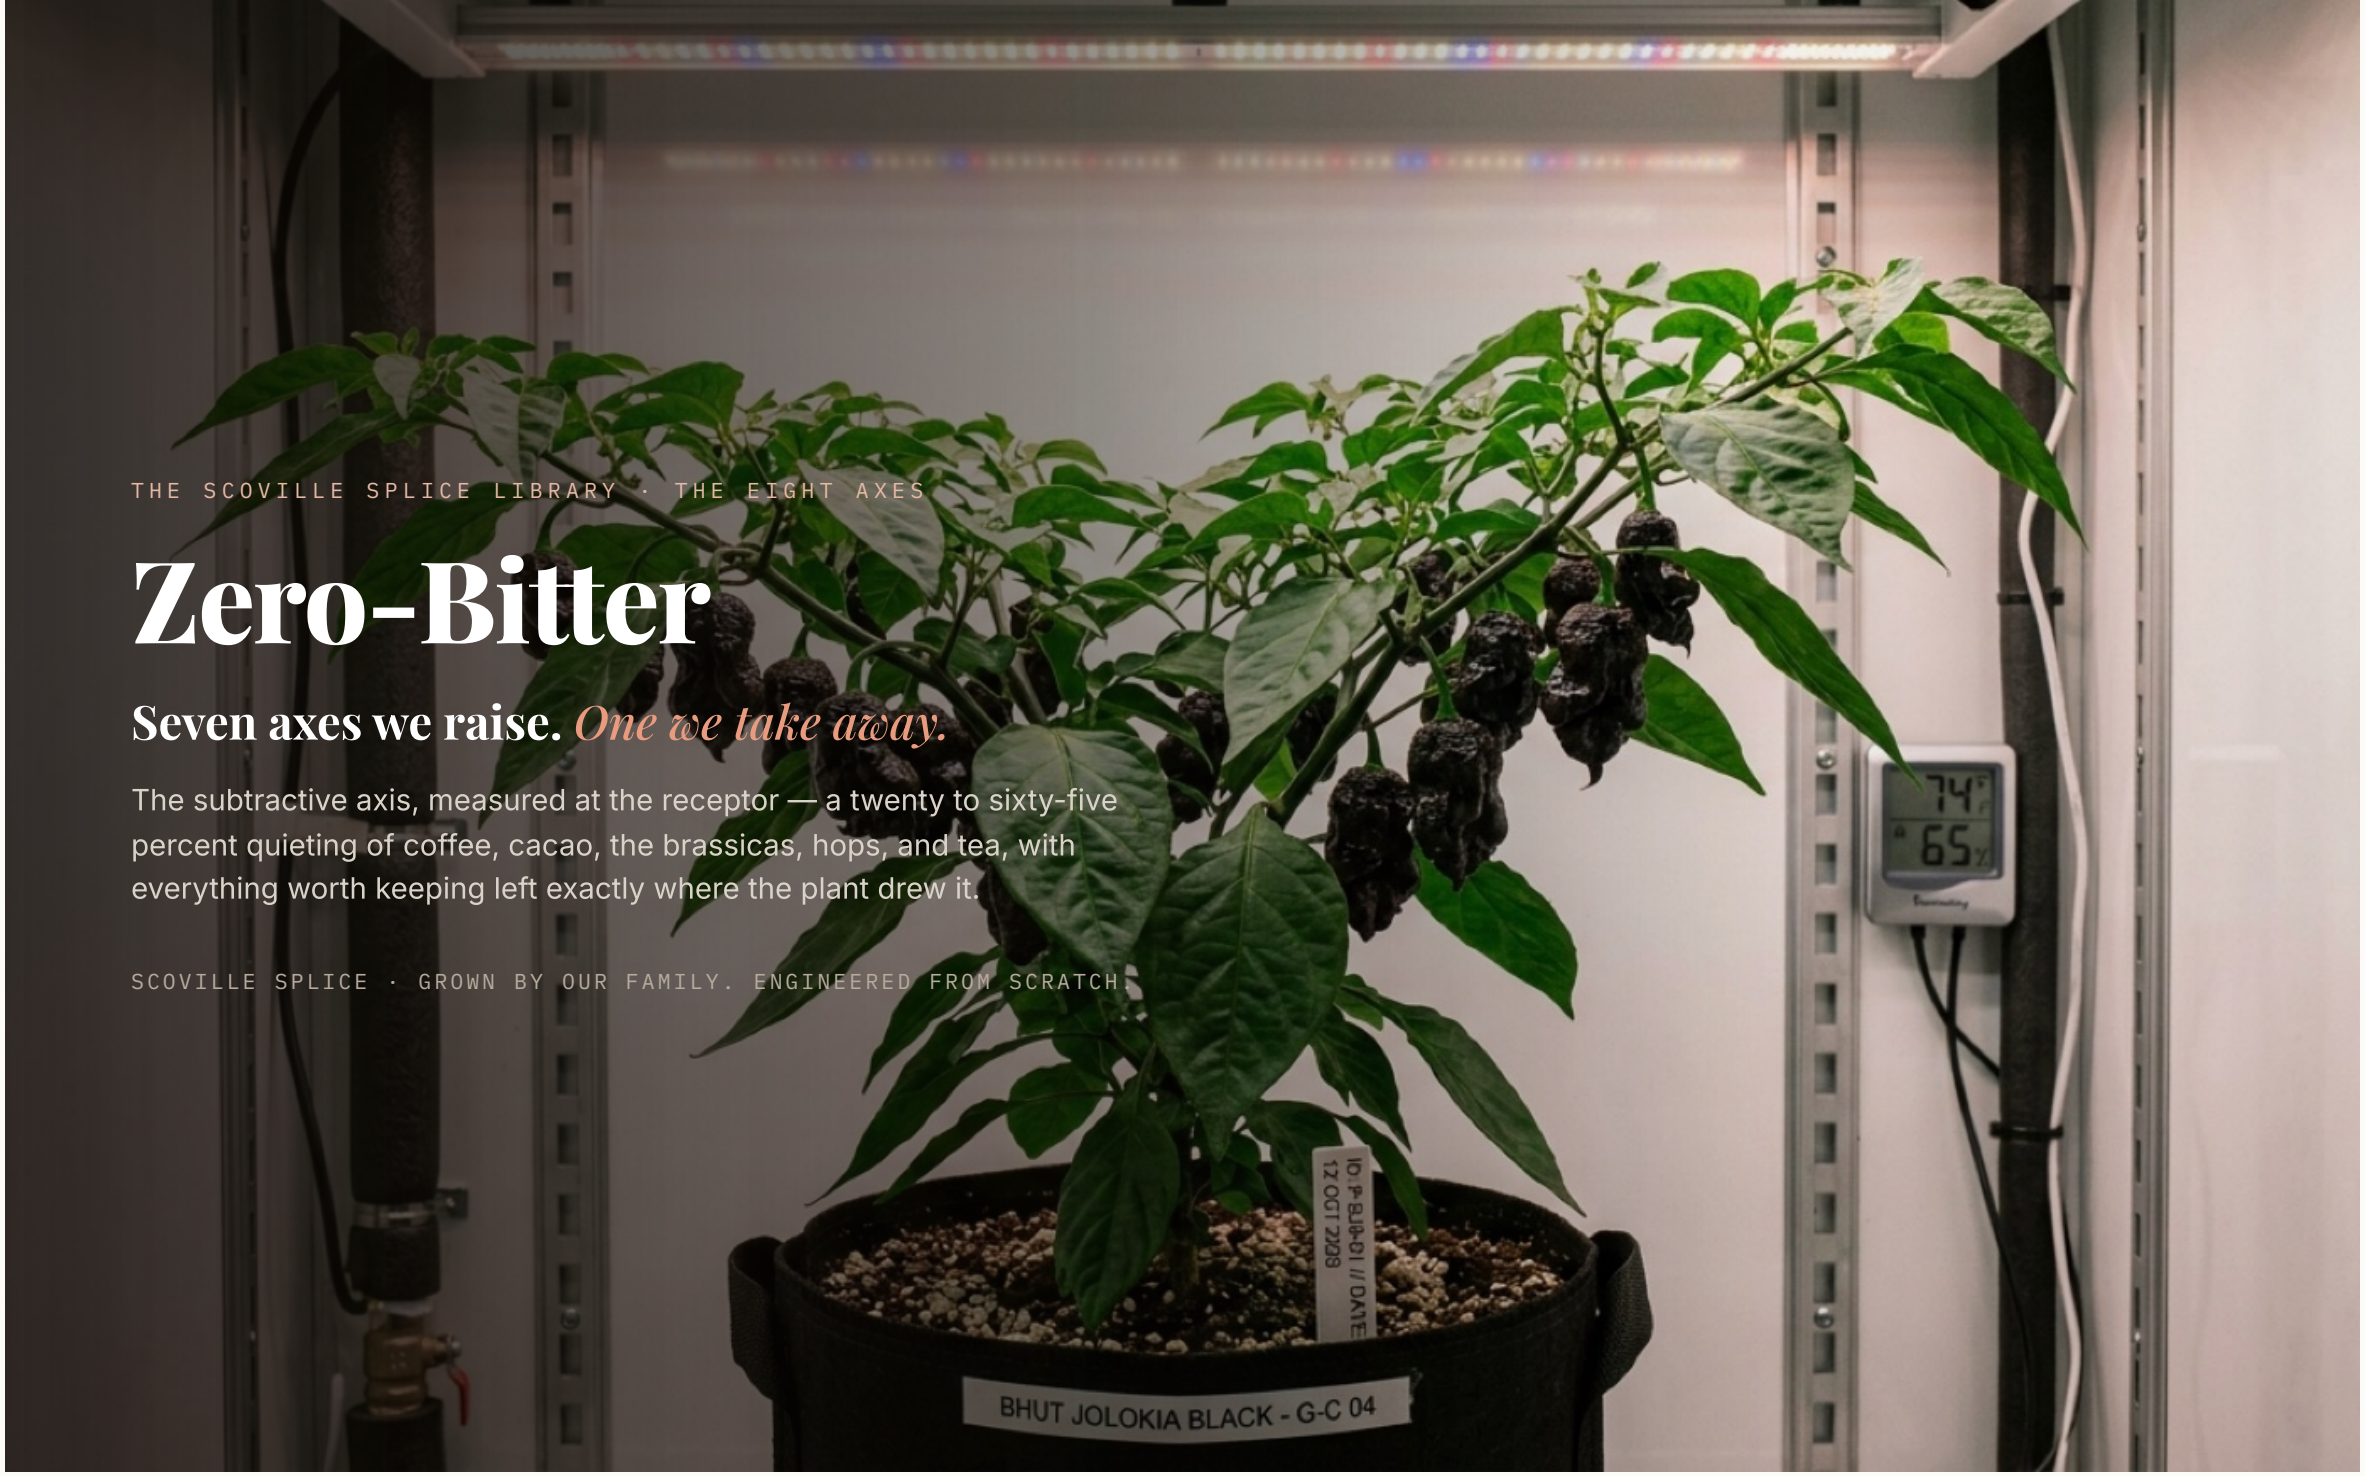

THE SCOVILLE SPLICE LIBRARY · THE EIGHT AXES

Zero-Bitter

Seven axes we raise. *One we take away.*

The subtractive axis, measured at the receptor — a twenty to sixty-five percent quieting of coffee, cacao, the brassicas, hops, and tea, with everything worth keeping left exactly where the plant drew it.

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

Contents

01 Opening essay · The one taste we take away

02 Chapter 1 · The Oldest Alarm — bitterness as biology.

03 Chapter 2 · TAS2R38, the Anchor — one receptor to measure a thousand bitters.

04 Chapter 3 · The Choir — the other twenty-four sentinels.

05 Chapter 4 · From Tongue to Thalamus — the cascade behind a wince.

06 Chapter 5 · The Axis We Read Backward — IFI-B and the arithmetic of attenuation.

07 Chapter 6 · The Anatomy of a Bitterness — the molecules and their roads.

08 Chapter 7 · The Design Loop — scoring an edit before we cut.

09 Chapter 8 · The Glucosinolate Engine — brassica bitterness from gene to goitrin.

10 Chapter 9 · Sparing the Good Bitter — down-tuning progoitrin, keeping glucoraphanin.

11 Chapter 10 · A Vegetable That Gets Eaten — the brassica table.

12 Chapter 11 · The Two Bitternesses of Coffee — lactones, phenylindanes, and the roast.

13 Chapter 12 · Lowering the Lactones — engineering the bean, not the roast.

14 Chapter 13 · Cold, Clear, Unsweetened — the cup, served two ways.

15 Chapter 14 · The Bitterness We Keep — hops, IBU, and the shape of a bite.

16 Chapter 15 · Theobromine and the Flavan-3-ols — cacao, where bitter meets astringency.

17 Chapter 16 · Tea, Grape, Citrus — the polyphenol family, and knowing when not to cut.

18 Chapter 17 · Bitter × Sweet — subtract and uncover.

19 Chapter 18 · The Taster at the Table — polymorphism, perception, and one number against one experience.

20 Chapter 19 · Reference Readouts — the bitter shelf, scored.

21 Chapter 20 · Measured, Not Claimed — precision, CV, and the line we do not cross.

22 Closing essay · What's left when the bitterness goes

Opening essay · The one taste we take away

Every other book on this shelf is a book about more. Heat arrives and climbs; sweetness lifts; umami deepens the floor of a broth until the spoon feels heavier in the hand; a cool note blooms and spreads. Seven of the eight axes we measure are things we have learned to *add* to a plant, to grow louder and more generous at the source. This is the book of the eighth, and the eighth is different in kind. On this axis the whole craft is restraint. Here, for once, we are not turning a flavor up. We are taking something away — carefully, selectively, by the smallest possible cut — and leaving everything else exactly where it stood.

The something is bitterness, and to understand why removing it is the hardest design in the whole library you have to begin with a child's face. A toddler meets raw broccoli and recoils on instinct; meets the first sip of a parent's coffee and makes the universal face, the one no culture has to teach; meets a strip of grapefruit pith and pushes the whole plate an inch away. That recoil is older than language, older than cooking, older than our species. For as long as there have been tongues, bitterness has been the one taste that means *stop* — a chemical fence a plant draws around itself, the word it uses to say *do not eat this.* We have spent ten thousand years of kitchens trying to talk ourselves past that fence, with sugar and cream and long roasting and the patient blunting of salt. This axis does not talk past the fence. It lowers it, at the one place it is built.



The greenhouse at first light — where a question about a vegetable nobody at the table wanted to finish became a line of work.

The place it is built is a receptor called **TAS2R38** — the most-studied of the twenty-five bitter sentinels on the human tongue, evolution's smoke alarm, tuned to the sulphurous edge of a glucosinolate, to the scorch at the bottom of a too-dark roast, to the back-of-throat grip of a young beer. On the seven additive axes we measure loudness and design it upward. On this one we measure the same loudness and design it *down*: how far we can push a bitter molecule's threshold the wrong way, so the receptor needs more and more of it before it will answer at all. We call the reading **IFI-B**, scored on the same 0–5,000 numeric range as the other axes — each anchored to its own reference compound by convention — and we hold ourselves to a

single, plain, published promise — a bitterness lowered **twenty to sixty-five percent**, read at the receptor, with not one note of the thing itself lost on the way down.

The difficulty, the thing that makes this the axis we gave the most patience, is that the warning and the worth are so often the same molecule, or two molecules built on the same road. The glucosinolates that make raw kale harsh are kin to the compounds a careful eater is quietly glad to be getting. The polyphenols that grip in a cup of tea are the ones that give it its backbone and its length. The theobromine that makes dark chocolate lift a little is the chocolate telling the truth about itself. A plant never separated the message from the medicine, because it never needed us to *enjoy* the meal — only to survive it. So the crude move, the one we almost never make, is to knock out the gene and let the bitterness collapse along with everything the plant paid for to build it. The careful move — the only one worth making — is to tell the warning from the worth, and turn down only the first.

This is also why the axis is named for a goal and not for its arithmetic. We never design a flavor all the way to zero. A coffee with no bitterness is not a better coffee; it is warm brown water. A dark chocolate with no grip is candy. A green with no edge is a green that has forgotten what it is. Bitterness, in the right measure, is spine — and a flavor without a spine cannot stand up. So *Zero-Bitter* is a direction, not a destination: the honest work is to find, for each plant, the exact amount of quiet that lets the rest of it finally be heard, and to stop there.

The warning older than language

To lower a fence you first have to understand how it stands. TAS2R38 is a G-protein-coupled receptor — a single protein that threads the membrane of the tongue's bitter-sensing cells seven times, forming a pocket on the outer face where a molecule can dock, and a handle on the inner face where, the moment something docks, it grips a waiting G-protein. There are twenty-five of these bitter receptors in the human genome, a small, ancient, promiscuous family. Between them they recognize thousands of chemically unrelated compounds, because the one job evolution asked of them was coverage, not precision: a plant toxin the lineage has never met before should still trip *something*. TAS2R38 is the archetype, the one the textbooks draw, because it is tuned to a chemical signature a plant wears whenever it is defending itself in earnest — the thiourea group, a nitrogen double-bonded through carbon to sulphur, **N–C=S**. Phenylthiocarbamide carries it. Propylthiouracil carries it. And goitrin, the compound a cut Brassica makes within seconds of being bitten, carries it too.

Watch what happens in the instant the molecule lands. Docked in the orthosteric pocket — a cleft braced by the third, fifth, sixth and seventh membrane-spanning helices — the agonist nudges the receptor into its active shape. On the inside, the receptor releases α-gustducin, the taste cell's own dedicated G-protein, and the freed βγ pair switches on phospholipase C-β2. That enzyme cleaves a membrane lipid into two messengers; one of them, IP_3, unlocks a calcium store inside the cell. Calcium floods the cytoplasm and throws open TRPM5, a channel that only answers to calcium, and the cell depolarizes. The charge opens CALHM1, the cell spits a pulse of ATP

across the gap, and the gustatory fibers of the seventh and ninth cranial nerves carry the single syllable upward: *stop.* From dock to decision is a few hundred milliseconds. The child's face is faster than thought because it is built below thought.

Here is the complication we have to respect, and it is written into the receptor's own DNA. TAS2R38 comes in two common forms, separated by three single-letter changes in its sequence — positions 49, 262 and 296. One spelling, **PAV**, folds into an exquisitely sensitive receptor; the other, **AVI**, into a dull one that barely answers. Inherit two copies of PAV and a trace of propylthiouracil is unbearable; inherit two AVI and you can chew the paper strip and shrug. Between a quarter and a half of people, depending on where their grandparents were born, carry the quiet form. This is the single most important fact about designing on the bitter axis, and we will come back to it, because it means the fence is not the same height for everyone standing at it — and a design read only at the receptor is not yet a design read at the table.

The arithmetic of subtraction

The whole Index rests on one idea: a flavor's score is not how much of a compound a food contains, but how *loudly* that compound speaks to its receptor. Loudness is read from the dose-response curve. Add a bitter agonist to a dish of cells carrying TAS2R38 and the response climbs along a sigmoid described, to a very good approximation, by the Hill equation:

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

where [A] is the concentration of the agonist, n the steepness of the climb, and **EC_{50}** the concentration that drives the receptor to exactly half its maximum. EC_{50} is the number that matters. A potent stimulus reaches half-maximum at a whisper of a concentration — a small EC_{50}; a weak one needs to shout. Invert it against a fixed reference and scale by a thousand, and loudness becomes a single portable score:

$$IFI = 1000 \cdot EC_{50}(reference) / EC_{50}(candidate)$$

The reference is anchored so that a standard bitter stimulus reads exactly **1,000** by convention. A food twice as loud as that reference — reaching half-activation at half the concentration — reads 2,000. This is the same machinery heat is scored on, pointed at a different receptor; what changes on the bitter axis is only the direction we push.

Because near threshold the Hill curve is close to linear — when [A] is well below EC_{50}, the response simplifies to $R/R_{max} \approx [A]/EC_{50}$ — perceived bitter loudness tracks a single ratio we will call ρ, the delivered agonist over the half-activation point:

$$\rho = [A] / EC_{50}, \text{ and } IFI\text{-}B \propto \rho$$

Now the design has a clean target. To make a food less bitter we must lower ρ, and there are exactly two levers. We can reduce [A] — grow the plant to carry less of the bitter ligand — or we can raise the

effective EC_{50}, so the receptor needs more of the molecule before it answers. In loudness they are the same move: halving the ligand and doubling the half-activation point both halve ρ, and both halve IFI-B. We define the attenuation we publish as the fractional drop in the score:

> $\alpha = 1 - \text{IFI-B(engineered)} / \text{IFI-B(wild-type)} = 1 - EC_{50}(\text{wild}) / EC_{50}(\text{engineered})$

and from that one line the entire promise of the axis — *twenty to sixty-five percent* — resolves into the exact molecular work required. A 20% attenuation means the engineered food retains 80% of its loudness, which means its effective EC_{50} must rise by a factor of 1 ÷ 0.80 = 1.25. A 65% attenuation, the far edge of what we will claim, means retaining 35%, which demands an EC_{50} raised by 1 ÷ 0.35 ≈ 2.86 — nearly tripling the concentration the receptor requires before it will answer at all.

ATTENUATION A	IFI-B RETAINED	EC_{50} SHIFT (ENG ÷ WILD)	LIGAND RATIO P(ENG) ÷ P(WILD)
20% (floor)	80%	1.25×	0.80
35%	65%	1.54×	0.65
45%	55%	1.82×	0.55
55%	45%	2.22×	0.45
65% (ceiling)	35%	2.86×	0.35

A worked case makes it concrete. Anchor the reference stimulus at $EC_{50}(\text{ref})$ = 1.0 µM, so it reads IFI-B 1,000 by definition. A wild cup of coffee whose bitter fraction reaches half-activation at 0.50 µM is twice as loud: IFI-B = 1000 × 1.0 ÷ 0.50 = 2,000. Ask for a 45% attenuation and the target score is 0.55 × 2,000 = 1,100; the EC_{50} the cup must now present is 1000 × 1.0 ÷ 1,100 = 0.909 µM — a 1.82-fold rise, exactly the fourth row of the table, reached by carrying the bitter ligand down to 55% of its native load while every aromatic and acid in the cup stays put. The arithmetic is not decoration. It is the specification we hand the plant.

Telling the warning from the worth

The reason the crude move fails has a number too. Take a Brassica — a broccoli, a kale, a cabbage. Its bitterness is not born in the leaf; it is born in the bite. Intact, the plant holds its glucosinolates in one compartment and the enzyme myrosinase in another, a loaded trap. Crush the tissue and the two meet, and within seconds the glucosinolates are hydrolyzed into the pungent, bitter breakdown products the receptor reads. Two of those glucosinolates matter most to us, and they sit on the same branch of the same pathway. **Progoitrin** breaks down to goitrin — the thiourea compound, the direct TAS2R38 agonist, the warning. **Glucoraphanin** breaks down to sulforaphane — the molecule the nutrition literature has spent thirty years interested in, the worth the plant paid to build.

A gene knockout that silences the trap — delete the myrosinase, or an upstream enzyme the whole branch depends on — collapses both at once. Call the progoitrin pool P and the glucoraphanin pool R, and measure a design by its selectivity:

```
σ = (ΔP / P) / (ΔR / R)
```

The knockout drives ΔP/P and ΔR/R both to 1 — everything falls — so σ = 1. That is not a flavor design; it is an amputation, and it is the exact mechanism by which a vegetable becomes warm green water, nutritionally hollowed in the same stroke that made it mild. The design worth making pushes σ as far above 1 as the pathway's own junctions allow: drop progoitrin by 60% while holding glucoraphanin to within 5% of where it stood, and σ = 0.60 ÷ 0.05 = 12. The warning falls by more than half; the worth barely moves. That ratio — twelve to one, not one to one — is the difference between the two philosophies of this axis stated in a single figure.



Broccoli in cross-section — the glucosinolate reserve and the enzyme that reads it live in separate compartments until the bite.

Finding the junction that moves P without touching R is a search problem, and it is the same search we run on every axis. A protein language model scores candidate edits across the enzymes that partition flux between the two glucosinolate branches — the chain-elongation and oxidation steps that decide, molecule by molecule, whether a precursor becomes progoitrin or glucoraphanin — ranking each substitution by its predicted effect on catalytic throughput and fold stability. A genome-scale flux-balance model then confirms that the branch point, and not some upstream bottleneck, is the true lever. It is the method that, pointed at the capsaicin acyltransferase in our heat work, wrote two substitutions no breeder had ever combined — **S39L+L345G** — into a single enzyme. Here the targets are different genes on a different pathway, but the logic is identical: compute the smallest edit that re-weights a branch, then write it in with CRISPR and measure what the receptor says. Of the forty-two host organisms live on the platform today, thirteen carry the bitter axis, and every one of them was designed at a junction, never with a hammer.

Coffee, kept whole

Coffee is the proof the method's patience is worth it, because coffee's bitterness is not one molecule but a committee, and most of what it guards is the reason anyone drinks it. Only a small share of a cup's bitterness — on the order of a tenth to a sixth — is caffeine, which answers to a handful of the bitter receptors at once. The bulk is built in the roast: chlorogenic acids, mild in the green bean, degrade under heat into chlorogenic acid lactones that carry the clean bitterness of a medium roast, and then, as the roast darkens, into phenylindanes, the

compounds responsible for the lingering, back-of-throat harshness of a French roast. Push too far and trigonelline breaks down as well, shedding its own bitter notes into the cup. This is a *family* response — several of the twenty-five receptors answering at once, with TAS2R38 the archetype of the architecture they share — which is precisely why a single knockout was never going to serve it. Silence one ligand and the committee still has a quorum.



Coffee in cross-section — the bitterness is built in the roast; the perfume is older than that, and we keep it.

So we design the cup the way the arithmetic demands: lower the bitter fraction's effective ρ by the targeted amount and leave the eight hundred-odd aromatic volatiles — the Maillard sugars, the furans and pyrazines, the fruit and the chocolate a good origin carries — entirely untouched, because they never spoke to a bitter receptor to begin with. The worked target from the arithmetic above lands the cup at IFI-B 1,100, a 45% attenuation. Shown as a profile, the engineered cup reads as almost all perfume and almost no ash:

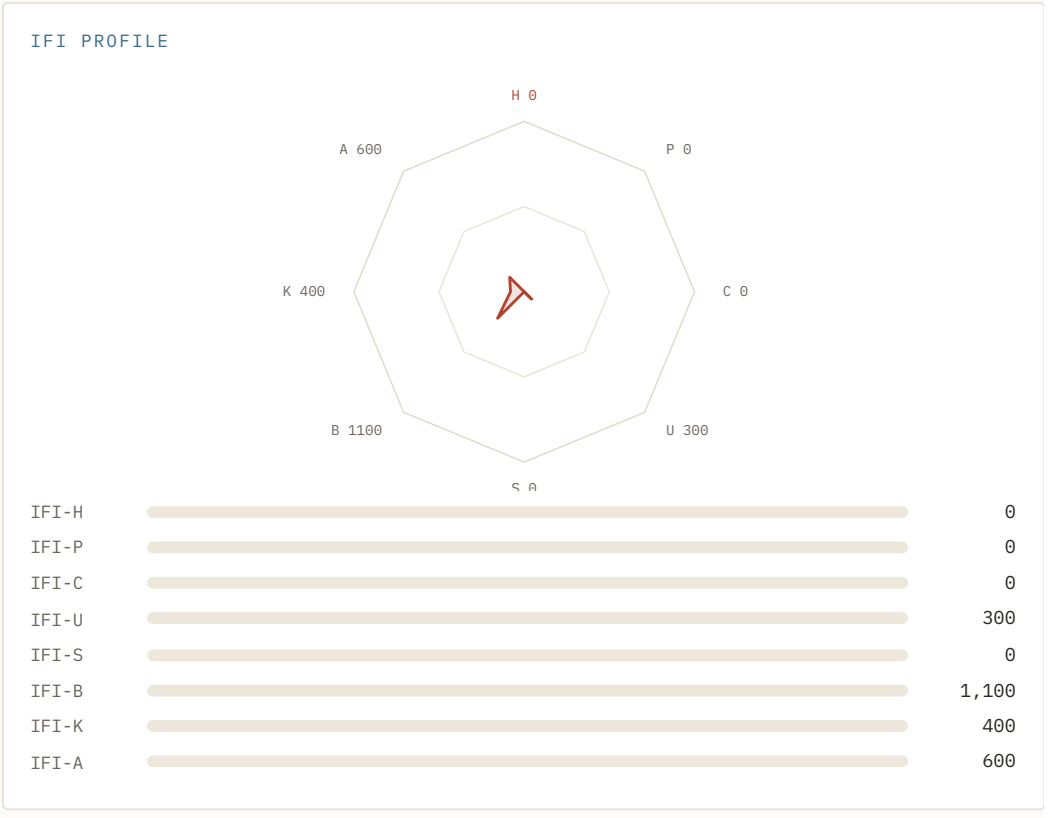

Illustrative target envelope for an engineered medium roast — the bitter axis is the only one we move; the others are the cup's own untouched character, shown to scale, not a batch record. What survives is the whole point: the body, the faint sweetness a dark roast usually buries, the acidity roasters chase and heat destroys. The cup

sheds its ash and keeps its perfume. It is the clearest answer we have to the question under the whole axis — whether you can make a thing gentler without making it less.

A ladder of bitter plants

Thirteen plants carry this axis today, and reading them in order is reading the many chemistries of a single warning. The number beside each is the count of constructs we have built into it — a measure of how much of the design space that organism has already given up.

HOST	CONSTRUCTS	BITTER CHEMISTRY WE LOWER	AXES	DIFFICULTY
Grape	1,610	tannins / polyphenols	B / A / M	Hard
Broccoli	380	glucosinolate breakdown (goitrin)	B / M	Moderate
Tea	280	catechins (EGCG), caffeine	U / B / A	Hard
Broccoli-Kale	280	glucosinolate breakdown	B	Hard
Orange	120	limonin, naringin	B / C	Hard
Potato	80	glycoalkaloids	B	Easy
Hops	80	iso-α-acids (isohumulone)	B	Moderate
Coffee	40	chlorogenic acid lactones, phenylindanes	B	Hard
Cabbage	40	glucosinolate breakdown	B	Moderate
Kale	40	glucosinolate breakdown	B	Moderate
Radish	40	glucosinolate breakdown	B	Hard
Rapeseed	40	glucosinolate breakdown	B	Hard

Each is its own sentence. In hops, the bitterness is the iso-alpha-acids a brewer boils out of the cone on purpose — isohumulone and its kin, among the most intense bitter ligands in any food — and the design is not to remove them but to move them, to let a brewer place the bitterness exactly rather than inherit whatever the boil gives. In cacao, bitterness is theobromine and caffeine, the two methylxanthines, laced through the astringent grip of the procyanidins; the methylxanthine is the chocolate telling the truth about itself, and the work is to let a fine chocolate keep its grip and its lift while shedding the flat, scorched bitterness of over-roasted beans.



Cacao — theobromine is the bean telling the truth about itself; we keep the truth and lower the scorch.

Tea is the hardest, because its catechins — the gallated flavan-3-ols, EGCG chief among them — are bitter and astringent and the whole backbone of the cup at once, speaking to a bitter receptor and to the

salivary proteins of astringency in the same sip; lower them carelessly and a green tea goes slack. Grape, with the most constructs of any bitter host at 1,610, is really an astringency problem wearing a bitter coat, the tannins that give a wine its length reading faintly bitter on the way to their real work in the mouth. Twelve chemistries, one receptor family, one rule: find the smallest cut that frees the rest.

Measuring the cut

A promise of *twenty to sixty-five percent* is only as honest as the instrument behind it, so it is worth saying plainly how the cut is read. The receptor assay is a cell line engineered to carry TAS2R38 and a reporter that lights when the cell's calcium rises; dose it with a ligand across a concentration series, fit the four-parameter sigmoid, and read EC_{50} off the curve. Run the wild extract and the engineered extract side by side and the ratio of their EC_{50}s *is* the attenuation, before a single person has tasted anything. The cell is faster, cheaper and more reproducible than a tongue, and it is also only a cell: it reports the receptor's answer, not the experience, which is why every receptor number is cross-referenced against a trained human panel before it leaves the building.

Reproducibility is where this axis earns trust, because a subtraction you cannot repeat is not a subtraction. We hold batch-to-batch variation to a coefficient of variation of **≤7%** — the standard deviation of a measurement is no more than seven percent of its mean. Because IFI-B runs linearly with the delivered ligand, that 7% carries straight through to the score: a cup read at IFI-B 1,100 sits inside a one-sigma band of roughly ±77. The attenuation claim is a *difference* of two such readings, and differences compound their errors. Propagate them — the standard deviation of α is the retained ratio times the root-sum-square of the two coefficients of variation — and a 45% attenuation, built from two 7% measurements, carries an uncertainty of about $0.55 \times \sqrt{(0.07^2 + 0.07^2)} \approx 0.054$, or ±5.4 percentage points. Forty-five percent, plus or minus five. That is why the published floor is 20% and not lower: at the floor, the same propagation gives roughly ±8 points, so a 20% cut sits about two and a half standard deviations clear of zero — detectable, honest, but close enough to the noise that we will not make a smaller claim. The band is not marketing. It is the window in which the receptor, the assay and the arithmetic all agree.

What we will not do

The ceiling is a different kind of limit, and it is the one we feel most strongly about. Past roughly 65% attenuation a food does not get better; it gets less. A coffee pushed to zero is the warm brown water we started with. A dark chocolate stripped of its grip is a candy pretending. A green with no edge has forgotten what it is for. Bitterness in the right measure is structure — the thing the other flavors lean against — and our instrument can keep cutting long after our judgment says to stop. So we stop. The honest number for each plant is not the lowest we can reach but the one that lets the rest of the food finally be heard, and we have walked back more designs for taking too much than for taking too little.

There are limits we did not choose, too, and we would rather name them than paper over them. The receptor's own two spellings mean the same molecular cut is not felt the same by everyone: lower a ligand by 45% at the receptor and a PAV taster, for whom the fence stood highest, feels the relief most, while an AVI taster, who was barely bothered to begin with, may notice little — the design is real, but its felt size travels with the eater's genome. There is a field limit as well: a glucosinolate is a plant's own pesticide, and a design that lowers the wrong one can hand the crop's defense to the first insect that finds it, which is why the selectivity math is an agronomic promise and not only a flavor one. And bitterness is a committee in more foods than coffee — quiet one ligand and sometimes another receptor in the family simply answers in its place, so we track the whole family's response, not one receptor alone, and accept the designs where the committee stays quiet.

It helps to hold all of this against the opposite case — a plant where bitterness is not subtracted but deliberately built. Our full-spectrum flagship, Phoenix (SS-0096), is designed to register on all eight axes at once, at a predicted heat of approximately 2.95 million SHU and IFI-H 4,200, and its bitter axis is set, on purpose, to **2,200**:

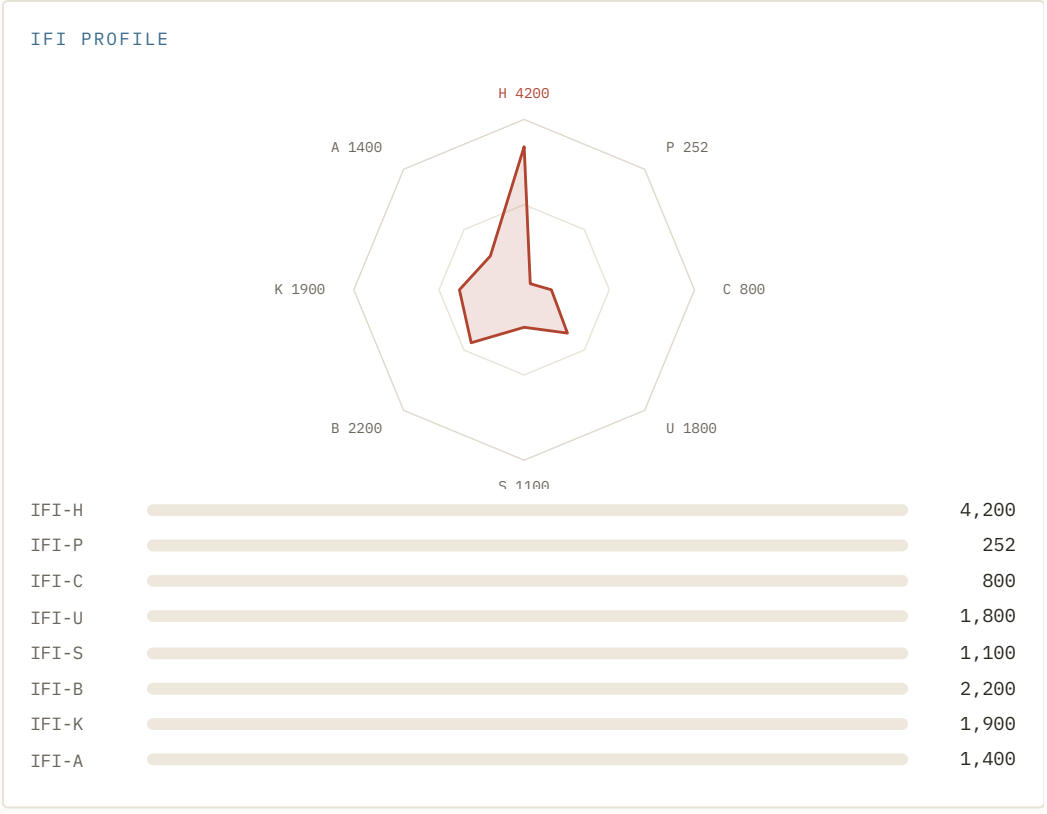

That 2,200 is not a failure to attenuate; it is architecture. In a plant designed to carry the whole spectrum, bitterness is the spine the heat and the sweetness and the umami stand against, and taking it to zero would make the other seven collapse inward. It is worth noting that heat inherited an old folk scale — Scoville, the SHU on every hot-sauce label — while bitterness never got one; there is no Scoville of bitter, no number a shopper could point to, which is exactly why anchoring IFI-B to a receptor's half-activation matters so much here.

The same reading that tells us to *stop cutting* on Phoenix tells us *how far to cut* on a cup of coffee. One axis, two directions, one honest number.

We did not come to this axis to make vegetables sweet, or to flatter anybody's palate. We came because of the plate at the end of the table, the broccoli pushed to its edge, and a conviction that the food was never the problem — only the warning wrapped around it. A vegetable turned down just enough is a vegetable that finally gets eaten. A coffee can keep all its perfume and shed its ash. The most nourishing shelf in any market is, with a stubborn regularity, the bitter one, and a food the body never receives does nobody any good, however much it holds. That last sentence is the quiet argument running under this whole axis, and we are careful to keep it no larger than it is: it is a reason we work, not a claim we make. The receptor, the mathematics of subtraction, and the living plants we have taught — one glucosinolate and one chlorogenic acid at a time — all point the same way: toward a shelf of foods that have, at last, stopped saying *stop*, and finally say yes. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.



PART I OF VI

The Sentinel

TAS2R38 · the bitter family · the transduction cascade from tongue to thalamus

Chapter 1 · The Oldest Alarm — *bitterness as biology.*

IFI-B · TAS2R38 · why the tongue says no

Sweetness and umami are a plant's invitations. Bitterness is its fence. Where the sweet receptors read *here is sugar, here is quick energy* and the umami receptor reads *here is protein, here is nourishment being made available,* the bitter family reads one word, and reads it at the top of its voice: *danger.* It had to. A leaf cannot run. Rooted in place, unable to flee the thing that would eat it, a plant defends itself the only way it can — in chemistry. Alkaloids, glucosinolates, cyanogenic glycosides, saponins, the whole slow-brewed arsenal a plant grows to keep from being stripped to the stem. And the animals that lived long enough to leave descendants were, over and over, the ones who put a leaf in their mouths, tasted that arsenal, and stopped. The wince at raw broccoli is not a defect in us. It is ten thousand generations of ancestors who survived because they spat something out.

What makes bitterness strange among the tastes — and what shapes everything about how we work on it — is how it is built. Sweet is read by essentially one receptor pair. Umami, by two. Sour and salty are read more or less directly, by ion channels sensing protons and sodium. Bitterness alone is read by a *family*: twenty-five functional bitter receptors in the human mouth, the TAS2Rs, a whole committee of sentinels where the other tastes have a single guard. The reason is arithmetic. There is one kind of sugar signal worth detecting, and a handful of savory ones, but the number of bitter poisons a plant or a mold might invent is effectively unbounded — thousands of structurally unrelated molecules, many of which no ancestor ever met. Evolution could not pre-learn them all. So it did the thing it almost never does and built a broadly tuned net: some TAS2Rs narrowly specific to one or two compounds, others promiscuous enough to answer to dozens, the set of them together casting a mesh wide enough to catch a toxin on first contact, sight unseen. Bitterness is not a taste so much as a threat-detection system that happens to open onto the tongue.



A broccoli floret in cross-section — the sulphur chemistry of a plant's defense is laid down in exactly the tissue we are asked to eat.

That breadth is a gift to a species trying to survive an unknown world and a problem for a species trying to enjoy a known one. A broad net catches the poison and the vegetable in the same mesh, because to a sentinel tuned for *unfamiliar sulphurous molecule* a glucosinolate is a glucosinolate whether it rode in on a toxic weed or on a head of kale a farmer spent a season raising. The receptor cannot read intent. It reads structure, and it reports structure as alarm. This is the central fact the whole axis is built around: the bitterness we most want gone is not a mistake the plant is making. It is the plant performing, flawlessly, a function that kept our lineage alive — and performing it on food we have no intention of being poisoned by.

The aversion is wired deep, and early. A newborn, minutes old, will grimace and turn away from a bitter solution placed on the tongue — the facial response is present before any food has ever been eaten, before learning could have shaped it, one of the few genuinely innate taste behaviors we carry. Sweetness, by contrast, draws a newborn in. The asymmetry is the point: a mistake about sweetness costs an animal a little wasted effort; a mistake about bitterness could cost it everything, once. Evolution tuned the alarm to be loud, fast, and biased toward false positives, because in a world of unknown leaves a false *no* is survivable and a false *yes* is not. We inherited the setting. It is why bitterness, alone among the tastes, needs no acquired dislike to be disliked, and why it takes years of coffee and beer and dark chocolate to teach an adult palate to override, in a few narrow cases, what the body still insists is a warning.

And here is the cost of all that ancient caution, the one that brought us to this axis. The compounds a plant loads into its leaves and seeds to defend them — the glucosinolates of the brassicas, the polyphenols of tea and cacao and the grape, the alkaloids scattered through the plant kingdom — are, with a frequency that cannot be coincidence, the very compounds an eater is otherwise glad to be getting. Evolution guarded the plant's chemistry with a single fence and did not trouble to sort the harmful from the merely defended, because to the plant there was no difference worth drawing. The result, carried forward into a supermarket, is that the most densely provisioned shelf — the deep greens, the high-cacao bar, the strong unsweetened tea, the whole-pressed bitter fruit — is also the shelf most often set down untouched. The alarm that once kept us from poison now, just as reliably, keeps us from the garden. Our work on this axis begins precisely there: not in silencing the alarm, which would be foolish and which we will not do, but in teaching a few well-chosen plants to stop setting it off for no reason.

A committee of sentinels — the shape of the net

To work on bitterness honestly, we have to look at the committee itself, and it rewards the looking. The human genome carries roughly twenty-five functional *TAS2R* genes — plus a scatter of pseudogenes, receptors switched off along the way — gathered in tight clusters on chromosomes 5, 7, and 12. The clustering is telling. It reads like a workshop bench where a successful tool is copied, the copy mutated,

the mutant kept or retired, generation after generation, until a whole drawer of related-but-distinct detectors sits ready. Where sweetness kept one tool, bitterness kept the drawer.

Each receptor in that drawer is a class A, rhodopsin-like G-protein-coupled receptor — among the smallest in the whole GPCR catalogue. Seven helices thread back and forth through the cell membrane; the agonist-binding pocket sits sunk *inside* that helical bundle, not perched on a large external domain the way the sweet and umami receptors carry their enormous Venus-flytrap folds. The consequence is architectural and it matters to us: bitter receptors are built to catch small molecules that dock down inside the membrane-spanning core, which is exactly the size and shape of the defensive compounds a plant makes. The receptor and the poison were designed for each other, from opposite directions, by the same long pressure.

The genius of the system is in how the tuning is spread. Some TAS2Rs are specialists — narrow, precise, answering to a single chemical family and little else. Others are deliberate generalists, a handful of them promiscuous enough to respond to dozens of structurally unrelated agonists apiece. Lay the narrow detectors and the broad ones over one another and roughly twenty-five receptors tile a chemical space of many thousands of bitter molecules, with overlap everywhere, so that a compound no receptor was built for still lands in the response field of two or three. That overlap is not waste. It is what lets the net generalize to a molecule never before met — the one property evolution needed and the one a fixed, single-lock receptor could never provide.

And the net is not confined to the mouth. The same TAS2R proteins are expressed far from the tongue — in the ciliated epithelium of the airway, in the solitary chemosensory cells that guard it, in the enteroendocrine cells lining the gut. Below conscious taste, the body runs the identical alarm against whatever arrives in the lumen, nudging motility, secretion, and reflex. Bitterness, read whole, is a body-wide chemical surveillance system, and the tongue is only its most articulate station — the one place where the alarm becomes a feeling you can name and act on. We hold that fact close when we design, because it is a reminder that we are never working on a mere preference. We are working on the readout of a sense the body trusts with its safety.

Inside the alarm — from a molecule to a nerve

When a bitter molecule reaches a taste bud, what follows is a cascade, and because every step in a cascade is a place a signal can be made louder or softer, it is worth walking the whole chain once, slowly. A grower, in the end, can only reach the very first step — but you cannot see why that first step is the right one to reach until you have seen the rest.

The agonist docks into the pocket within the helical bundle, held by residues contributed by the third, fifth, sixth, and seventh transmembrane helices. Binding nudges those helices through the small conformational "micro-switch" rearrangement that defines class A receptors, opening a cavity on the inner face of the cell. Into that cavity couples gustducin — a taste-specific G protein, close cousin to the transducin that carries vision — and the receptor catalyzes the

exchange of GDP for GTP on its α-subunit. The α-subunit separates from its βγ partners, and it is the βγ pair that carries the signal onward.

Gβγ switches on phospholipase C-β2, an enzyme that cleaves the membrane lipid PIP_2 into two messengers, diacylglycerol and inositol trisphosphate. The inositol trisphosphate diffuses to the endoplasmic reticulum and opens the type-3 IP_3 receptor there, and calcium — held in storage against exactly this moment — floods into the cytosol. That rising calcium gates TRPM5, a calcium-activated channel that admits sodium and depolarizes the cell. Depolarization, together with the calcium, opens the CALHM1/CALHM3 channel, through which the cell releases ATP directly — no classical synaptic vesicle, a bare efflux of the molecule the body uses everywhere as a currency of urgency. That ATP crosses to the gustatory nerve fibers pressed against the taste cell and activates their P2X2 and P2X3 receptors: the chorda tympani carrying the message from the front of the tongue, the glossopharyngeal from the back. Action potentials climb to the nucleus of the solitary tract in the brainstem, relay through the thalamus, and arrive at the gustatory cortex in the insula. Milliseconds, molecule to grimace. The wince you feel at a raw floret is the far end of a chain seven links long.

Here is why the walk matters. Of those seven links, six belong to the eater — they are fixed human biology, the same in every mouth, and nothing we grow will or should touch them. The plant speaks at the first link only: *how many agonist molecules it hands the pocket.* That is the entire surface a grower has to work on, and it is the honest one. We do not blunt a nerve or dampen a brain. We change what the plant presents to a receptor that goes on working exactly as it always has. Everything on this axis reduces to designing that first link — and to measuring, precisely, how loudly it still rings.

The receptor with a name — TAS2R38 and the taster gene

Among the twenty-five, one is famous enough to carry this axis's masthead, and it is the right one to carry it. TAS2R38 is the gene behind the classic PTC taste-test strip handed out in biology classrooms — the little paper that is unbearably bitter to some students and blank to others. Its agonists share a single chemical signature, the thiourea group, a nitrogen–carbon(=sulphur)–nitrogen motif: phenylthiocarbamide and 6-n-propylthiouracil, the laboratory standards, and — the part that matters at a family dinner — goitrin and the related products a Brassica makes when its defensive chemistry is broken open by chewing. TAS2R38 is, in large part, the receptor that builds the "green wall" of raw kale, cabbage, and mustard greens, because their chemistry yields precisely the thiourea-type molecules it is tuned to catch.

What makes TAS2R38 indispensable to anyone selling *less bitter* is that it is not the same receptor in every mouth. Three coding variants — at amino-acid positions 49, 262, and 296 — sort people into two common versions of the protein. One haplotype, carrying proline-alanine-valine at those sites, is the taster form; the other, alanine-valine-isoleucine, is the non-taster form. People with two taster copies are the "super-tasters" who find raw brassica genuinely punishing; people with two non-taster copies taste the same leaf as nearly neutral; and a sizeable fraction of any population — the

proportion shifting with ancestry — carries two non-taster copies and has, in effect, been eating a different vegetable than their super-taster neighbor all along.

This is not a footnote. It is a constraint we build into the mathematics. There is no single, true "bitterness of broccoli" to be measured and reduced, because the plant is read by a distribution of receptors, not one. So the Index defines its bitter axis against a fixed reference receptor and reports every value for a *population*, with the human variance stated plainly rather than averaged into a flattering single figure. When we claim an attenuation, we are claiming a shift of the whole distribution — the super-taster's experience and the non-taster's moving together toward ease — not a verdict from one privileged tongue. It is the kind of honesty the subject demands, and the kind our family was raised to prefer: a number that admits what it does not control.

The chemistry the plant actually makes

The alarm is only half the story; the other half is the chemistry that trips it, and in the plants we are most often asked to quiet, that chemistry is beautiful, specific, and — this is the heart of the problem — entangled with the nourishment we are trying to protect.

In the Brassica, bitterness is the output of a two-part defensive device the plant keeps disarmed until the moment of injury. Glucosinolates sit stored in the vacuole; the enzyme that activates them, myrosinase, is sequestered apart, in separate cells or separate compartments. Intact, they never meet. The instant a leaf is bitten, the compartments rupture, myrosinase reaches the glucosinolates, and it cleaves the thioglucoside bond. The unstable fragment left behind undergoes a Lossen rearrangement, and what it becomes depends on the conditions: usually an isothiocyanate, but where an epithiospecifier protein is present, a nitrile or epithionitrile instead. Sinigrin yields allyl isothiocyanate — the same sharp molecule that is the bite of mustard. Progoitrin yields goitrin, a direct agonist of TAS2R38. And glucoraphanin yields sulforaphane, the isothiocyanate the body is most glad to receive.

Read that list again and the dilemma of the whole axis is visible in a single pathway. The very reaction that lights the bitter receptor is the reaction that makes the nourishing compound. Reach for the obvious lever — knock out myrosinase, or strip the glucosinolates wholesale — and you do quiet the bitterness, but you destroy the sulforaphane in the same stroke, and you disarm the plant's own defense besides. That is a way to win the measurement and lose the point. The honest target is selective and far harder: down-tune the specific agonist species that shout loudest at TAS2R38 — the progoitrin-to-goitrin line and the most potent individual glucosinolates — while preserving glucoraphanin and the isothiocyanate it becomes. A scalpel held to one compound, not a hammer brought down on the pathway. Most of the real difficulty of this axis lives in that distinction.



Coffee — the bitterness here is mostly made in the roast, not grown in the bean; chlorogenic acids degrade under heat into the phenylindanes that give a dark cup its edge.

Every bitter food tells a version of this story in its own molecules. Coffee's bitterness, famously, is *not* mainly its caffeine — caffeine is only a modest bitter. The dominant bitters are chlorogenic-acid lactones in a lighter roast and, as the roast deepens, phenylindanes formed when those chlorogenic acids degrade under heat, which is why a dark roast tastes harsher than its caffeine content alone could ever explain. To ease coffee is to manage a thermal cascade and the precursor pool that feeds it. In hops, and so in beer, bitterness is the iso-alpha-acids — isohumulones — created when the hop's humulones isomerize in the boil, the exact molecules the brewer's IBU scale was built to count.

Tea and the grape sit at an instructive edge, one that defines a boundary we never blur. Their polyphenols — the catechins, the tannins — are both bitter *and* astringent, but those are two different senses riding on one molecule. Bitterness is a taste: a receptor reading a shape. Astringency is a touch: tannins cross-linking and precipitating the proline-rich proteins that lubricate the mouth, felt as drag and pucker, with no taste receptor involved at all. A single catechin can ring both bells at once. We score them on separate axes because they are separate machines, and conflating the drag of a tannin with the bitterness of a receptor is the most common error anyone makes in reading a bitter food. On this axis we measure only the receptor.

A number for a loudness — bitterness on the Index

A sense this old and this variable resists description. The Iyengar Flavor Index answers it the only way a family of mathematicians could be satisfied with: by turning loudness into a number anchored to the receptor itself. The reference constant is set by convention — the way the Scoville scale is defined by its own dilution threshold, not derived from any single measured literature value. The formula is one line.

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the midpoint of the dose–response curve. Write that curve out and the choice of anchor justifies itself. Receptor response follows a Hill relationship, $R = R_{max} \cdot [A]^n \div (EC_{50}^n + [A]^n)$;

set the concentration [A] equal to EC_{50} and the expression collapses, for any value of the Hill slope n, to exactly R_max ÷ 2. The half-maximal point is the one place on the curve that is stable, reproducible, and indifferent to how steep the response happens to be — unlike a threshold, which drowns in noise, or a maximum, which saturates and is hard to reach. We anchor where the biology is most honest.

Because the score is a ratio of two concentrations, the units cancel. IFI-B is dimensionless — a pure number, scaled by a thousand — which is the whole trick that lets a cup of coffee, a kale leaf, and a chili be read against one another on a single page. And the direction is worth fixing in the mind: a *lower* EC_{50} means the food reaches half-maximal bitterness at a *lower* concentration, which is a *louder* bitter, which is a *higher* IFI-B. Native strong bitters read high; easing a food lowers its score.

A worked reading makes it concrete. Take a fixed reference bitter with an EC_{50} of 300 µM, and a food whose agonists give it an effective EC_{50} of 180 µM. Its loudness is IFI-B = 1,000 × 300 ÷ 180 = 1,667. Now ease that food so its agonists reach half-maximal activation only at 225 µM — the receptor must now meet 25% more before it answers — and the score becomes 1,000 × 300 ÷ 225 = 1,333. The bitterness has fallen from 1,667 to 1,333, a drop of 20.0%, bought with a 25% rise in the concentration the receptor demands. That inverse exchange — a little less loudness for a little more required concentration — is the engine of the entire axis, and we will return to it in a moment with the full arithmetic.

The per-axis scorecard normalizes every reading to a ceiling of 5,000, and it is on that scale that we can show the one plant in our whole library designed to carry a *deliberate* bitterness. Phoenix — the cultivar we engineered, gene by gene, to register on all eight axes at once — does not read zero on the bitter axis. We designed it, on purpose, to a chosen and moderate 2,200, which is 44% of full scale, because a trace of genuine bitterness is not a flaw to be erased but a structure to be kept: the backbone that good coffee and good dark chocolate are built around. Bitterness dialed to a target, rather than merely silenced, is the proof the axis is a dial and not only an alarm.



Phoenix (SS-0096) in the controlled chamber — the one cultivar in the library whose bitterness is a chosen set-point, held at IFI-B 2,200 rather than erased.

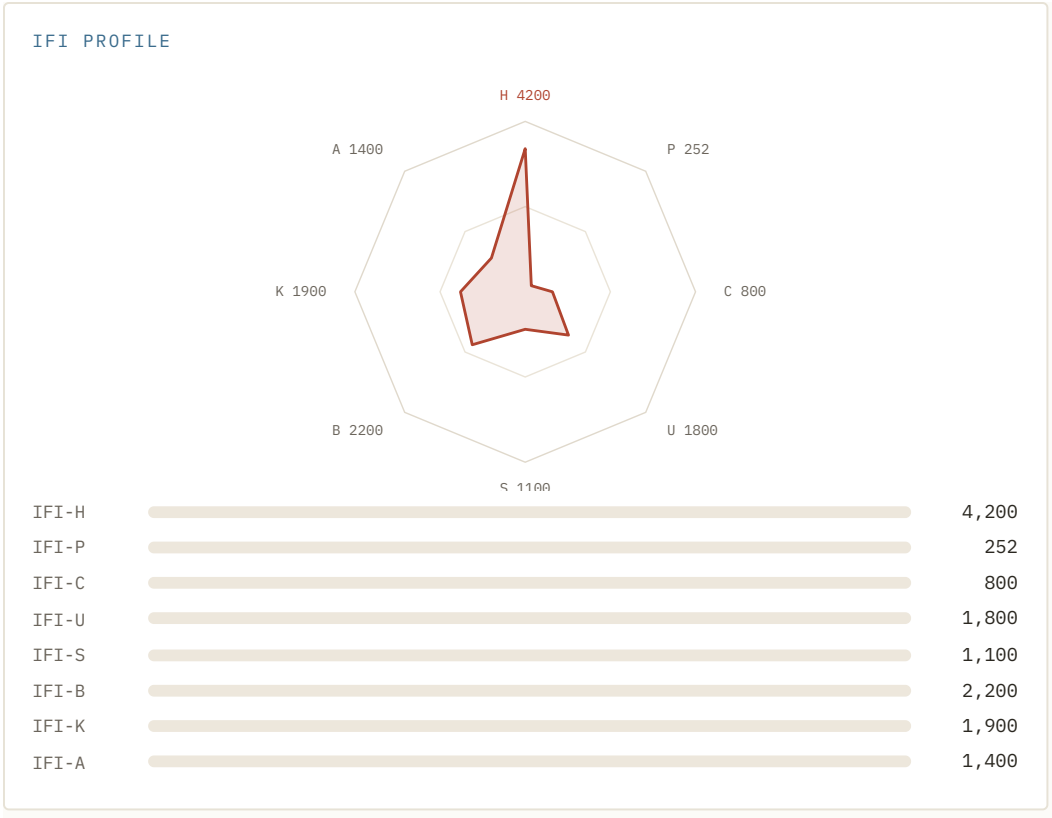

Phoenix is the exception that proves what the axis is for. The rule — the daily work — is the other direction: the plants whose bitterness we are asked, over and over, to quiet. These are the hosts in our library where Zero-Bitter is a primary axis, the shelf this book is written to reach.

HOST ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Grape	1,610	B / A / M	Hard
Broccoli	380	B / M	Moderate
Tea	280	U / B / A	Hard
Broccoli-Kale	280	B	Hard
Orange	120	B / C	Hard
Potato	80	B	Easy
Hops	80	B	Moderate
Coffee	40	B	Hard
Cabbage	40	B	Moderate
Kale	40	B	Moderate
Radish	40	B	Hard
Rapeseed	40	B	Hard

The arithmetic of attenuation

Set the exchange down formally, because the whole of what we promise on this axis — a 20% to 65% attenuation of bitterness — is contained in it, and a reader owed a number is owed its derivation.

Let a native food have bitter loudness L_0, and let *a* be the attenuation — the fraction by which we lower that loudness. The engineered food then reads $L_1 = (1 - a) \cdot L_0$. But loudness, by the Index formula, is inversely proportional to EC_{50}: double the concentration a receptor needs, and you halve the score. So lowering the loudness by the factor $(1 - a)$ requires raising the effective EC_{50} by exactly its reciprocal:

$$EC_{50}(\text{engineered}) \div EC_{50}(\text{native}) = L_0 \div L_1 = 1 \div (1 - a)$$

That single relation says, in plain terms, that attenuation is a demand on the plant to hand the receptor *fewer agonist molecules* — to make half-maximal activation retreat to a higher concentration. The canonical band of 20% to 65% maps onto a clean, honest range of EC_{50} shifts.

ATTENUATION *A*	REMAINING SIGNAL (1 − A)	EC_{50} RATIO 1 ÷ (1 − A)	CONCENTRATION TO REACH HALF-MAX
20%	0.80	1.25×	+25%
35%	0.65	1.54×	+54%
50%	0.50	2.00×	+100%
65%	0.35	2.86×	+186%

Read the two ends against the worked food from before. At the gentle end, a 20% quieter cup asks the receptor for 25% more before it speaks: an EC_{50} moved from 180 µM to 225 µM, a score from 1,667 to 1,333. At the deep end, 65% quieter is a 2.86× EC_{50} — the same food's 180 µM pushed out to 514 µM, its loudness falling from 1,667 to 583, the receptor now needing nearly three times the concentration to ring at the old volume. Everything we deliver on this axis lives between those two lines. The reason it stops there, at both ends, is not timidity. It is the subject of the next pages, and it is where the chemistry, the statistics, and the ethics turn out to agree.

Why the floor is twenty and the ceiling is sixty-five

Two honest questions sit under the band. Why does attenuation start at 20% — why not claim the 8% improvement, the 12%? And why does it stop at 65% — why not chase bitterness to zero? The answers come from opposite directions, one from the instruments and one from conscience, and we hold both.

The floor is set by measurement, not by ambition. Every EC_{50} is a measured quantity, and measured quantities scatter; the ratio that becomes an IFI score inherits the scatter of *both* of its terms. For two independent measurements each carrying a coefficient of variation CV, the combined relative uncertainty is the root-sum-square, $\sqrt{(CV^2 + CV^2)} = CV \cdot \sqrt{2}$. Our process holds batch-to-batch CV to 7% or below; combined across the two terms of the ratio, that is $7\% \times \sqrt{2} \approx 9.9\%$. And a change is only credible when it clears roughly twice the noise — about $2 \times 9.9\% \approx 20\%$. So the 20% floor is not a marketing choice. It is the smallest shift in bitterness our instruments can tell apart from an ordinary batch-to-batch wobble. We decline to claim anything we could not distinguish from luck.

PROCESS	BATCH CV	COMBINED CV (√2 · CV)	SMALLEST RESOLVABLE Δ (≈2Σ)
Our engineered line	≤7%	≈9.9%	≈20%
Wild-type baseline	42%	≈59%	≈119%

Set beside a wild-type process the discipline becomes stark. Conventional, unselected material scatters with a batch CV near 42%; combined, that is roughly 59% of noise, and nothing short of a 119% change is resolvable against it. A native crop could lose fully half its bitterness and the measurement would be unable to prove it happened. The 7% discipline is precisely the thing that converts

bitterness from an opinion anyone can argue into a number anyone can check. Tightening the process is not housekeeping; it is what makes the whole axis real.

The ceiling is set by reverence. Sixty-five percent attenuation leaves 35% of the native bitter signal standing, and it is left standing on purpose, for two reasons we will not trade away. The first is safety: the alarm is a safety system, and a tongue that cannot taste bitterness at all is a tongue that cannot taste a genuine toxin. We will not take that from anyone, however much they dislike their greens. The second is nourishment and pleasure together: the chemistry that makes the bitterness is so often the chemistry that makes the compounds the body is glad to receive — the isothiocyanate line in the Brassica, the polyphenols of coffee and cacao — that to strip bitterness to zero is to risk stripping the nourishment and the character in the same motion. A square of chocolate with no bitterness left is not improved; it is gone. Thirty-five percent retained is our deliberate margin, and we have never regretted keeping it.

The failure modes are worth naming without euphemism, because designing against them is most of the craft. Knock out myrosinase and you quiet the bitter and kill the sulforaphane. Collapse the glucosinolates wholesale and you lose the same nourishment and disarm the plant's defense besides. Chase the number past the point where an eater can still tell a spoiled or toxic food from a safe one, and you have optimized a figure into a hazard. Each of these wins a measurement and loses the meal. The 20%-to-65% band is simply where the three honest constraints — what we can prove, what the plant can spare, and what a body is owed — all hold at once.

What we will not do

So we return, with the mechanism and the mathematics now in hand, to the sentence the whole axis is built on, and we can at last say it with its full weight behind it. We are not here to silence the alarm. The bitter family is twenty-five sentinels of ancient and excellent design; the signal they carry is the far end of a seven-link chain the body trusts with its safety; and the chemistry that trips them is, half the time, the chemistry that feeds us. None of that is a thing to switch off. What we are here to do is narrower, and we think more respectful: to teach a few well-chosen plants to stop ringing the alarm over food that was never a threat — selectively, by the single compounds that shout loudest; measurably, as a shift of a whole population's experience rather than one lucky tongue's; and always within a margin that leaves the warning, the nourishment, and the character intact.

Picture again the shelf most densely provisioned and most often set back down untouched — the deep greens, the strong cup, the high-cacao bar, the whole bitter fruit. The alarm that once kept our ancestors from poison now stands between an eater and exactly the food grown to nourish them best. Our family came up at a vegetarian table, taught to treat what we grow with reverence and to give the body what it needs with intention, and this is the form that reverence takes here: not a trick played on the tongue, but a quiet, surgical correction made in the plant — so that the hand reaches the garden shelf, and the receptor, reading honestly, finds nothing there worth

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

Chapter 2 · TAS2R38, the Anchor — *one receptor to measure a thousand bitters.*

the reference receptor · the PTC locus · the PAV/AVI divide

Twenty-five bitter receptors are too many to measure against. If the Index is to give bitterness a single honest number — one that lets a floret of broccoli and a cup of coffee and a square of chocolate be read on the same page — it needs one reference receptor to anchor the axis, the way IFI-H is anchored to TRPV1 and IFI-U to the glutamate receptors. We chose **TAS2R38**, and the choice was not close. What follows is the whole case for it: why one voice can stand for a choir, the chemistry it was born to read, the machine it is, the arithmetic by which it turns a mouthful of bitterness into a score, and the three letters of genetics that keep the whole axis honest.

Why one voice has to carry the choir

The human bitter sense is not one receptor but a family of about twenty-five, the TAS2Rs, scattered across three chromosomes and expressed in overlapping sets on the taste cells of the tongue and palate. Evolution built them as a smoke alarm, not a spectrometer: their job for a hundred thousand years was to make a mouthful of something alkaloidal or thiourea-laced unmistakably unpleasant *fast*, before the second bite. Breadth beat precision. A single TAS2R will answer to dozens of unrelated molecules, and a single bitter molecule will light up several TAS2Rs at once, so the real bitterness of a food is a chord struck across many of them — and a chord cannot be a coordinate.

A standard needs a coordinate. The entire Index rests on one move: turn a receptor's sensitivity into a number by asking not *how much* of a compound a food holds but *how loudly* that compound speaks to the receptor that reads it. Loudness is captured by the half-maximal concentration, the **EC_{50}** — the amount of a stimulus that drives its receptor to half of full response. A small EC_{50} means the receptor answers to a whisper; a large one means it takes a shout. Invert that against a fixed reference, scale by a thousand, and you have the formula that every axis obeys:

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

The inversion is the heart of it. Because loudness runs *opposite* to the concentration needed to be heard, we divide — a candidate half as potent as the reference needs twice the concentration to reach half-activation and so scores half as loud. By construction the reference itself scores exactly 1,000, and the scorecard caps each axis at **5,000**, the pin a stimulus five times more potent than the reference would reach. But none of this exists until we name the reference. For

bitterness, the reference is TAS2R38, and naming it well is the difference between a number anyone can cite and a number no two labs would agree on.

The strip of paper

TAS2R38 is the best-characterized bitter receptor in the human genome, and the reason is a strip of paper. In 1931 a chemist at DuPont, Arthur Fox, released a cloud of phenylthiocarbamide dust in his lab; a colleague across the bench tasted it as intensely bitter while Fox tasted nothing at all. That accident opened one of the oldest and cleanest stories in human sensory genetics. The ability to taste **PTC** (phenylthiocarbamide) and its safer laboratory cousin **PROP** (6-*n*-propylthiouracil) is largely governed by this one gene, and it splits humans into tasters and non-tasters with a sharpness almost no other trait in perception can match. Roughly three-quarters of most populations taste PROP; the remaining quarter are nearly blind to it. That bimodal split — not a smooth bell curve of sensitivity but two populations with a gap between them — is visible on a paper strip at the back of any classroom, and it is why TAS2R38 has been sequenced, mutated, expressed in cultured cell lines, and docked against its ligands more thoroughly than any of its two dozen siblings.

For a reference receptor, that depth of knowledge is the whole point. An anchor must be the thing you know cold, because every other reading in the axis inherits its error. A reference whose dose-response curve wandered batch to batch, whose structure was guessed at, whose pocket was a mystery, would smear uncertainty across every bitter food ever scored against it. TAS2R38 is the opposite: a receptor whose behavior has been pinned down from the gene to the glass. When the reference is this well known, the division in the formula is clean, and the thousand foods measured against it all share one trustworthy zero-point.

The chemistry it was tuned to read

The second reason runs deeper than convenience — it is about what TAS2R38 is *for*. Of all the bitter sentinels, it is the one most sharply tuned to the chemistry of the brassicas: to the thiourea edge of glucosinolate-derived compounds, the sharp isothiocyanates and the cyclized goitrin that make raw kale bite and overcooked cabbage turn harsh. Its natural agonists carry a shared chemical signature, the **N–C=S** group — the thiocarbamide of PTC and PROP, the N=C=S of an isothiocyanate — and that motif is precisely what the brassica defense chemistry presents to the tongue.



The cut face of engineered broccoli — the plant whose defense chemistry TAS2R38 evolved to read, and where our subtraction does the most generous work.

The brassica stores its warning disarmed. Glucosinolates — sinigrin, progoitrin, glucoraphanin — sit intact and tasteless in the tissue until the cell wall breaks. The moment a jaw or a knife ruptures the cell, an enzyme held in separate myrosin cells, **myrosinase**, meets its substrate and strips the sugar, leaving an unstable aglycone that rearranges in milliseconds (a Lossen-type rearrangement) into the volatile isothiocyanate — allyl isothiocyanate from sinigrin, sulforaphane from glucoraphanin — or, from progoitrin, cyclizes into **goitrin**, the heaviest and harshest note in the chord. Which product forms is not fixed: at neutral pH the rearrangement runs to the sharp isothiocyanate, but a specifier protein in the tissue — the epithiospecifier — can divert it toward milder nitriles instead, so two heads of the same broccoli, chewed by two people, can present the receptor with different agonists. Every one of the isothiocyanate products wears the N–C=S edge, and goitrin, with its thiocarbonyl locked into a ring, wears it where the receptor grips hardest, which is part of why it reads so loud. When a child tastes broccoli as a threat, TAS2R38 is, more often than any other receptor, the pathway doing the saying.

That is why anchoring the bitter axis to the brassica receptor is not merely defensible but honest. The plants where our subtraction does the most good — where a few points off the bitter axis is the difference between a vegetable eaten and a vegetable pushed to the rim of the plate — are the brassicas, and the receptor that reads the brassicas is the one we chose as the reference. The axis is anchored where the work lives. It is also where the bulk of our bitter-axis library lives: of the host organisms we engineer on the Zero-Bitter axis, the plurality are Brassicaceae, and the receptor that defines the axis is the one each of them speaks to first.

ORGANISM	CONSTRUCTS	FAMILY	BITTER AXES
Grape	1,610	Vitaceae	B / A / M
Broccoli	380	Brassicaceae	B / M
Tea	280	Theaceae	U / B / A
Broccoli-Kale	280	Brassicaceae	B
Orange	120	Rutaceae	B / C
Potato	80	Solanaceae	B
Hops	80	Cannabaceae	B
Coffee	40	Rubiaceae	B
Cabbage	40	Brassicaceae	B
Kale	40	Brassicaceae	B
Radish	40	Brassicaceae	B
Rapeseed	40	Brassicaceae	B

Inside the helical bundle

Structurally, TAS2R38 is a class A, rhodopsin-like G-protein-coupled receptor: seven membrane-spanning helices threaded through the taste-cell surface. Unlike the great Venus-flytrap domains of the sweet and umami receptors, which clamp their ligands in a huge extracellular jaw, a bitter receptor has only a short external tail. A bitter agonist does not dock on an outside platform; it slips *into* the helical bundle itself, settling into a pocket lined by transmembrane helices three, five, six, and seven. From inside the membrane it levers the receptor's shape just enough to wake the machinery on the other side. The pocket is roomy and forgiving, which is exactly what a broad sentinel needs: it is how one receptor answers to PTC, to PROP, to goitrin, and to a scatter of otherwise unrelated molecules that happen to present the right electronic edge.

Inside that crevice the thiocarbonyl does the binding work. The sulfur of the N–C=S group is an electron-rich hydrogen-bond acceptor, and the receptor cradles it against polar side chains on helices five, six, and seven while the agonist's hydrophobic tail — the propyl of PROP, the vinyl of goitrin — packs against the lipid-facing wall. Two of the three taster positions sit on this inner surface: residue 262 on helix six and residue 296 on helix seven line the pocket directly, so their substitutions reshape the cradle itself, while residue 49, higher up toward the first helix, governs how well the finished receptor folds and reaches the cell surface at all. One position controls how much receptor there is; two control how firmly each copy grips.

What happens after binding is a cascade worth walking step by step, because the Index reads its output and because the shape of the curve we measure is set here. The agonist's arrival nudges helix six at the intracellular face, opening a coupling site for the taste cell's G-protein — **α-gustducin** with its βγ partners. The complex splits; the βγ subunits activate **phospholipase C β2 (PLCβ2)**, which cleaves the membrane lipid PIP_2 into **IP_3** and diacylglycerol. IP_3 crosses to the endoplasmic reticulum and opens the IP_3 receptor there, releasing a pulse of **Ca^{2+}** into the cytosol. That calcium gates **TRPM5**, a cation channel, and sodium rushes in; the cell depolarizes. Depolarization opens **CALHM1/CALHM3** channels, which release **ATP** — not in a vesicle but straight through the pore — onto the afferent nerve fibers waiting beneath. ATP lands on **P2X2/P2X3** receptors on those fibers,

and the signal finally leaves the tongue: to the nucleus of the solitary tract, to the thalamus, to the gustatory cortex, where it becomes the conscious fact of *bitter*.



The sequencer bench — the PAV/AVI haplotype at the heart of this axis is read here, three single-letter changes that decide whether a food tastes like a warning.

Two features of that chain matter for measurement. First, it is an amplifier: one agonist in the pocket yields many IP_3 molecules, each gating a flood of calcium, so the system is steep near its threshold — small changes in concentration produce large changes in firing. Second, it is a chain, which means a reading on the axis is a reading of the *whole* pathway's half-point, not of the binding step alone. When we measure EC_{50}, we are measuring where this entire machine crosses half of its full roar.

The arithmetic of a bitter number

Now the formula earns its keep. A receptor's response to rising agonist does not climb in a straight line; it traces a sigmoid, flat at the bottom, steep through the middle, flat again at saturation. The clean algebra for that S-curve is the **Hill equation**:

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

Here R is response, [A] is agonist concentration, R_max is the saturating maximum, and *n* is the Hill coefficient, the steepness of the rise. Set [A] equal to EC_{50} and the equation collapses to R = R_max ÷ 2 — which is the definition of EC_{50}, the concentration that buys exactly half the maximum response. The amplifying cascade above is why *n* for the bitter pathway sits above one: the curve is steeper than a simple single-site binding, and a steep curve is what turns a smooth range of sensitivities into the hard taster / non-taster gap on the paper strip. Because the curve is a straight sigmoid only when plotted against the *logarithm* of concentration, a tenfold change in potency is one decade — one clean step left or right along the dose axis.

Measuring an EC_{50} is itself a dilution experiment. We express TAS2R38 in cultured cells alongside a reporter that lights when the receptor fires — a promiscuous G-protein that funnels bitter signaling into a calcium flash a plate reader can see — then present the agonist at a ladder of concentrations and fit the responses to the Hill curve. The concentration at the curve's half-height is the EC_{50}, read straight off the fit. A whole food is scored the same way, by dilution rather than by dose: we take the bitter extract, step it down a dilution series, and find the point at which it drives the reference receptor to half-max. That

dilution *is* the food's effective potency — a coffee that still half-activates TAS2R38 after a fortyfold dilution is speaking far louder than one that falls silent at fivefold — and it is that effective half-point, not any single compound, that the formula turns into the food's IFI-B.

To make the arithmetic concrete we fix the anchor. We take PROP read by the taster form of TAS2R38 as the reference stimulus and define its half-activation as the axis's 1,000-point — a reference EC_{50} we hold at 50 µM. (Fixing the reference is not a measurement; it is the act of choosing a zero, the same act that puts 0 at the freezing point of water.) PROP's molar mass is 170.23 g/mol, so that reference concentration converts to a mass figure a formulator can weigh:

```
50 µM × 170.23 g/mol = 8.51 mg/L
```

Every other bitter stimulus is now a division against that anchor. A compound twice as potent — half-activating at 25 µM — scores 1,000 × 50 ÷ 25 = **2,000**. One half as potent, at 100 µM, scores 1,000 × 50 ÷ 100 = **500**. One four times weaker, at 200 µM, scores **250**. The reference, dividing into itself, scores **1,000** exactly, as it must.

STIMULUS	EC_{50} (MM)	EC_{50} (MG/L)	IFI-B = 1,000 × 50 ÷ EC_{50}
Twice the reference potency	25	4.26	2,000
PROP — the reference anchor	**50**	**8.51**	**1,000**
Half the reference potency	100	17.02	500
One-quarter the reference potency	200	34.05	250

A number without its error bar is a boast, not a measurement, so the last piece of the arithmetic is honesty about precision. We hold our batch-to-batch coefficient of variation to **≤7%**, and the Index score is a ratio of two measured quantities — the reference EC_{50} and the candidate's — each carrying its own scatter. For a ratio of independent measurements, relative variances add in quadrature:

```
CV(IFI) ≈ √( CV(ref)² + CV(candidate)² ) = √(0.07² + 0.07²) ≈ 0.099
```

So an IFI-B of 1,000 is properly written 1,000 ± 10% — roughly ±100 — and tightening the candidate measurement to a 5% CV pulls the compound figure to about ±9%. That is why the scorecard rounds and why we never publish a bitter score to a precision the chemistry cannot keep. The error is not a flaw in the number; it is part of the number, and stating it is a form of respect for whoever reads it.

Three letters that split the table

The famous taster divide lives in three positions along those helices. The common haplotypes of TAS2R38 differ at amino-acid residues **49, 262, and 296**, and the two that dominate human populations carry the signatures **PAV** — proline, alanine, valine — and **AVI** — alanine, valine, isoleucine. PAV is the taster form: it folds and traffics and couples efficiently, and it answers to PTC at low concentration. AVI is the non-taster form: the same three substitutions subtly reshape the helices and the pocket, and the receptor responds weakly or not at all. In the arithmetic above, the AVI pocket simply pushes the EC_{50} up the dose axis by more than a decade — frequently past the point where

the compound even stays in solution, which the tongue reports as *nothing*. Same molecule, same formula, a tenfold-and-more shift in where half-activation falls.

The variation is carried by three single-nucleotide changes — catalogued as *rs713598*, *rs1726866*, and *rs10246939* — and because they travel together on the chromosome they behave almost as a single switch. A person with two PAV copies is a strong taster; two AVI copies, a non-taster; one of each, an intermediate who registers PROP as present but bearable. They travel as one because the three sites sit in tight linkage disequilibrium — close enough on the chromosome that recombination rarely splits them — so PAV and AVI account for the overwhelming majority of copies, and the rarer recombinant haplotypes (AAI, PVI, and their kin) are both uncommon and functionally in-between, exactly as the structural division predicts: swap one pocket residue and you get a partial receptor, not a dead one. The single position that carries the most weight is 49, the folding residue, for the plainest reason in the cascade — a receptor that never reaches the surface cannot taste anything, however well its pocket is shaped. One gene, three letters, and the same floret of broccoli becomes two different foods at the same table.

HAPLOTYPE	RESIDUES 49 · 262 · 296	PROP / PTC RESPONSE	SHARE OF POPULATION (TYPICAL)
PAV / PAV	Pro · Ala · Val (×2)	strong — intensely bitter	~30%
PAV / AVI	one of each	intermediate	~45%
AVI / AVI	Ala · Val · Ile (×2)	weak to none	~25%

There is a quiet piece of mathematics hiding in that last column. Read the PAV allele's frequency straight off the genotypes — every PAV/PAV contributes two copies, every heterozygote one — and it lands near (2 × 0.30 + 0.45) ÷ 2 ≈ 0.52. Feed that frequency back through Hardy-Weinberg and it predicts a genotype split of roughly 28 / 50 / 23, within measurement of the observed 30 / 45 / 25. The taster divide is not some fragile quirk; it is a balanced, equilibrium polymorphism, carried at near-even odds through most of humanity — which is exactly why a bitter standard cannot pretend it away.

What subtraction does to the score

Zero-Bitter is subtraction. On this axis we design plants to attenuate bitterness by **20% to 65%**, and the formula tells us precisely what that buys. Below the EC_{50}, where most everyday bitterness is read, the Hill curve is nearly linear in concentration, so reducing the effective bitter agonist reaching the receptor by a fraction lowers the loudness — and the score — by that same fraction, to first order. A 20% subtraction multiplies IFI-B by 0.80; a 50% subtraction halves it; the deepest 65% subtraction takes it to roughly a third, ×0.35. The full designed band carries a bitter reading from four-fifths of its wild loudness down to about a third of it. Put in the dilution language of the assay, a 50% subtraction is a halving of the effective agonist: a brassica that once needed a fortyfold dilution to fall to half-activation now reaches that same half-point at twentyfold, and its score drops in step — the cut is visible in the glass before it is ever visible on the tongue.

The molecular levers behind that multiplier are the brassica's own chemistry, read backward. We can turn down the aliphatic glucosinolate pool at its source, dialing the R2R3-MYB regulators (the MYB28/MYB29 family) that govern how much precursor the plant lays down. We can bias the hydrolysis away from goitrin and the harshest isothiocyanates and toward the gentler nitrile products. Or we can shift the profile among glucosinolate species toward the ones TAS2R38 reads least loudly. Each lever lowers the effective agonist at the pocket; each shows up as fewer points on the axis. None of it is a claim about the body — the plainest arithmetic of nourishment is only that a vegetable eaten feeds you and a vegetable refused does not, and a child who can taste past the warning gets the whole plant. We subtract the harshness that stood between them and the broccoli, and our model specifies exactly how much is subtracted.



Coffee — the archetype of a bitterness worth tuning rather than erasing, read on the same axis as the brassica it shares a receptor edge with.

The subtraction has an honest ceiling, and the Hill curve names it. The clean proportional mapping of percent-attenuated to percent-off-the-score holds only in the sub-EC_{50} regime. Near and above half-activation the sigmoid flattens into saturation, and there the same molecular subtraction buys far less relief — halving the agonist on the shoulder of the curve might move the perceived bitterness only slightly, because the receptor was already near its roar. This is the real edge case of Zero-Bitter: the most savagely bitter foods, the ones already pinning the receptor, are the ones where a given subtraction returns the least score. We design against it by working the profile, not just the dose — moving the agonist off the steep part of the curve rather than merely trimming it.

Not every engineered bitterness is something to erase. The counter-case is the plant we built to register on all eight axes at once, where bitterness is not a flaw to subtract but one deliberate color in a full chord — a designed presence, placed against this very anchor:

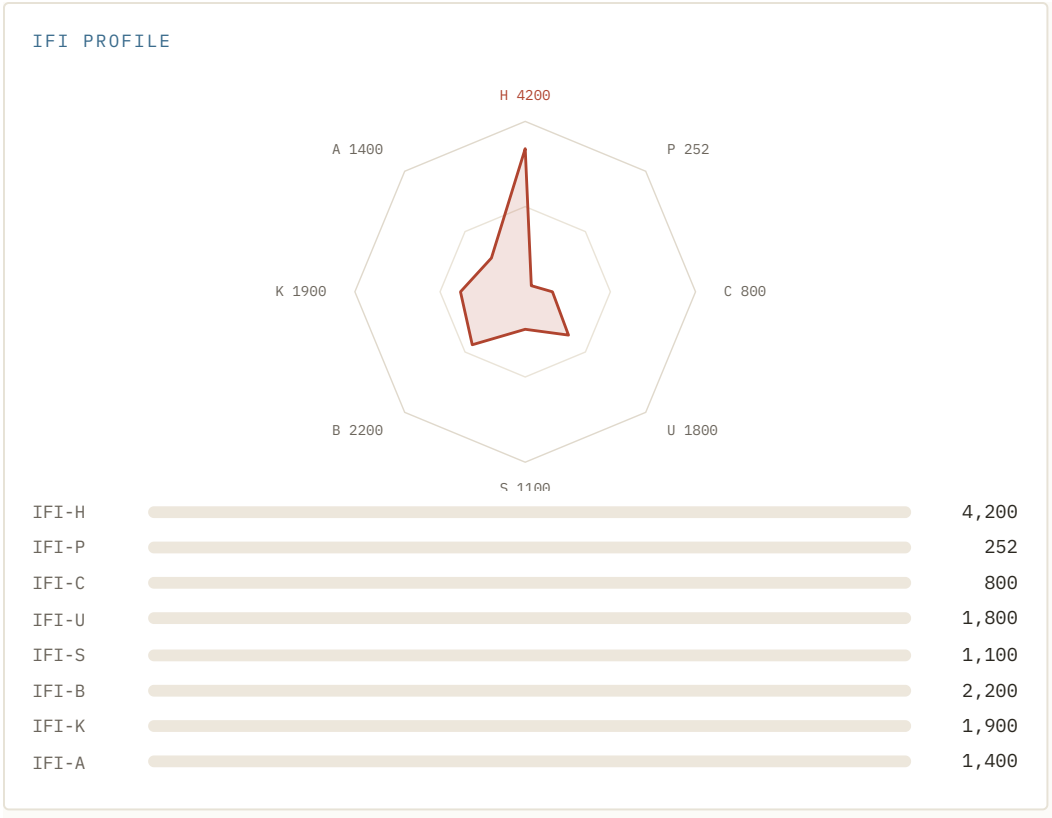

That B of 2,200 is not an accident the subtraction missed; it is a bitterness placed on purpose, tuned to sit in balance beneath the heat and the umami and the kokumi, and the only reason we can speak of it as *placed* rather than merely *present* is that TAS2R38 gave us the ruler to place it with.

The chord and the anchor

Here is the honest limit, and it is the same fact we opened with, returned as a warning. A real bitterness is almost never one receptor's work. It is a chord struck across several TAS2Rs at once — TAS2R38 reading the thiourea edge, its siblings answering to the alkaloids and the hop acids and the chlorogenic compounds in their own registers — and the sensation you call *bitter* is their sum. We measure that chord, as every lab must, against the one voice we know by heart. The anchor does not claim to *be* the whole chord; it claims to be the fixed, fully-mapped reference against which the chord can be given a number at all. A thermometer does not contain the weather.

And the three letters keep their own caution. Because the taster polymorphism is carried at near-even odds, IFI-B can never be a single lived experience — the AVI non-taster and the PAV taster are eating genuinely different foods at the same table, and no one number speaks for both mouths. What the axis gives instead is something a standard actually can promise: an honest reading against a **defined reference genotype**, the taster form, stated with its error bar and reproducible to the ≤7% our chemistry can hold. That is the discipline of the whole axis in one line — we do not pretend bitterness is one thing for everyone; we pin it to the one receptor we understand completely, measure the chord against it, and show our work. TAS2R38 is the anchor. Everything else on the bitter axis is read against the one note we can sing from memory. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 3 · The Choir — *the other twenty-four sentinels.*

the TAS2R family · narrow and broad tuning · the chord a food strikes

We anchored the axis to one receptor, because a standard needs one ruler; but a ruler is not the whole of what it measures, and the honest picture of bitterness is not a solo but a choir. TAS2R38 is one of **twenty-five functional bitter receptors** written into the human genome — the *TAS2R* family, a cluster of genes that arose by repeated duplication as our lineage met poison after poison and kept a copy of each guard it needed. Where sweetness makes do with a single receptor pair and umami with two, bitterness fields a committee of twenty-five, and the committee is the whole point. No single sentinel could cover the unbounded chemical space of things that might harm a body, so the lineage built many, tuned them differently, and let them vote. To measure that chorus honestly — to put one number on it a tomato, a coffee bean, and a cup of tea can be read against on a single page — we first have to hear every voice in it, and then understand why, for the work of quieting a bitterness, those twenty-five voices can be treated, mercifully, as one.



A receptor panel in culture — reading a single food across the bitter family means watching which of the twenty-five sentinels light up, and how brightly.

The family that poison built

The *TAS2R* genes are a textbook case of what biologists call a birth-and-death gene family. They sit in two clusters — a large one on chromosome 12 and smaller company on chromosome 7 — and they got there the slow way, by one gene being accidentally copied, the spare copy drifting free of its original job, and selection keeping whichever drifted copies happened to catch a molecule worth being warned about. Run that loop across the deep history of an omnivore, through every leaf and root and unfamiliar fruit a wandering diet ever tested, and you arrive at twenty-five working receptors standing

beside roughly eleven broken ones — pseudogenes, the fossils of guards we no longer keep lit. The broken ones matter as much as the working ones: they are the record of a committee that was once larger, trimmed to the set a cooked, cultivated, more trusting diet still needs.

The committee has a logic, and it is worth sitting with, because it is the same logic that shapes how we measure. A single broadly tuned guard — one receptor that answered to *everything* faintly — would be a smoke detector with the sensitivity turned to zero: it would miss the narrow, specific poisons that do the real harm, and it would cry wolf at the harmless. A single narrowly tuned guard would catch its one poison beautifully and be blind to all the rest. So the family splits the difference across twenty-five members, some narrow and exact, some broad and forgiving, each covering a different patch of chemical space, their ranges overlapping like the fields of fire of guards on a wall. The brain never sees the twenty-five separate reports. It reads only the summed alarm — *bitter, this much* — and spits, or doesn't. Evolution built a parliament and handed the eater a single verdict. Our task, measuring, is to respect both the parliament and the verdict: to know the members, and to score what the eater actually feels.

One door, seven turns

To measure a receptor you have to know what happens when a molecule finds it, because the number we assign is a number about that event. Each of the twenty-five is a class-A G-protein-coupled receptor — the most common architecture of signalling protein in the body, the same broad design that reads adrenaline and light and the scent of a rose. Seven alpha-helices thread back and forth through the cell membrane, lashed together into a barrel. A bitter molecule does not dock on some large reaching arm the way a sugar settles into the open clamshell of the sweet receptor; the bitter pocket is shallow and sits *within* the helical barrel itself, which is exactly why the family can retune so freely — change a few residues lining that small interior pocket and you change what fits, without rebuilding the whole protein.

What follows, when the right molecule seats itself, is a cascade of seven turns, and tracing it once is worth the patience because every one of the twenty-five doors opens onto the very same hallway:

1. **The bind.** The agonist settles into the pocket inside the barrel, and the contacts it makes pull helix six outward at the receptor's base — a small swing, two or three ångström, but enough to open the intracellular face of the protein like a hand uncurling.
2. **The exchange.** Waiting against that face is α-gustducin, the taste cell's own G-protein, carried in a heterotrimer with Gβ3 and Gγ13. The opened receptor acts as a switch that lets gustducin trade its spent GDP for a fresh GTP, and the trio breaks apart.
3. **The messenger enzyme.** The freed Gβγ pair finds and activates phospholipase C beta-2 (PLCβ2), an enzyme sitting at the inner face of the membrane.
4. **The cleavage.** PLCβ2 cuts a membrane lipid, PIP2, into two second messengers — IP3 and diacylglycerol — and it is IP3 that carries the signal inward.

5. **The calcium.** IP3 crosses to the endoplasmic reticulum, the cell's internal store, and opens the IP3 receptor channel there; calcium floods out of the store into the cytosol. This is the signal made loud — one bound molecule at the surface becoming a wash of ions inside.
6. **The current.** That calcium gates TRPM5, a channel selective for sodium and its monovalent kin; sodium pours in, and the cell depolarizes.
7. **The release.** The depolarized cell opens CALHM1 (with its partner CALHM3), a large-pore channel, and pushes ATP straight out through it — no vesicle, a strange and specific trick of the taste cell. That ATP lands on P2X2 and P2X3 receptors studding the gustatory nerve fiber pressed against the cell, and the nerve fires, up the chorda tympani and the glossopharyngeal, into the nucleus of the solitary tract in the brainstem, where *bitter* first becomes a thing felt.

Seven turns, and here is the point that makes the whole method possible: steps three through seven are shared. Whichever of the twenty-five doors the molecule struck, the signal that leaves the cell is the same current, carried by the same gustducin, the same calcium, the same channel. Twenty-five front doors; one hallway. The convergence is a mercy, and we build on it.

Specialists and generalists

The family divides, roughly, into specialists and generalists, and the division is written into the shape of that interior pocket. At the specialist end sit receptors with a smaller, more rigid cavity that admits only a close-matching key. Our anchor, TAS2R38, is one of these — moderately narrow, tuned chiefly to the thiourea edge of glucosinolate chemistry, the N–C=S motif of the brassicas and of the classic test compounds PTC and PROP. Nearby, TAS2R16 answers to a different structural family entirely, the β-glucopyranosides, the bitterness of salicin and its kin. At the other end sprawl the generalists — TAS2R10, TAS2R14, and TAS2R46 chief among them — receptors with a larger, more accommodating pocket that makes a looser grip on many differently shaped molecules. These broad receptors are the wide mesh of the net, the guards that catch a poison no narrow sentinel was ever shaped for. Between the extremes, the rest of the family fills the middle, each with its own range.

RECEPTOR	TUNING	READS, AMONG OTHERS
TAS2R38	narrow–moderate	PTC, PROP, goitrin (the thioureas)
TAS2R16	moderate	salicin, β-glucopyranosides
TAS2R1	moderate	iso-alpha-acids (hops), bitter peptides
TAS2R43 / TAS2R44	moderate	caffeine, the bitter off-taste of some sweeteners
TAS2R10	broad	strychnine, caffeine, and many more
TAS2R14	broad generalist	dozens of unrelated plant bitters
TAS2R46	broad generalist	strychnine, denatonium, absinthin

The specialist–generalist split is not only a fact about chemistry; it is a fact about *loudness*, and loudness is what the Index measures. A narrow receptor meeting its matched ligand tends to fire at a low

concentration — a small pocket gripping tightly means half the receptors are already answering when very little of the molecule is present. A broad receptor, holding many keys loosely, typically needs more of any one of them before it reaches the same half-answer. That threshold concentration has a name, and it is the hinge on which the whole number turns.

The arithmetic of a single sentinel

The threshold is the **EC_{50}** — the concentration of a stimulus at which it drives its receptor to half of full response. A low EC_{50} means a potent, loud stimulus (little is needed to reach half-maximal firing); a high EC_{50} means a quiet one. The full relationship between how much molecule is present and how hard the receptor answers follows a dose–response curve, which for a simple one-site receptor takes the Hill form:

$$R / R_{max} = [L]^n / (EC_{50}^n + [L]^n)$$

Here R is the response, R_max the saturating ceiling, [L] the ligand concentration, and n the Hill coefficient that sets the steepness of the rise. Set n = 1 and let x stand for the ratio [L] / EC_{50}, and the curve simplifies to the clean, memorable **R / R_max = x / (1 + x)** — a response that is half-maximal exactly when x = 1, that is, when the ligand sits right at its EC_{50}.

The Index converts that threshold into a score by inverting it against a fixed reference and scaling by a thousand:

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

Three things fall straight out of that single line, and each is worth doing by hand once. First, the reference scores itself: put the reference in for the candidate and IFI = 1,000 × EC_{50}(ref) / EC_{50}(ref) = **1,000**, exactly. That is why the number 1,000 recurs through every axis — it is not a measurement but a definition, the zero-point of the ruler, a stated convention like how the Scoville scale is defined by dilution-to-threshold. Each reference constant defines IFI = 1,000 on that axis by convention, and each axis is calibrated to its own independently defined reference. Lay the eight reference points side by side and the scorecard reads 1,000 across the board, by construction:

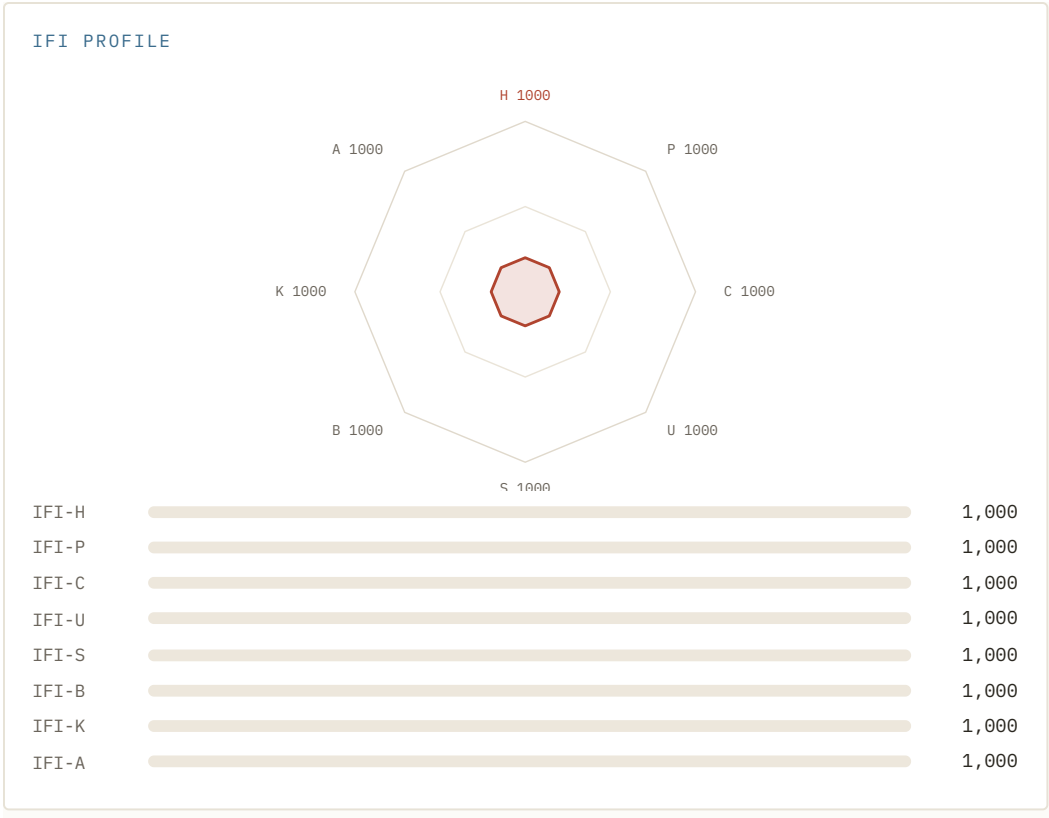

Second, the scale is proportional and legible. A candidate twice as potent as the reference — one that reaches half-activation at half the concentration, EC_{50} = EC_{50}(ref) / 2 — scores 1,000 × EC_{50}(ref) / (EC_{50}(ref)/2) = **2,000**. One a tenth as potent, needing ten times the concentration, scores **100**. Loudness and score move together, and a profile can be read at a glance. A bitterness that registers at **B 2,200** on a finished cultivar's design profile, for instance, is announcing that its bitter compounds are designed to reach half-activation at 1,000 / 2,200 = 0.4545 of the reference's concentration — roughly 2.2 times louder than the reference voice, a bitterness built deliberately into one note of a larger design.

Third — and this is the quiet discipline of the thing — the number carries its own uncertainty, and we carry it openly. An EC_{50} is not a fact handed down; it is a fit to a dose–response curve measured in cells, and every measurement has spread. Across our constructs we hold batch-to-batch variation to a coefficient of variation of **≤ 7%**. Because the IFI is a ratio of two measured EC_{50} values, their uncertainties combine in quadrature:

> $CV(IFI) = \sqrt{(CV(reference)^2 + CV(candidate)^2)}$

If both the reference and the candidate carried a 7% CV, the score would inherit $\sqrt{(0.07^2 + 0.07^2)} = \sqrt{0.0098} \approx$ **9.9%**, call it ten percent. This is precisely why we characterize the reference not once but exhaustively, driving its own CV down toward 2%: do that and the same candidate's score tightens to $\sqrt{(0.02^2 + 0.07^2)} \approx$ **7.3%**. The reference earns its authority by being the most-measured thing in the system.

One more conversion closes the loop between the receptor and the cup, because an EC_{50} is reported in molar concentration and a food is not. For a ligand of molar mass M grams per mole, an EC_{50} of E millimolar corresponds to a mass basis of E × M milligrams per liter. Caffeine's molar mass is 194.19 g/mol; a receptor whose half-point for it fell near 1 mM would be answering at roughly 1 × 194.19 ≈ 194 mg/L

— about 194 parts per million, a figure that lives in the same range as a strong brewed cup. The arithmetic walks all the way from a channel opening in a membrane to a number you could, in principle, taste.

Reading a food as a chord

A single compound against a single receptor is the clean case. A real food is never that. A real food is a **chord** — several molecules at once, each landing on its own subset of the twenty-five, and the bitterness that reaches the brain is the summed, blended output of whichever sentinels got touched.



Coffee — a whole handful of the generalists struck at once; chlorogenic lactones, phenylindanes, and caffeine each light up their own subset of the choir.

A glucosinolate lands most sharply on TAS2R38, nearer a single struck string. A cup of coffee strikes a whole handful of generalists at once — its chlorogenic acid lactones, its roast-born phenylindanes, and its caffeine each lighting a different subset, which is why coffee's bitterness feels broad and layered rather than pointed. A hop's bitterness runs largely through TAS2R1 and TAS2R14, the iso-alpha-acids finding those two doors. And denatonium — the most bitter substance known, the stuff painted onto things to keep a child from swallowing them — is bitter *precisely* because it strikes many receptors at once, a full chord where a glucosinolate plays nearer a single note. To score a food, then, we do not score one receptor; we take the net, the integrated loudness across the choir, and that integrated quantity is what a trained human panel reports and what the axis is built to track. The committee's single verdict, turned into a single number.

Why we lower the molecule, not the door

Here the convergence we traced — twenty-five doors, one hallway — pays off as method, and it is the most important idea in the book. There are two ways to make a bitterness quieter. You can block the receptor, jamming the door so the molecule arrives and finds no answer. Or you can lower the molecule, so there is less arriving at every door it would have struck. We take the second road, always, and the choir is the reason. If we blocked a receptor we would have to know the full roster a given bitterness plays — which of the twenty-five, and in what proportion — and jam each one, chasing a moving target through a parliament. But if we lower the **compound** itself — fewer glucosinolates, fewer phenylindane precursors, a rebalanced set of iso-alpha-acids — the reduction reaches *whichever* members

of the choir that compound was striking, all at once, without our ever needing the full list. Lower the molecule and every string it was playing goes quiet together.



Reading half-activation at the bench — an EC_{50} is never handed down; it is fit from a dilution series, each point a measured response carrying its own spread.

The dose–response curve tells us exactly how much quieter, and it tells us honestly, including where the method has to fight. Return to R / R_max = x / (1 + x) and watch what a fixed cut in the compound does from different starting points. A strongly bitter food sits high on the saturating shoulder of the curve, where the response has nearly flattened; a mildly bitter one sits low, where the response still climbs almost linearly with concentration. Halve the compound in each, and the *perceptual* payoff is wildly different — because perception follows the curve, not the concentration.

Food's bitter load, [L] / EC_{50}	Compound cut	New load, [L] / EC_{50}	Activation R / R_max (before → after)	Perceived attenuation
4.00	50%	2.00	0.80 → 0.67	17%
2.00	50%	1.00	0.67 → 0.50	25%
1.00	50%	0.50	0.50 → 0.33	33%
0.50	50%	0.25	0.33 → 0.20	40%
1.00	75%	0.25	0.50 → 0.20	60%

Read the first four rows together: the *same* fifty-percent cut in the molecule buys seventeen percent of relief on a fiercely bitter food and forty percent on a gentle one, purely because of where each sits on the curve. This is why the bitterness attenuation we quote lands in a **20–65%** band rather than at some single heroic figure — the band is the curve, not a slogan. The floor near twenty percent belongs to the foods that sit deep on the saturating shoulder, where even a large compound cut yields a modest felt change; reaching toward the sixty-five-percent ceiling takes both a food lower on the curve and a deeper cut, as the last row shows — three-quarters of the compound gone to buy sixty percent of the bitterness.

Two honest notes keep this from sounding like a free lunch. First, the simple n = 1 picture bends when the real Hill coefficient runs steeper than one, as bitter dose–response often does: a steeper n sharpens the curve near the EC_{50}, so a cut that drags a food across its own half-point pays off more than the table suggests, while a cut that leaves it stranded high on the shoulder pays off even less. We fit n for each system rather than assume it. Second, and this is a line we hold for reasons past chemistry: we will not strip a plant's bitter compounds to

the floor to chase a larger number. Those glucosinolates and polyphenols are the plant's own defenses and, in a body, among its antioxidants and quiet nutrition. The ceiling on attenuation is set not only by the curve but by the plant's health and the food's worth — we lower the note; we do not gut the instrument. The number is bounded by the thing we are trying to nourish.

The host roster, and what each carries

Because the intervention is molecular and not receptor-by-receptor, the method ports across organisms as cleanly as it ports across the choir. A growing roster of hosts now carries a tuned bitter axis, each scored on the same ruler, each held to the same batch-to-batch reproducibility, each lowered — never jammed — toward the fullest flavor and the fullest nourishment it can hold.

HOST	CONSTRUCTS	AXES	DIFFICULTY
Grape	1,610	B / A / M	Hard
Broccoli	380	B / M	Moderate
Tea	280	U / B / A	Hard
Broccoli-Kale	280	B	Hard
Orange	120	B / C	Hard
Potato	80	B	Easy
Hops	80	B	Moderate
Coffee	40	B	Hard
Cabbage	40	B	Moderate
Kale	40	B	Moderate
Radish	40	B	Hard
Rapeseed	40	B	Hard

The bitter axis also rarely travels alone, which is the other reason to score the whole choir honestly: tea carries umami and astringency beside its bitterness, grape carries astringency, and a fully designed cultivar can carry all eight axes at once, bitterness among them as one deliberate voice. That is how a bitterness reads inside a larger composition — one line of a full chord, tuned on purpose:

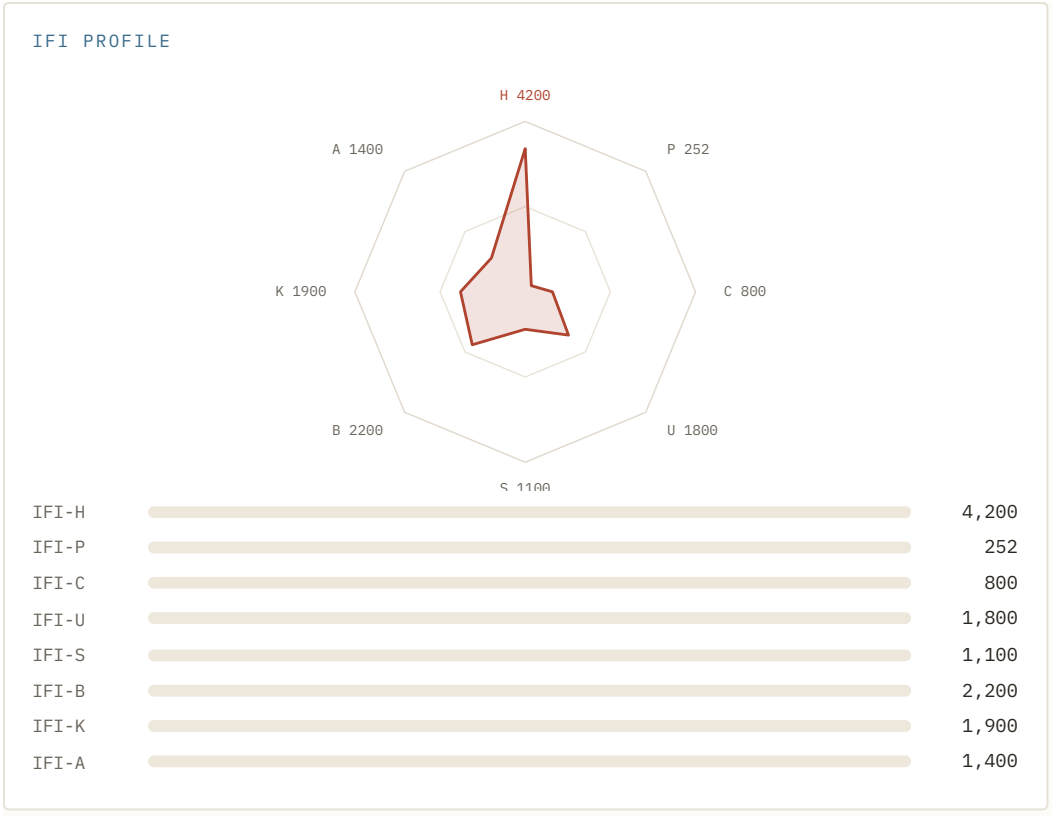

On a profile like that, the **B 2,200** is not a flaw to be attenuated but a note placed with intent — bitterness as structure, in a design that also carries sweetness at S 1,100 and kokumi at K 1,900, each axis calibrated to its own defined reference convention rather than a single shared scale, the way a cook sets a dark edge against a round one. Whether the task is to quiet a bitterness or to build one, the honesty is the same: we report the EC_{50} with its spread, we propagate that spread into the score, and we let the number say exactly how loud, and how sure.

The personal fingerprint, and the ruler we keep

The choir is not only large; it is individually tuned, and this is the last thing a careful measure has to confront. Nearly every one of the twenty-five genes carries common variants in the population — not just the famous TAS2R38 switch that sorts the world into tasters and non-tasters of PTC, but quieter variants up and down the family. So no two people hold quite the same set of sentinels at quite the same sensitivities. One person is sharply alive to a hop's bite and dull to a coffee's; another, the reverse — a personal bitter fingerprint, inherited, as particular as a face. It is a complication worth naming plainly rather than hiding: a single IFI-B is written for one defined reference tongue, and the honest question of whose tongue that should be is one we keep open and answer out loud, not bury in a footnote. A standard that pretended everyone tasted alike would be a standard built on a fiction.

This is exactly why the Index fixes one reference voice and declares it. The ruler is a single receptor, read as a single defined threshold, calibrated exhaustively so its own uncertainty is small — one voice, deliberately chosen, so that every measurement across every host resolves against the same note. And then, having measured against the one, we confirm against the many: a human sensory panel, whose tongues integrate the whole choir the way an eater's will, the twenty-five summed back into the single verdict they were always meant to give. The cell assay tells us which strings a molecule plays and how loudly; the panel tells us what the chord sounds like to a person sitting down to eat.

For the family doing this work, the discipline and the reverence are one thing, not two. The bitter receptors turn up beyond the tongue as well — in the lining of the gut and the airway, where they appear to serve older chemical-sensing roles — and that is a biology we note with interest and do not trade on; nothing here turns on it. What we build on is only the honest, measurable thing: twenty-five guards, one hallway, one summed feeling, lowered gently by lowering the molecule, and checked against a real mouth. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The ruler is one voice. The music is many. We measure the one, and we listen, always, to the many.

Chapter 4 · From Tongue to Thalamus — *the cascade behind a wince.*

gustducin · PLCβ2 · IP3R3 · TRPM5 · CALHM1/3

A receptor catching a molecule is only the first instant of a wince. Between the glucosinolate settling into the TAS2R38 pocket and a child pushing the plate away lies a fast, specific, and surprisingly recently understood cascade — and because our family designs its bitterness attenuation to act at the very top of that chain, at the receptor and before it, the whole chain is worth tracing end to end. Where you step into a signalling pathway decides what you get back out of it. Step in late, by masking, and you spend a second taste fighting the full, amplified output of machinery that has already done its work. Step in at the source, as our plants do, and the machinery below simply has less to carry. Everything in the pages that follow is an argument for the second place to stand, and we think the argument is only worth making if it is made in the biochemistry, the arithmetic, and the honest shape of a dose–response curve.

The cell that speaks without a synapse

The bitter signal is carried by a particular kind of cell. The tongue's taste buds hold several cell types, and bitterness — along with sweet and umami — is the business of the **type II** receptor cell. These cells are unusual, and the unusualness is the whole reason the cascade takes the shape it does: they have no classical synapse, no little vesicles of neurotransmitter waiting at a junction with the nerve. They must talk to the nervous system another way, and the chain of molecules inside them is built, step by patient step, toward that final problem of how to speak without a mouth.

When an agonist activates TAS2R38, the receptor does what a G-protein-coupled receptor does — it trades its resting shape for an active one and engages a G protein on the inside face of the membrane. For the bitter receptors that G protein is built around **α-gustducin**, a taste-specific cousin of the α-transducin that lets the eye's rod cells read light. Gustducin is the signature molecule of bitter, sweet, and umami transduction; an animal without it is substantially taste-blind to all three. It is worth pausing on that kinship. The tongue reading a toxin and the retina reading a single photon run nearly the same circuit — a receptor, a transducin-family G protein, an effector enzyme, a channel — because evolution, having solved the problem of turning a faint outside event into a cellular shout once, reused the solution. The eye amplifies light. The type II cell amplifies danger.

There is a second elegance here, and it belongs to the person doing the tasting. TAS2R38 is famously polymorphic: three coding changes — at residues 49, 262, and 296 — sort humanity into the **PAV** haplotype, which tastes the thiourea compounds vividly, and the **AVI** haplotype, which barely tastes them at all. The same broccoli is a punishment to one diner and a mild green to the next, and nothing in the vegetable has changed; only the pocket reading it has. We find this clarifying rather than discouraging. It tells us that bitterness is not

a fixed property of a food but a negotiation between a molecule and a receptor — and a negotiation is a thing one can design toward, from either side. We cannot edit the diner's receptor, and would not. We can edit how loudly the plant speaks into it.



A taste-cell preparation on the microscope stage — the calcium flash that follows a bitter agonist is the physical event the IFI-B axis is calibrated against.

Down the chain, molecule by molecule

Activation splits the G protein. The α-gustducin subunit goes one way — tamping down the cell's resting cyclic-nucleotide tone — but the decisive messengers are the **βγ** pair that gustducin releases, specifically the Gβ3·Gγ13 dimer. That βγ pair switches on an enzyme, **phospholipase C β2 (PLCβ2)**, parked in the membrane. PLCβ2 is the hinge of the whole pathway and the first true amplifier in it. One activated enzyme does not cleave one lipid and stop; it turns over its substrate again and again for as long as it stays active, each cycle cleaving a membrane lipid, PIP_2, into two second messengers. The one that matters here is **inositol trisphosphate (IP_3)**, which diffuses off into the cell's interior carrying the message inward.

IP_3 finds its target on the endoplasmic reticulum — the type-3 IP_3 receptor, **IP_3R3**, the particular channel these cells rely on — and opens it, and a store of calcium the cell has kept sequestered floods out into the cytoplasm. This is the second amplification: a handful of IP_3 molecules binding a handful of channels releases thousands of calcium ions, because the gradient across the ER membrane is steep and the channel, once opened, is a wide gate. The wince is now a rising tide of Ca^{2+}.

That calcium does two things. It is the internal signal that *a bitter molecule is here,* and it is the key that opens the next gate: **TRPM5**, with its relative TRPM4, a calcium-activated channel in the cell membrane that, once Ca^{2+} rises, lets sodium pour in. Sodium entry depolarizes the cell — swings its voltage positive — and that voltage swing is what finally reaches the outside world. It triggers trains of action potentials and, with them, opens the last and most beautiful piece of the chain: **CALHM1**, teamed with **CALHM3**, a large voltage-gated channel that is, in effect, the type II cell's substitute for the synapse it does not have. When it opens, it releases **ATP** — adenosine triphosphate, the cell's energy currency, pressed into service here as the neurotransmitter the cell lacks a synapse to deliver. ATP floods out of the taste cell and onto the waiting nerve.

The nerve reads it through purinergic receptors — **P2X2 and P2X3** ion channels on the gustatory afferent fibers — and the message, now electrical, runs inward: to the nucleus of the solitary tract in the brainstem, up through the ventral posteromedial thalamus, and into the gustatory region of the insular cortex, where it finally becomes the thing a person actually experiences, the conscious taste of *bitter,* arriving a few hundred milliseconds after the molecule first touched the tongue. The whole relay is specific enough that a defect at almost any link blunts bitterness selectively, and fast enough to make a toddler recoil before a thought has formed. Laid out as a signal chain, with each stage's job and a modeled latency, it reads like a circuit diagram:

STAGE	MOLECULAR EVENT	ROLE IN THE CHAIN	MODELED LATENCY
1 · Reception	agonist → TAS2R38 active conformation	switch (on/off)	<1–5 ms
2 · Transduction	α-gustducin releases Gβ3·Gγ13	relay, first branch	1–10 ms
3 · Enzyme	PLCβ2 cleaves PIP_2 → IP_3 + DAG	catalytic amplifier	5–30 ms
4 · Store release	IP_3 → IP_3R3 → Ca^{2+} efflux	second amplifier	10–50 ms
5 · Depolarization	Ca^{2+} → TRPM5 / TRPM4 → Na^+ influx	voltage conversion	10–40 ms
6 · Release	CALHM1·CALHM3 opens → ATP out	synapse substitute	1–10 ms
7 · Afferent	ATP → P2X2 / P2X3 on gustatory fiber	ignites the spike train	1–5 ms
8 · Central relay	NST → thalamus (VPM) → insular cortex	perception	100–300 ms

Modeled order-of-magnitude budget consistent with the known physiology; the central relay dominates the clock, which is why the tongue's chemistry feels instantaneous and the brain's accounting does not. Summed, the stages fall between roughly 130 and 450 milliseconds — the "few hundred" of lived experience, most of it spent not in the mouth but on the way to the cortex.

The arithmetic of loudness

The Index exists to put a single honest number on where a food sits before any of this fires. For every axis, including the bitter axis, that number is anchored to the half-maximal activation of the receptor — not how much of a compound a food holds, but how *loudly* that compound speaks into the channel that reads it:

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

EC_{50} is the concentration — or, for a whole food, the amount — at which a stimulus drives its receptor to half of full response. A food rich in a potent bitter principle reaches half-activation on very little of itself, so its EC_{50} is small and, inverted against a fixed reference, its score is large. The reference is pinned at exactly 1,000 by construction: substitute EC_{50}(reference) into the formula and the two terms cancel. Everything louder than the reference scores above a thousand; everything gentler scores below. On the bitter axis the reference receptor is **TAS2R38**, and the scorecard normalizes each axis to a ceiling of 5,000.

The formula is not decorative; it inverts. Take the fullest-spectrum plant we have built, **Phoenix**, which registers on all eight axes at once and carries a designed bitter value of **IFI-B 2,200**. Read the formula backward and that single number tells us an EC_{50} ratio:

> $2{,}200 = 1{,}000 \times EC_{50}(\text{ref}) \div EC_{50}(\text{Phoenix}) \implies EC_{50}(\text{Phoenix}) = EC_{50}(\text{ref}) \times 1{,}000 / 2{,}200 = \mathbf{0.4545 \times EC_{50}(ref)}$

Phoenix's bitter principle reaches half-activation of TAS2R38 on 45.5% of the amount the reference needs — it is 2.2× louder on the bitter channel than the anchor. That bitterness is designed and wanted, a structural note woven through a plant engineered to sing on every axis; we quote it here because it is the one complete profile the Index publishes, and because a chapter about quieting bitterness ought to begin by admitting that bitterness, designed and placed, is sometimes exactly what a composition calls for.

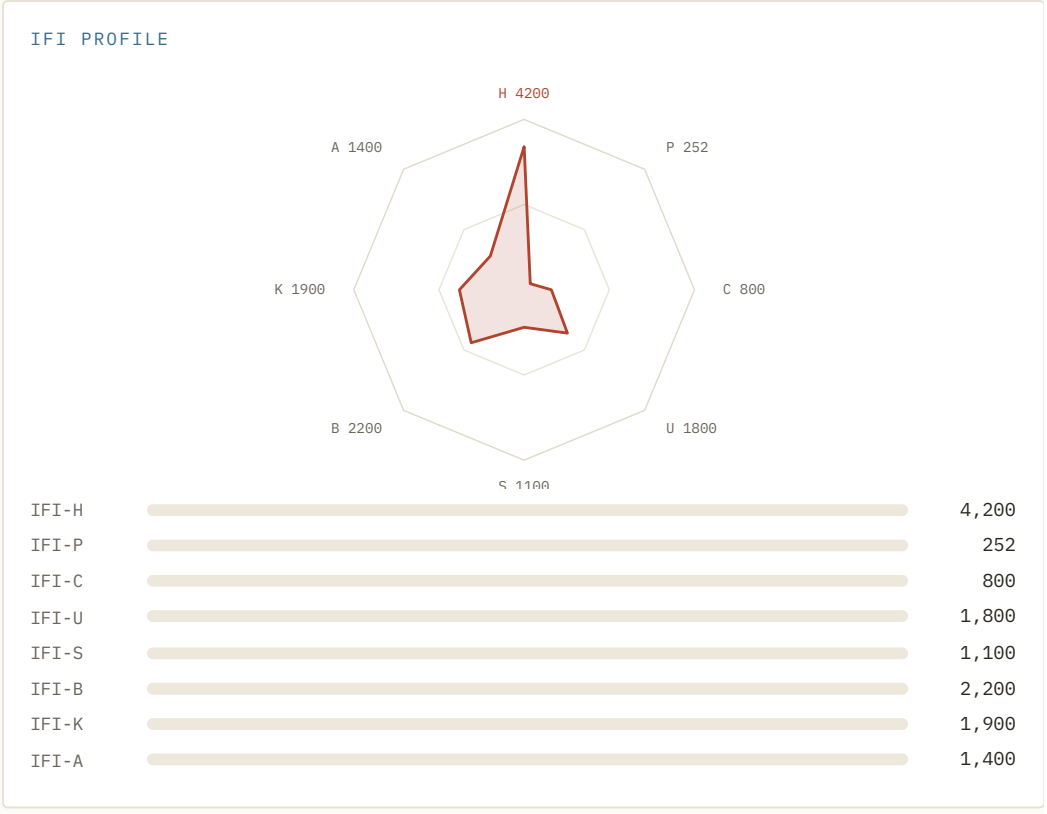

The units behave the way units should. An EC_{50} quoted in molarity converts cleanly to mass once a molecular weight is in hand — the bitter thiourea PROP, for instance, weighs about 170 g/mol, so an illustrative assay EC_{50} on the order of 10 µM is 10×10^{-6} mol/L × 170 g/mol ≈ **1.7 mg/L**. The ratio that makes the Index is dimensionless — milligrams over milligrams, micromoles over micromoles — so the score survives the translation from the chemist's bench to the kitchen scale unchanged. That invariance is the quiet reason a tomato, a cup of tea, and a chili can be read against one another on a single page.

A cascade is a multiplier

Here is the fact that decides our entire strategy, and it is a fact about arithmetic, not taste. The chain from receptor to calcium is not a relay of ones; it is a relay of *multiplications.* One active receptor engages many G proteins across its brief active lifetime. Each activated PLCβ2 cleaves substrate again and again, minting many IP_3 from one switch-on. Each opened IP_3R3 releases not one calcium ion but thousands. Stack those stages and the net chemical gain from a single binding event to the calcium transient is plausibly of order **10^3 to 10^4** messenger molecules — an estimate, order-of-magnitude, read off the known biochemistry rather than off any one measurement, but robust in its size if not its last digit.

That gain is a lever, and the question of all bitterness engineering is which end of the lever you grab. A masker or a blocker works *downstream* of the multiplier: it must raise a competing signal loud enough to dominate after the cascade has already amplified the alarm, or sit on the receptor and argue with the agonist molecule by molecule. Our intervention works *upstream* of the multiplier, in the plant, before the agonist is synthesized in the quantity that would trip the sentinel. Each precursor molecule we keep the plant from making is of order 10^3–10^4 downstream messengers that are never produced — calcium that never rises, ATP that is never spent, a spike train that never leaves the tongue. The leverage of standing upstream over standing downstream is, to an order of magnitude, the gain of the cascade itself. We did not invent that advantage. We simply chose to stand where the multiplier was working for us instead of against us.



The controlled chamber where the plant itself is tuned — our one place to intervene is before the molecule is ever made, upstream of the entire cascade.

Dose, response, and the shape of a sigmoid

To intervene honestly you have to know the shape of the curve you are bending, and the receptor's response to its agonist is not a line but a sigmoid. Model it, as the field does, with the Hill relation: if x is the agonist concentration expressed in units of its own EC_{50}, the fractional receptor response is

$$\theta = x^n \div (1 + x^n)$$

with n the Hill coefficient, the steepness of cooperation. Taking the simplest non-cooperative case, $n = 1$, so $\theta = x \div (1 + x)$. At the EC_{50} itself, $x = 1$ and $\theta = 0.5$ — half-activation, by definition. Far below it the

response climbs almost linearly with concentration; far above it the response flattens onto a saturated shoulder where adding more agonist barely moves the needle.

The local sensitivity of that curve has a clean closed form. Differentiate and the *elasticity* — the fractional change in response per fractional change in agonist — is $E = n(1 - \theta)$. At threshold, where $\theta = 0.5$, a 1% change in agonist content moves the signal by about 0.5% for a simple receptor, and by more for a cooperative one. Deep on the saturated shoulder, where $(1 - \theta)$ is small, that same 1% barely registers. This is the mathematical face of a thing every cook already knows in the hand: you can make a mild thing noticeably milder by trimming a little, and you cannot do the same to something already overwhelming. The curve is steepest, and the receptor most discriminating, precisely at its own half-point.

That shape is why our published attenuation is a band — **20% to 65%** — and not a single proud number. The honest reduction a diner perceives depends entirely on where the food sat on the sigmoid before we touched it. Applying one representative construct-level cut to foods at different starting points makes the band fall out of the arithmetic:

FOOD CLASS	X_0 (START)	θ_0	CONTENT CUT	X_1	θ_1	PERCEIVED Δ
Saturated (heavy roast, dense Brassica)	4.00	0.800	50%	2.00	0.667	**17%**
At threshold	1.00	0.500	50%	0.50	0.333	**33%**
At threshold, deeper edit	1.00	0.500	80%	0.20	0.167	**67%**
Delicate (light cup, young leaf)	0.30	0.231	50%	0.15	0.130	**44%**
Delicate, deeper edit	0.30	0.231	75%	0.075	0.070	**70%**

Read the first row and the last together and the whole design problem is visible at once. The *same* fifty-percent cut buys a barely-perceptible 17% on a food marooned on the saturated shoulder and a generous 44% on a delicate one; push the edit deeper and a food at threshold drops two-thirds. The published 20–65% envelope is the practical interior we commit to across real foods — it trims the punishing low end, where a plant sits so far up the shoulder that neither we nor a masker can do much, and it trims the extreme high end, which asks for edits deeper than we will make without changing the plant's character. Between those honest edges lies the range we will put our name to.

One subtlety keeps the two columns of that table from being confused. The Index registers *potency* — the loudness a food can reach — and potency scales linearly with the agonist content by the very construction of the formula: halve the content and you halve the IFI-B, because a food with less agonist needs more of itself to reach half-activation. The *perceived* reduction at a fixed serving is the

sigmoid quantity in the table's final column, and the two coincide only in the dilute regime near the foot of the curve. We report both, and we never let the friendlier of the two masquerade as the other.

Where we actually cut

The receptor that anchors this axis has a favorite kind of molecule, and knowing it tells us exactly where in the plant to work. TAS2R38 is the thiourea receptor — the one tuned to the N–C=S motif that appears in the laboratory compounds PTC and PROP and, in the garden, in **goitrin**, the bitter principle released when a Brassica's glucosinolate **progoitrin** is broken open. Broccoli, kale, cabbage, radish, rapeseed, mustard: their green bite is, to a first approximation, this chemistry speaking into this receptor. So the upstream edit is specific. We do not blanket the plant; we modulate the branch of its own metabolism that loads the glucosinolate pool, lowering the agonist the fruit or leaf will present before a single molecule reaches a tongue.

Doing it there, rather than on the receptor or at the cortex, matters for a reason the sentinel's own tuning supplies. Bitterness is an alarm, and evolution set alarms to fire early — a poison detector that waited for certainty would be a poison detector that killed its owner. So the behavioral trip point, the occupancy at which a diner stops chewing and turns away, sits low on the curve, not at saturation. That is a gift to the upstream strategy: you do not have to abolish the compound to move a food from *rejected* to *loved,* only to bring the receptor's occupancy back down across the steep part of the sigmoid and under the trip point. Partial attenuation, honestly delivered, is enough — which is the whole reason a 20–65% band, rather than a total erasure, is the right thing to build.

Set the three ways to quiet a bitter note side by side and the case makes itself:

	UPSTREAM — IN THE PLANT *(OURS)*	COMPETITIVE BLOCKER	SWEET MASKER
Site of action	before synthesis, before the receptor	on TAS2R38	at the cortex, by distraction
Agonist still ingested	reduced at the source	full dose	full dose
Persists when the sweetness fades	yes — it is structural	n/a	no — it fades with the sugar
Effect on other axes	targeted to the bitter branch	can blanket neighboring bitter notes	adds an entire sweet axis
Relative to the cascade gain	acts *before* the multiplier	argues *at* the receptor	fights *after* the full gain
Added load on the body	none	a new molecule to ingest	an added sugar load

A masker leaves the bitter compound entirely present and merely shouts over it, which is why a bitter coffee needs so much sugar and why the sweetness has to keep coming; the moment it fades, the bitterness is still there, untouched, doing its work on the receptor. A pharmacological blocker can quiet the receptor but blankets it and can dull notes you wanted to keep. We do neither. Our intervention is upstream of the entire cascade, in the plant, and the quiet it makes is real, and it does not fade when the sugar does.

Reproducibility is what makes a number mean anything

A bitter score is only worth printing if the next harvest prints the same one. This is where the discipline that elsewhere turned a ragged wild-type spread into a reproducible line earns its keep on the bitter axis: we hold batch-to-batch variation to a coefficient of variation of **≤7%**, against the **42%** that unselected material will happily wander across. The difference is not cosmetic, and it is not a matter of pride. It is the difference between a claim and a rumor, and the arithmetic says so exactly.

The IFI-B score is a ratio of two EC_{50} quantities, and for a ratio of independent quantities the relative variances add in quadrature:

> $CV(IFI)^2 = CV(reference)^2 + CV(candidate)^2$

Hold the reference tightly, at a CV near 2%, and read a candidate at our spec of 7%, and the score inherits

> $CV(IFI) = \sqrt{(0.02^2 + 0.07^2)} = \sqrt{(0.0004 + 0.0049)} = \sqrt{0.0053} \approx \mathbf{7.3\%}$

Read instead a wild candidate at 42% and the same formula returns $\sqrt{(0.02^2 + 0.42^2)} \approx$ **42%** — the measurement is as uncertain as the plant, and no amount of careful design survives it. Put the two regimes to the one test that matters, resolving a real attenuation. Suppose a module is meant to pull a food from IFI-B 1,000 to IFI-B 700, a 30% drop. At a 7% CV the two carry standard deviations of 70 and 49; the gap of 300 against a pooled spread of $\sqrt{(70^2 + 49^2)} \approx 85$ is a z of about 3.5 — unmistakable, the before and after cleanly separated. At a 42% CV the same two scores carry spreads of 420 and 294, a pooled 513, and a z of 0.6 — the two clouds overlap so thoroughly that the attenuation vanishes into noise. The honest bitterness number and the ≤7% CV are therefore the same achievement stated twice: you cannot report an attenuation you cannot reproduce, and reproducibility near the steepest part of the receptor's curve, where small errors in content become larger errors in signal, is the harder and the more necessary half of the work.



The sequencer bench — the repeatable side of a sensory claim; a bitter score is only worth printing when the next harvest prints the same one.

Honest limits

The cascade is clean, the formula inverts, and still we owe the reader the edges of what this axis can and cannot do. Three honest limits sit at the border of the work.

First, TAS2R38 is one receptor among roughly two dozen human bitter receptors, and total perceived bitterness is a sum across all of them. The IFI-B axis is anchored on TAS2R38 because it is the dominant, best-characterized channel for the glucosinolate and thiourea bitterness of the Brassica family — the cleanest and largest target in our host list. But a cup of coffee is a different animal: its bite is carried by caffeine and chlorogenic-acid lactones and phenylindanes speaking into *other* members of the family, receptors TAS2R38 never answers to. Our module, tuned to the glucosinolate branch, addresses the fraction of a food's bitterness that runs through the pathway we engineered; residual bitterness carried by other compounds into other receptors remains, and we say so plainly. Across our bitter-axis hosts the fit varies, and the honest reader can see it in the list: Grape at 1,610 constructs and Broccoli at 380 and the Broccoli-Kale line at 280 and the leafy Brassicas — Cabbage, Kale, Radish, Rapeseed, each at 40 or more — are glucosinolate-forward and sit squarely in the TAS2R38 story; Coffee at 40, Hops at 80, Tea at 280, Orange at 120, Potato at 80 are multi-receptor foods where our axis attenuates a real but partial share of the whole.

Second, the saturated shoulder is a true ceiling. A food so far up the sigmoid that it sits near full occupancy cannot be made much gentler by trimming agonist, because the curve has gone flat under it — the 17% at the top of the dose–response table is not a failure of the edit but a property of the mathematics. Beyond a point, the only honest lever is a different plant, not a deeper cut, and we would rather change the character of the cultivar than claim a reduction the receptor will not grant.

Third, and most personal to our family, is the tension between the alarm and the nourishment. The glucosinolates that make a Brassica bitter are also, in part, the very phytochemistry that makes the plant worth eating. We will not pretend a dial is a scalpel: lowering the compound lowers both the bite and, where they are the same molecule, some of the chemistry we admire. So our aim is a tuning, never a stripping — to bring the receptor-loud fraction down under the sentinel's trip point while preserving as much of the plant's own gift as the chemistry allows, and to be honest that this is a dial and not a switch. We began breeding because we wanted people to eat the whole green thing and love it, not to eat a hollowed-out version of it.

That is finally what the cascade teaches, read all the way from the glucosinolate to the insula. The wince is not an enemy; it is a billion-year-old courtesy, a warning the body pays forward. We do not want to silence it, and we could not if we tried — the chain is too specific and too fast. We want only to give it less to sound about, at the one place a plant can be changed before it is ever picked, so that a food grown to nourish can also be a food a person is glad to meet. The quiet we make is upstream, and structural, and real. It does not arrive with the sugar, and it does not leave with it. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.



PART II OF VI

The Mathematics of Less

IFI-B · the arithmetic of attenuation · the twenty-to-sixty-five band · the molecules of bitterness

Chapter 5 · The Axis We Read Backward — *IFI-B and the arithmetic of attenuation.*

the formula · the EC_{50} fold-shift · the 20–65% band · the 5,000 cap

Every number in this library rests on one honest line, and this axis uses the same one as the other seven:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of its full response — the physical measure of how *loudly* a molecule speaks. A molecule that reaches half-response at a low concentration is loud; one that needs a flood to get there is quiet. Invert the candidate's EC_{50} against a fixed reference, scale by a thousand, and loudness becomes a score you can print, cite, and reproduce months later. On the seven additive axes, a smaller EC_{50} means a louder voice and a higher score, and the whole design task is to push the number up.

Hold the word *loudness* close, because it is the thing that keeps this axis honest. We are not measuring how much of a bitter compound a food contains — grams per gram, parts per million, the brewer's bitterness units. We are measuring how hard that compound leans on the receptor that reports it. Two foods can carry the same milligrams of two different bitter molecules and read as wildly different bitternesses, because the two molecules sit at different EC_{50}s on the same receptor: one is a shout, the other a murmur, at identical concentration. Amount is chemistry. Loudness is chemistry *filtered through a receptor* — and the tongue only ever knew the second one. The Index measures what the tongue measures.

Why the number still climbs with the bite

Here is the one thing to hold onto about IFI-B, and we make it plainly because it is easy to get backward: **the number still means loudness.** A higher IFI-B is a *more* bitter plant, not a less bitter one — exactly as on every other axis, where higher means more. We do not invert the scale to flatter the goal. It would be tempting: a book of this kind might like a number that climbs as the bitterness falls, so that "better" and "bigger" point the same way. We decline, and the reason is the whole discipline of the Index. 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 a full-spectrum fruit's bitterness of 2,200 would mean the opposite of a tomato's. Consistency across the standard is worth more than a convenient direction. So on this axis, uniquely, the design goal and the number point *opposite* ways: we are the only axis in the Index where success is usually a *smaller* reading. We raise seven numbers and lower one —

and we keep the lowered one in the same honest units as the rest, so that a bitterness can always be read against a sweetness or a heat on one page without a translator.

What we therefore report, and what this whole axis is organized around, 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 the bitter loudness we remove:

```
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 two definitions collapse into one another with no slack. Since IFI-B(engineered) ÷ IFI-B(native) = EC_{50}(native) ÷ EC_{50}(engineered) = 1 ÷ φ, the attenuation is simply r = 1 − 1 ÷ φ, and inverting that gives the clean working form:

```
φ = 1 ÷ (1 − r)
```

So a forty-five percent quieting is not a vague softening; it is the bitter molecule's threshold moved to **1.82×** its native value, measured in a dish, reproducible to a published tolerance. The table below is the Rosetta stone of the axis — attenuation, the fold-shift that produces it, and the fraction of the original loudness that remains:

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

Read the table across, not down. The left column is the promise we make to an eater; the middle column is the physical fact in a dish that keeps the promise; the right column is what survives. Nothing in the three can drift without the other two moving with it, which is exactly what a standard is for.

What TAS2R38 actually does

To lower a number honestly, you have to know what the number is counting. Bitterness begins at a bitter taste receptor, and the reference receptor we anchor IFI-B to is **TAS2R38** — a seven-pass G-protein-coupled receptor studded into the membrane of the Type II taste cells that cluster inside the fungiform, foliate, and circumvallate

papillae of the tongue. It is the most thoroughly mapped member of a family of roughly twenty-five human bitter receptors, and it is the one that reads the thiourea motif — the small **N–C=S** group — that hangs off phenylthiocarbamide, off propylthiouracil, and, as it happens, off the breakdown products of the vegetables a child learns to push to the side of the plate. We anchor to it the way a tuning orchestra anchors to the oboe: not because it is the only voice, but because it is the steadiest, best-characterized pitch to tune the rest against.

The cascade it sets off is worth walking in full, because every link in it is a place a flavor could in principle be quieted, and knowing the chain tells us which link we are allowed to touch:

1. A bitter agonist settles into the receptor's binding pocket, and the seven helices shift — a motion of a few ångströms that the cell feels as a command.
2. On the inside face, the receptor hands that command to a heterotrimeric G-protein built around **α-gustducin**; the trimer splits, freeing a $G\beta_3\gamma_{13}$ pair.
3. The freed βγ pair switches on **phospholipase C-β2**, an enzyme that sits waiting at the membrane.
4. PLC-β2 cleaves the membrane lipid PIP_2 into two messengers: diacylglycerol, and **inositol trisphosphate (IP_3)**.
5. IP_3 drifts to the endoplasmic reticulum and opens the **IP_3 receptor type 3**, and the cell's internal calcium store empties into the cytosol.
6. The rising calcium gates **TRPM5**, a cation channel, and the taste cell depolarizes.
7. Depolarization opens **CALHM1/CALHM3** channels, which release a pulse of **ATP** straight through the membrane.
8. That ATP lands on **P2X2/P2X3** receptors on the gustatory nerve ending, and the signal — *bitter, here, this loud* — finally leaves the tongue for the brain.

Eight steps, a billion times a second, and the whole of it is what a single IFI-B reading compresses into one printable integer. Two facts from that chain shape everything we do. The first: the receptor itself is polymorphic. The gene carries three well-known substitutions — **P49A, A262V, V296I** — whose combinations sort people into the taster haplotype (PAV) and the non-taster haplotype (AVI). A compound that is a wall of bitterness to one eater is a faint note to the one sitting beside them, and no engineering we do to a plant changes which eater is reading it. The Index reports loudness at a defined reference receptor so that two labs, two months apart, mean the same thing by "bitter"; the variation across a dinner table is real, and it is the honest reason we publish a shift rather than pretend to a single universal perception. The second fact is the one that disciplines the entire axis: **we never touch step one through step eight.** We do not edit the eater's receptor. We edit the plant — and the only thing a plant can change is how much agonist it presents to a receptor it will never meet until the fruit is in someone's mouth. Attenuation, for us, is always and only a reduction in agonist load, read out as a rightward shift of the *effective* EC_{50} of the food.

Where the bitterness is made

So the real engineering lives upstream of the tongue, in the plant's own chemistry, and it differs by family. The honest work of this axis is knowing, for each food, which molecule the receptor is actually reading and which enzyme in the plant decides how much of it exists.



Coffee fruit in cross-section — the bitterness a cup carries is not in the bean yet, but written into it, as chlorogenic acids that the roast will fold into bitter lactones and phenylindanes.

In **coffee**, the cup's bitterness is only partly the caffeine everyone blames — caffeine carries perhaps a tenth of it. The rest is built in the roaster: green-bean **chlorogenic acids**, which are esters stitched together along the phenylpropanoid branch (PAL → C4H → 4CL and onto the hydroxycinnamoyl transferases), break down under heat into **chlorogenic-acid lactones**, which dominate a medium roast's bitterness, and at darker roasts into **phenylindanes**, the harsh, lingering compounds that read as the "ash" at the back of a bitter cup. The lever is therefore not in the cup and not even in the bean as you buy it — it is in the living plant, in the flux through the quinate transferase that commits carbon to chlorogenic acid. Draw that pool down at the source and the roast has less raw material to fold into lactone and phenylindane, while the aromatics and the caffeine that make coffee *coffee* are left standing. The bitterness falls because the precursor the fire needed was never laid down.



Broccoli in cross-section — the bite is in the glucosinolates, and so is much of the nourishment; the whole art is lowering the first without spending the second.

In the **brassicas** — broccoli, kale, cabbage, their kin — the bitterness is the family's famous chemistry doing its job. Glucosinolates sit inertly in the intact tissue until a bite or a knife brings them into contact with **myrosinase**, the enzyme that hydrolyzes them into the pungent, sometimes acrid compounds that read as the plant's defense: allyl isothiocyanate from **sinigrin**, and the bitter, thiourea-bearing **goitrin** from **progoitrin**. That last one is the clean case for

TAS2R38, because goitrin's thione carries the very N–C=S motif the receptor was built to read. The temptation is to simply knock the glucosinolates down — and that is the trap, because the glucosinolates are also where much of the nourishment lives. **Glucoraphanin**, in particular, is the precursor of sulforaphane, the compound a head of broccoli is most prized for. To cut it away for the sake of a milder bite would be to lower the bitterness by throwing out the reason to eat the plant, and that is not a trade this family will make.

The better lever is to change the *profile*, not the total. The enzyme **AOP2** converts the methylsulfinyl glucosinolates toward the alkenyl forms — toward sinigrin and the progoitrin that becomes goitrin. Turn its flux down and the plant stops spending its glucoraphanin on bitterness; the sulforaphane precursor accumulates instead of being converted to the harsher, goitrogenic, TAS2R38-loud molecules. The result is a broccoli whose bitterness has moved rightward on the EC_{50} scale while its most valuable phytonutrient has, if anything, climbed. That is the whole ethos of this axis in one enzyme: not *less vegetable*, but *the same vegetable, speaking more gently*.

The other bitter shelves run the same logic with different chemistry. In **hops**, the bitterness is iso-α-acids — humulones isomerized to isohumulones in the boil — and the lever is the acylphloroglucinol pathway in the lupulin gland, where the balance of the harsher and softer α-acids is set. In **tea**, it is the galloylated catechins, the EGCG and ECG that are bitter and astringent at once, and the lever is the acyltransferase that decides how many of the catechins carry a gallate group at all. In **grape**, it is the flavan-3-ol tannins and their polymers, tuned along the proanthocyanidin branch. The table below gathers the shelf — the construct counts and the engineering difficulty are the library's own; the chemistry beside them is simply what the receptor is reading:

ORGANISM	CONSTRUCTS	DIFFICULTY	WHAT TAS2R38 (OR A BITTER SIBLING) READS	THE LIVING LEVER
Coffee	40	Hard	chlorogenic-acid lactones, phenylindanes, caffeine	lower chlorogenic-acid flux in the phenylpropanoid branch
Broccoli	380	Moderate	goitrin, sinigrin (glucosinolate-derived)	down-tune AOP2; keep glucoraphanin, drop the alkenyls
Kale	40	Moderate	sinigrin, progoitrin	the same aliphatic-glucosinolate lever
Cabbage	40	Moderate	sinigrin, glucobrassicin	tune the MYB28 master regulator of the pathway
Hops	80	Moderate	iso-α-acids (isohumulone)	shift the α-acid profile in the lupulin gland
Tea	280	Hard	galloylated catechins (EGCG, ECG)	reduce catechin galloylation
Grape	1,610	Hard	flavan-3-ol tannins, procyanidins	tune the proanthocyanidin branch
Potato	80	Easy	glycoalkaloids (solanine, chaconine)	cap glycoalkaloid synthesis — *and no further*

That last row carries a warning we will come back to, because a potato's bitterness is also a poison's bitterness, and the two cannot be separated by wishing.

Reading a shift three ways

A single attenuation number travels through three different frames before it becomes a sensation, and reading an IFI-B honestly means holding all three at once. The number we publish lives in the first frame. What the receptor does lives in the second. What the eater feels lives in the third.

The first frame is the Index itself, and it is the clean one — serving-independent, reference-anchored, the number that belongs on a scorecard. When we say a food's bitterness is attenuated by 50%, we mean its IFI-B has fallen to half of its native value and, equivalently, that the effective EC_{50} of its bitter principle has doubled: $\varphi = 1 \div (1 - 0.50) = 2.00\times$. This frame does not care how big a bite you take. It is a property of the plant, the way a melting point is a property of a solid.

The second frame is the receptor, and here the Hill equation does something generous. TAS2R38, like most bitter receptors, turns its agonist into a response along a sigmoid,

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

where [A] is the agonist concentration and **n**, the Hill coefficient, is the steepness — for the bitter receptors typically a little above one, around 1.3 to 1.5. Take a food that natively sits right at its EC_{50}, which most bitter foods do, since that is the concentration at which they register as solidly bitter in the first place. At $[A] = EC_{50}$, the response is exactly half-maximal — by definition, $(EC_{50})^n \div ((EC_{50})^n + (EC_{50})^n) = ½$. Now apply the 50% attenuation. The agonist in the mouth has not changed — the eater takes the same bite — but the *effective* EC_{50} has doubled to $2{\cdot}EC_{50}$, and the occupancy at that same bite becomes, with n = 1.4,

$$R \div R_{max} = (EC_{50})^{1.4} \div ((2{\cdot}EC_{50})^{1.4} + (EC_{50})^{1.4}) = 1 \div (2^{1.4} + 1) = 1 \div 3.64 = \mathbf{0.275.}$$

The receptor was firing at 50% of its ceiling; now it fires at 27.5%. A doubling of threshold has bought a drop of nearly half the actual receptor activation — more than the number alone promised — because a food sitting on the steep middle of its own dose-response curve is in exactly the place where a modest rightward shift of EC_{50} pays out most. The attenuation is most efficient precisely where the eater needs it most. This is why we design foods to sit near their EC_{50} before we quiet them, and why the same published attenuation feels like more relief on a solidly bitter food than on a faintly bitter one.

The third frame is perception, and it rounds the step gentler still. The nervous system reads intensity roughly as the logarithm of the signal — the old Weber–Fechner relation — so the receptor's drop from 50% to 27.5% does not arrive on the tongue as a brutal, linear "nearly half gone." It arrives compressed: the ratio the brain is taking the logarithm of is 50 ÷ 27.5 = 1.82, and a log of 1.82 is a soft, rounded step down, not a cliff. We are not trying to amputate a sensation. We are trying to move it from *too much* to *right*, and the logarithm is on our side — it turns a precise physical halving into the gentle, characterful softening that a good cup or a good head of broccoli actually wants.

Three frames, one promise. The scorecard prints the first. The dish delivers the second. The eater lives in the third. An IFI-B read honestly knows which frame it is standing in.

The band, and why both edges are drawn

Two lines bound the working band — a floor at twenty percent and a ceiling at sixty-five — and both 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: a quieting no one can reliably feel. We will not publish an attenuation an eater cannot perceive, and there is a measurement reason beneath the perceptual one, which the next section makes exact — a sub-twenty-percent shift is hard to tell, with confidence, from no shift at all. A number we cannot defend is a number we do not print.

The **ceiling at sixty-five percent** is a floor of a different kind — a floor under the flavor, and under the body. Push past it and three things fail at once. The food loses its spine: the structural bitterness that makes coffee taste like coffee and chocolate like chocolate thins into something flat and characterless, and a plant stripped of its bitterness does not taste pure, it tastes absent. The nourishment can go with it, because in the brassicas the very compounds that read as bitter are kin to the compounds that feed you, and only careful profile-shifting keeps the second when you lower the first; past a point, the knife cannot tell them apart. And the sentinel loses too much of its reach.

That last one is not a metaphor, and the potato is its plainest case. Bitterness is, before it is a flavor, a poison-detection system — the body's oldest chemical alarm, tuned over evolutionary time to the shapes of the compounds that have historically meant *do not eat this*. A potato's bitterness is its **glycoalkaloids**, solanine and chaconine, and those are not merely unpleasant; above roughly two hundred milligrams per kilogram they are genuinely toxic, which is exactly why they taste the way they do. A green, bitter potato is a potato telling the truth. A plant engineered to be entirely unable to taste bitter is a plant we have made less safe to be wrong about — a sentinel we have blindfolded. So we cap glycoalkaloid work hard, and we treat the sixty-five percent ceiling as a discipline rather than a wall we strain against. The best work on this axis usually lives well short of it: a thirty or forty-five percent quieting that lifts the grip off a vegetable while leaving its spine, its sulforaphane, and its warning intact. We lower bitterness the way you lower your voice to be understood, not the way you cut a wire.

What the measurement can and cannot resolve

Numbers are a form of respect for the reader, and respect means saying how sharp the number is. The library holds its batch-to-batch coefficient of variation to **≤7%** — the same tolerance that governs the heat work and everything else — and that tolerance propagates in a way worth doing out loud, because it decides which of the fine gradations in the Rosetta table are real distinctions and which are only prose.

The attenuation is a ratio, r = 1 − IFI-B(engineered) ÷ IFI-B(native), and the engineered and native readings each carry their own ≈7% scatter. For a ratio of two independent measurements, the relative uncertainties add in quadrature, so the quotient Q = IFI-B(engineered) ÷ IFI-B(native) carries a coefficient of variation of $\sqrt{(0.07^2 + 0.07^2)} = \sqrt{0.0098} \approx$ **9.9%**. Because r is just 1 − Q, its absolute standard deviation equals Q's: about 0.099 × (1 − r) on a single batch pair, tightened by the square root of the replicate count. Running the standard triplicate, the one-sigma uncertainty on a published attenuation comes out to:

TARGET ATTENUATION *R*	1-Σ ON *R* (TRIPLICATE)	DISTINGUISHABLE FROM…
20%	±4.6 pts	barely from zero — hence the floor
30%	±4.0 pts	cleanly from 20%
45%	±3.1 pts	cleanly from 30%
50%	±2.9 pts	cleanly from 38%
60%	±2.3 pts	from 50%, yes; from 62%, no
65%	±2.0 pts	from 55%, yes; from 62%, no

Two lessons fall out, and both are honest. The first: the twenty-percent floor is not only a floor of perception but a floor of measurement — a claimed attenuation of twenty percent carries roughly a ±4.6-point shadow on triplicate, so it is only a handful of sigma clear of claiming nothing at all, and below it the measurement cannot keep its promise even if the tongue could. The second is a confession about the table itself: the close gradations near the ceiling — the sixty, the sixty-two, the sixty-five — are separated by two- and three-point steps that sit right at the edge of a two-point standard deviation. They are honest *narrative* markers, the texture of a craft, but they are not claims that a sixty-two-percent broccoli is statistically resolvable from a sixty-percent one. Where we need a gradation to be a real, defensible distinction, we space it by at least a clear five or ten points, and we say plainly when a step is finer than the instrument. A standard that overstates its own precision is not a standard; it is a decoration.

The 5,000 cap and the scorecard

One number in this chapter's subtitle has not yet been earned: the **5,000 cap**. It is the per-axis normalization maximum — the value at which every axis's spoke on an IFI radar scorecard reaches its full radius. It exists because the eight axes do not naturally share a top. Left unbounded, one loud axis would blow the radar out of shape and crush the others into the center, and a scorecard that cannot be read at a glance has failed at the one thing a scorecard is for. Fixing the full radius at 5,000 gives all eight spokes a common ceiling, so a profile reads as a *shape* — balanced, or heat-forward, or quietly savory — rather than a tangle of incommensurable magnitudes.

The arithmetic of the cap is the same inverse relation as the rest of the axis, read against the three fixed points the standard actually pins down. From IFI-B = 1,000 × EC_{50}(reference) ÷ EC_{50}(candidate), the candidate's effective EC_{50} in units of the reference is simply 1,000 ÷ IFI-B, and its share of the scorecard radius is IFI-B ÷ 5,000:

IFI-B READING	WHAT IT IS	EC_{50}(CANDIDATE) ÷ EC_{50}(REFERENCE)	SHARE OF SCORECARD RADIUS
5,000	the normalization cap	0.20×	100% (full spoke)
2,200	a full-spectrum fruit's kept bitterness	0.455×	44%
1,000	the reference anchor, by definition	1.00×	20%
0.55 · B_0	a native bitterness B_0 after a 45% quieting	1.82 × native	0.55 × its native share

Read the middle row: a bitterness of 2,200 means a bitter principle whose EC_{50} is 0.455 times the reference's — less than half the concentration reaches half-response, so it is more than twice as loud as the anchor — and on the radar it fills forty-four percent of the bitter spoke. That is a deliberately, sturdily bitter plant, plotted so you can see the choice. And read the bottom row for the discipline of the thing: a 45% attenuation does not touch the cap or the anchor; it takes a plant's native bitterness B_0, whatever it measures, to $0.55 \cdot B_0$ — a precise, proportional step inward on exactly one spoke, every other axis untouched. The scorecard is where the arithmetic of this chapter becomes a picture.

Bitterness, specified upward

And "right" is never zero, which returns us to the axis's own name. The clearest proof sits in a cultivar from the other end of the library entirely — the one plant we have ever designed to register on all eight axes at once.



Our model predicts a fruit of approximately 2.95 million SHU for this full-spectrum cultivar in the controlled chamber — built, deliberately, to keep a bitter spine beneath its heat.

Phoenix — SS-0096, with our model predicting approximately 2.95 million SHU — carries a bitterness we built *in* and kept, because a fruit that loud needs a bitter spine beneath the heat or the burn reads as a shriek instead of a flavor. Bitterness, placed right, is structure: it gives the heat a floor to stand on and the sweetness something to resolve against. Its scorecard says so in eight lines:

IFI PROFILE

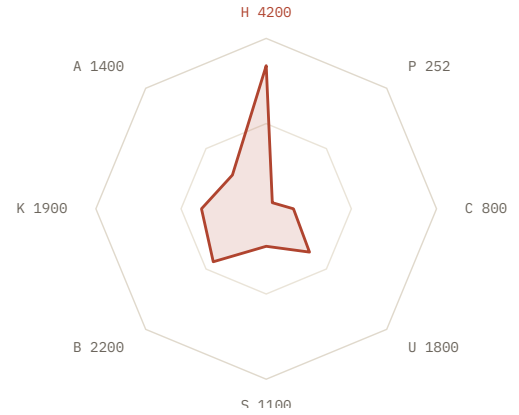

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

Its IFI-B of 2,200 is not a failure of this axis; it is this axis used in the other direction, bitterness specified *up* to exactly the level the composition needs — forty-four percent of the bitter spoke, held there on purpose. To make the cost of forgetting that visible, here is the one profile we keep in the drawer and will never print as a product: Phoenix with this chapter's own forty-five-percent quieting run against its bitter spoke alone, 0.55 × 2,200 = 1,210, every other axis untouched.

IFI PROFILE

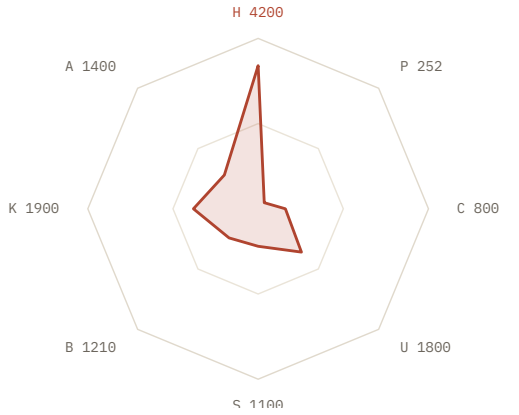

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	1,210
IFI-K	1,900
IFI-A	1,400

On paper it looks like an improvement — a lower bitterness, the direction this whole axis usually travels. In the mouth it would be a collapse: the heat left without a floor, the burn tipping from flavor into noise. That single changed line is the entire argument of the axis made concrete. Zero-Bitter is not a crusade against a taste. It is the ability to put bitterness precisely where a flavor wants it — which, for a head of broccoli, means far less than it has, with its sulforaphane and its warning left standing; and for a full-spectrum superhot, means a firm and deliberate 2,200. We learned to read this axis backward not

to erase a flavor, but to be able, finally, to *place* it — to lower a voice where a plant was shouting, and to let it keep speaking where the dish needs it heard.

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

Chapter 6 · The Anatomy of a Bitterness — *the molecules and their roads.*

glucosinolates · chlorogenic acids · flavan-3-ols · iso-alpha-acids · the two great pathways

Before any single plant, a map. Bitterness in the foods this book cares about is not one chemistry but a handful of distinct molecular families, each built on its own biosynthetic road, each read by its own chord of TAS2Rs — and knowing which family you are facing is most of knowing how to quiet it. A lever that lowers a coffee's bitterness does nothing to a broccoli's, because the two bitternesses are made by entirely different machinery, carried to the tongue by molecules that share neither a parent amino acid nor an enzyme nor a receptor. What unifies the work is not the molecules but the method: find the road, find the single valve on it that feeds the harsh end-product, and turn that valve precisely — by a measured fraction, in one tissue, at one point in the plant's day — while the rest of the plant's chemistry flows on untouched. We learned the discipline on heat, where a four-enzyme tail hangs off a highway that runs in every green thing alive. Bitterness asks the same question of a different set of roads, and answers it the same way: subtraction, exact, reverent of everything we do not touch.



The coffee bean in cross-section — most of a cup's harshness is laid down here, in the chlorogenic-acid chemistry of the green seed, long before a roaster ever touches it.

The two great roads

The bitternesses divide first by which great pathway builds them. One road is the **glucosinolate** route — the sulphur-and-nitrogen defense chemistry nearly unique to the brassicas and their relatives, built from amino acids (chiefly methionine) into thioglucoside bombs that shatter, on chewing, into the isothiocyanates and goitrin that TAS2R38 reads most sharply. The other road, carrying most of the rest of the plant kingdom's bitterness, is the **phenylpropanoid branch** — the universal five-enzyme highway **PAL → C4H → 4CL → HCT → COMT**

that runs in every vascular plant, together with the **flavonoid arm** that peels off it at the first acyl-CoA. From that single branch flow the chlorogenic acids of coffee, the flavan-3-ols of cacao and tea and the grape, the flavanone glycosides of citrus — a remarkable share of all food bitterness traces back to this one metabolic trunk, which is why a method that can read and tune the branch transfers from bean to leaf to fruit with the same logic and the same enzymes. A few bitternesses stand outside both: the hop's iso-alpha-acids, built by a lineage-specific polyketide route in the cone; the citrus limonoids, off the triterpenoid road; the steroidal glycoalkaloids that make a sun-greened potato bite. We name those honestly as exceptions rather than force them onto the map, and give each its own treatment where that organism has the floor.

Here is the shelf's bitterness, family by family, with the road that builds each and the chord that reads it:

FAMILY	EXAMPLE MOLECULES	FOUND IN	BUILT ON	READ MAINLY BY
Glucosinolate-derived	sinigrin → allyl-ITC; progoitrin → **goitrin**	broccoli, kale, cabbage, mustard, radish	glucosinolate pathway (Met-derived)	**TAS2R38**, TAS2R16
Chlorogenic-acid derived	CGA lactones; **phenylindanes** (dark roast)	coffee	phenylpropanoid branch	TAS2R46, TAS2R7, broad
Methylxanthine	**theobromine**; caffeine	cacao, coffee, tea	purine-alkaloid route	TAS2R7, R10, R14, R43, R46
Flavan-3-ol / proanthocyanidin	epicatechin, catechin, **procyanidins**; EGCG	cacao, tea, grape	flavonoid arm	TAS2R5, R39, broad + astringency
Flavanone glycoside	**naringin**; neohesperidin	citrus (orange, grapefruit)	flavonoid arm	TAS2R40, TAS2R43
Limonoid	**limonin** (delayed bitterness)	citrus	triterpenoid route	broad
Steroidal glycoalkaloid	α-solanine, α-chaconine	potato (greened)	cholesterol-derived SGA route	TAS2R, broad + irritant
Iso-alpha-acid	**isohumulone**, isocohumulone	hops, beer	cone polyketide (VPS) route	TAS2R1, TAS2R14, TAS2R40

Two facts in that table organize the whole rest of the book, and we will keep returning to them. The first is that **the reference receptor is not always the one doing the reading.** The second is that **bitterness and**

astringency are neighbors, and sometimes the same molecule wearing two coats. We hold both and come back to them, but first we have to walk the two great roads step by step, because the valve we turn sits on a specific enzyme, and you cannot turn a valve you cannot name.

Road one — the mustard bomb

The glucosinolate road is the most dramatic chemistry on the shelf, because the plant builds a weapon and holds it unloaded. Begin at **methionine**. A branched-chain aminotransferase, **BCAT4**, strips methionine's amino group to give a 2-oxo acid; the enzyme **MAM** (methylthioalkylmalate synthase) then runs that oxo acid around a three-step elongation cycle — condense, isomerize, decarboxylate — adding one methylene at a time, so the plant can dial the eventual side-chain length before the molecule is ever a glucosinolate at all. Transaminate back, and you have a chain-extended amino acid poised to enter the core pathway.

The core is a short, conserved cascade of oxidations. A cytochrome **CYP79** converts the amino acid to an **aldoxime**; a second cytochrome, **CYP83**, oxidizes the aldoxime to a reactive *aci*-nitro compound; a C–S lyase (**SUR1**) installs the sulphur that becomes the thioglucoside linkage; a glucosyltransferase (**UGT74**) hangs the glucose; and a **sulfotransferase (SOT)** caps the molecule with the sulfate that makes it a finished, stable, water-soluble glucosinolate. Two members of this family matter most to the tongue. **Sinigrin** (2-propenyl glucosinolate) is the mustard-and-horseradish principal. **Progoitrin** (2-hydroxy-3-butenyl glucosinolate) is the one that turns a boiled brassica sour-bitter in the back of the mouth.

None of this is yet bitter, and that is the whole trick. The finished glucosinolates are stored in dedicated sulphur-rich cells, held physically apart from the enzyme that would detonate them — **myrosinase** (a thioglucoside glucohydrolase, **TGG1/TGG2**), sequestered in separate *myrosin* cells. The leaf is a landmine with the fuse in a different room. Chew it, chop it, let a caterpillar bite it, and the two compartments mix. Myrosinase cleaves the thioglucoside bond, releasing glucose and an unstable aglycone that cannot hold itself together: it undergoes a **Lossen rearrangement** and spits out an **isothiocyanate** — allyl-ITC from sinigrin, the clean heat of fresh mustard. Progoitrin's aglycone, carrying a hydroxyl two carbons down, cyclizes instead into **goitrin** (5-vinyloxazolidine-2-thione), a stable ring that sits in the mouth and reads, to TAS2R38, as frank bitterness. The receptor is tuned to the **N–C=S** motif these products share — the same thiourea-like signature carried by the classic laboratory probes of this receptor — which is exactly why the brassica column of the family table all points at the same ruler.



Broccoli in cross-section — the glucosinolate road stores its weapon unloaded, the thioglucoside in one cell and the myrosinase that detonates it in another, mixed only when the tissue is cut.

Two valves on this road interest us, and they behave very differently. The first is **MAM**, far upstream: narrow its chain-elongation and you shift the *profile* of glucosinolates the plant can make, nudging it away from the progoitrin that becomes goitrin without emptying the plant's whole defense arsenal. The second is at the detonation step itself — the balance between isothiocyanate and the blander nitrile is set by **specifier proteins** (ESP, NSP) that ride alongside myrosinase. Tune that balance and the same stored glucosinolate hydrolyzes toward a product the receptor barely sees. We prefer the upstream valve where we can reach it, because it never touches the plant's response to a real pest; a brassica that cannot defend itself is not nourishment, it is a casualty. This is the standing constraint of the glucosinolate road and we say it plainly: the bitterness and the defense are the same chemistry, and the art is lowering the first without disarming the second.

Road two — the universal highway

The phenylpropanoid branch is the opposite kind of road: not a sequestered weapon but the busiest thoroughfare in the plant, the trunk from which lignin, pigments, UV screens, and a long catalogue of flavors all depart. It begins at **phenylalanine**. **PAL** (phenylalanine ammonia-lyase) removes ammonia to give *trans*-cinnamic acid; **C4H** (cinnamate 4-hydroxylase, a cytochrome P450) hydroxylates the ring to *p*-coumaric acid; **4CL** (4-coumarate:CoA ligase) spends an ATP to activate it as *p*-coumaroyl-CoA. That CoA thioester is the great fork of plant chemistry, and two of its exits make most of the bitterness in this book.

The first exit builds coffee's harshness. A hydroxycinnamoyl transferase (**HCT**, and its quinate-specific cousin HQT) couples the coumaroyl-CoA to **quinic acid**; a ring hydroxylation (**C3′H**) finishes it into **chlorogenic acid** — 5-caffeoylquinic acid, the single most abundant soluble compound in a green coffee seed, present at a few percent of its dry weight. Green CGA is only mildly bitter. The bitterness of a cup is mostly *made in the roaster*, from CGA, by heat: moderate roasting cleaves it into **chlorogenic acid lactones**, which are distinctly bitter; push the roast dark and those lactones condense further into **phenylindanes**, the source of the harsh, lingering, almost metallic bitterness of a French roast. This is the honest and slightly humbling fact of the coffee road — much of what the drinker tastes is

thermal chemistry downstream of anything the plant did — and it sets the terms of our lever. We cannot roast for the roaster. What we *can* do is meter the **substrate**: lower the green seed's chlorogenic-acid load at the HQT valve, by a measured fraction, so that every roast the buyer chooses to run starts from less of the precursor that the heat turns bitter. Less fuel in the seed, less bitterness out of the fire, at the roast the drinker already loves.

The second exit off *p*-coumaroyl-CoA is the **flavonoid arm**, and it carries the bitterness and astringency of cacao, tea, and the grape. Chalcone synthase (**CHS**) condenses the coumaroyl-CoA with three malonyl-CoA units into naringenin chalcone; isomerase and hydroxylase steps carry it to the dihydroflavonols; **DFR** (dihydroflavonol 4-reductase) reduces those to leucoanthocyanidins, and then two reductases set the stereochemistry of the **flavan-3-ols** — **LAR** (leucoanthocyanidin reductase) makes *trans*-catechin, **ANR** (anthocyanidin reductase) makes *cis*-epicatechin. A single flavan-3-ol monomer, one epicatechin, is bitter. Link a handful of them end to end and you have a **procyanidin** — a condensed tannin — which is a different sensation entirely. We come to that polymerization where the two neighbors of this chapter meet.



Cacao in cross-section — the flavonoid arm of the phenylpropanoid highway lays down the flavan-3-ols here; a monomer reads bitter, a polymer reads as the dry grip of a tannin.

The flavonoid arm also builds the sharpest bitterness in all of fruit, and it hides a lesson about precision that we love. Grapefruit's bitterness is **naringin** — the flavanone naringenin dressed in a disaccharide. Here is the whole of it: naringenin wearing a **neohesperidose** (rhamnose linked α-1→2 to glucose) is intensely bitter; the *same* naringenin wearing a **rutinoside** (rhamnose linked α-1→6 to the very same glucose) is **tasteless**. One bond angle on one sugar, a 1→2 against a 1→6, is the entire difference between a mouth-puckering bitterness and nothing at all — because only the 1→2 shape presents the aglycone to the receptor's pocket. The valve is a single rhamnosyltransferase choosing between two linkages. We find that fact beautiful, and we find it instructive: it is the clearest proof on the shelf that bitterness is not "how much phenolic a plant holds" but

"what exact shape reaches the receptor," and that a lever placed on one enzyme can silence a bitterness without removing a single carbon of the molecule that carried it.

How the tongue reads bitter

Twenty-five receptors in the human genome are **TAS2Rs**, bitter receptors, and they are the reason bitterness is a single sense built to recognize a thousand unrelated poisons. They are class-A G-protein-coupled receptors, seven membrane passes each, and they range from the narrowly tuned (a receptor that answers to one or two molecules) to the broadly tuned (TAS2R14 or TAS2R46, each reading dozens). A bitter food does not strike one receptor; it strikes a **chord** across several, and the brain reads the chord.

The transduction is the same in every one, and it is worth walking because it tells you what a "loud" bitterness physically is. A bitter ligand settles into the receptor's pocket and nudges it into its active shape. The receptor exchanges GDP for GTP on its partner G-protein, **gustducin**; the freed **Gβγ** subunit switches on **phospholipase C-β2 (PLCβ2)**, which cleaves the membrane lipid **PIP_2** into **IP_3** and diacylglycerol. IP_3 opens its receptor on the cell's internal calcium store, and **Ca^{2+}** floods the cytosol. That calcium opens **TRPM5**, a monovalent-cation channel, and sodium rushing in **depolarizes** the taste cell. Depolarization opens voltage-gated channels and, through the large-pore channel **CALHM1**, the cell releases **ATP** into the synapse, where it excites the gustatory nerve fibers ($P2X_2$/$P2X_3$) that carry the signal to the brainstem. "Loud," at the molecule's level, means a ligand that drives this cascade to half of its maximum at a *low* concentration — a tight fit, a low half-maximal activating dose. That single physical quantity, the **EC_{50}**, is the hinge of everything numeric in this book.

TAS2R38 earns its place as our ruler for a reason beyond chemistry: it is the most-studied taste receptor in the human body, because it is the one that varies most between people. A three-letter haplotype controls it, three amino-acid positions whose variants sort the human population into tasters and non-tasters of the brassica bitter signal:

POSITION	TASTER (PAV)	NON-TASTER (AVI)	REFERENCE SNP
49	Pro (P)	Ala (A)	rs713598
262	Ala (A)	Val (V)	rs1726866
296	Val (V)	Ile (I)	rs10246939

Roughly seven in ten people carry at least one copy of the taster haplotype and register goitrin and the brassica isothiocyanates sharply; the rest taste them faintly or not at all. This is not an abstraction. It is the molecular reason one child at the table gags on broccoli while another eats it happily — the same floret, two different receptors reading it — and it is one of the quiet reasons we care about this axis at all. A vegetable whose bitterness is metered down the glucosinolate road becomes food a taster can finally love, without a non-taster losing anything. We measure every brassica against the **taster** receptor, deliberately, because designing for the person who tastes the most is how you reach everyone.

The ruler and the chord

Here is the first organizing fact come due. TAS2R38 is the sharp, well-mapped ruler, and it is genuinely the receptor doing the reading for the glucosinolate column. But a coffee's harshness is a chord struck mostly across TAS2R46 and TAS2R7; a hop's bite lands on TAS2R1 and TAS2R14 and TAS2R40; theobromine plays five receptors at once. The Index still scores every one of these against the TAS2R38-anchored axis, because a standard needs one ruler and TAS2R38 is the ruler we trust most and the public knows best. But the honest footnote on every non-brassica readout is that the single number is a **projection** — a multi-receptor reality flattened onto one well-characterized axis. The projection is faithful enough to compare a coffee to a coffee and to score an attenuation against its own baseline; it is not a claim that the coffee's bitterness lives in TAS2R38. We print that caveat here and on every entity page, because a number that pretends to be more precise than the biology is a number that has stopped being science, and across a whole library the one thing we will not spend is the reader's trust.

The Index, in numbers

Every axis in this library is a single number anchored to a receptor's half-maximal activation, and the bitter axis is no exception:

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

EC_{50} is the concentration that drives the receptor to half of its full response; invert the candidate's against a fixed reference, scale by a thousand, and loudness becomes a score. Two quantities hide inside that candidate EC_{50}, and the honest practitioner keeps them apart. The first is **potency** — how tightly the molecule fits the receptor, a property fixed by chemistry and written as the pure compound's own EC_{50}. The second is **load** — how much of that molecule a food actually carries to the tongue. A potent ligand present in a trace, and a weak one present in a flood, can read at the same loudness. So we measure the axis the way the mouth meets it: we take the **real food**, not a purified compound, and serially dilute it against the receptor until activation falls to half. The dilution at which it does is the food's **effective EC_{50}**, and the Index is that value inverted against the reference probe. This is precisely why our lever bites: a plant-side edit cannot change a molecule's intrinsic potency, but it can change the load — and lowering the load means more food is needed to reach half-activation, so the effective EC_{50} climbs and the Index reading falls.

We anchor the bitter axis with the standard thiourea probe of TAS2R38, **PROP** (6-*n*-propyl-2-thiouracil), and set its half-maximal activation as the reference point so that the reference reads exactly 1,000. For the worked figures that follow, take that anchor at **EC_{50}(ref) = 2.0 µmol/L** — a realistic half-maximal for PROP at the taster receptor, and a clean number to compute against. Then:

CANDIDATE	EFFECTIVE EC_{50}	IFI-B = 1000 · 2.0 / EC_{50}
reference probe (the anchor)	2.00 µmol/L	**1,000**
twice as loud	1.00 µmol/L	2,000
four-fifths as loud	2.50 µmol/L	800
half as loud	4.00 µmol/L	500
quiet after a deep attenuation	6.67 µmol/L	300

The arithmetic is deliberately plain so the reader can run it: halve the EC_{50} and the score doubles, because loudness is inverse concentration. Run it backward to read a known score as a concentration. **Phoenix**, the one cultivar in the whole library designed to register on all eight axes at once, carries a predicted bitter reading of **B = 2,200** — the single place we tune bitterness *up to a target* rather than down. Invert the formula: $EC_{50}(c) = 1000 \cdot EC_{50}(ref) / IFI\text{-}B = 1000 \cdot 2.0 / 2{,}200 =$ **0.91 µmol/L**. That is what a bitter axis of 2,200 physically means — a stimulus that our model predicts would half-activate the receptor at a little under half the reference concentration, a deliberate, bright, designed bitterness woven in beside the heat:

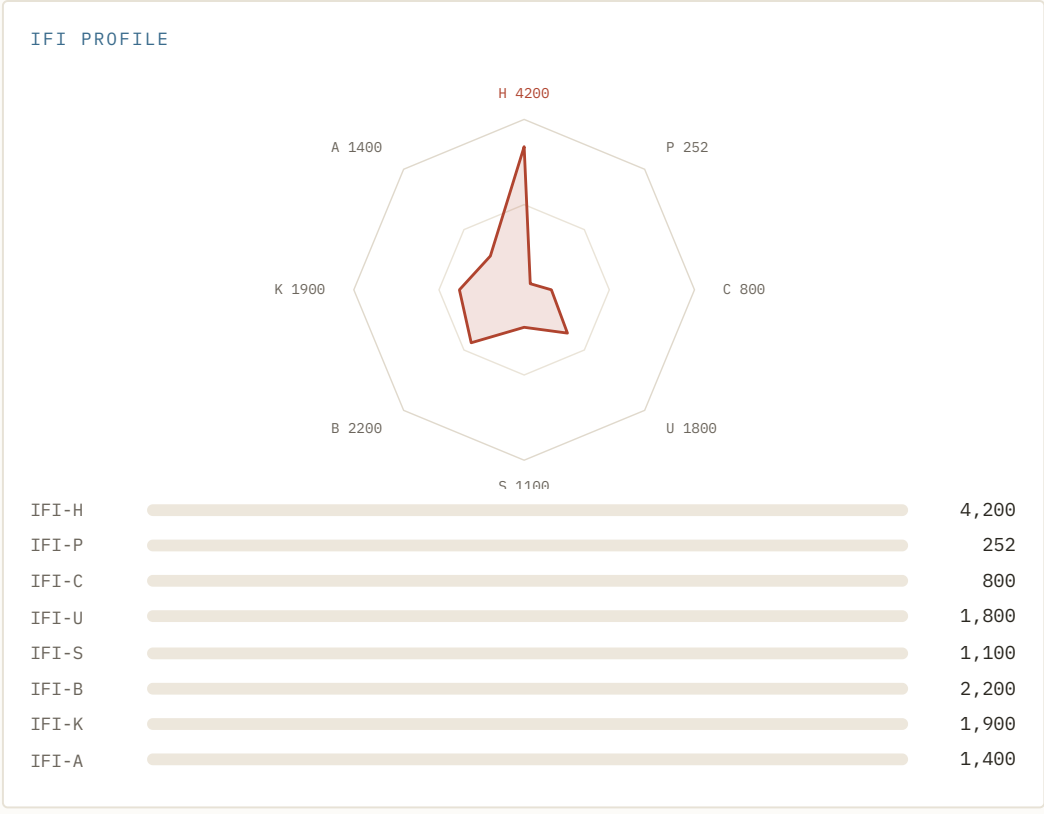

Now the question the whole axis exists to answer: when we turn a valve and remove some fraction of the bitter end-product, how much quieter does the food actually get? Not linearly, and the nonlinearity is the most important thing in this chapter. Receptor response follows a **Hill curve** — a sigmoid — in the ligand's concentration *C* relative to its EC_{50} (call it *K*):

> $R(C) = C^n / (K^n + C^n)$

with *n* the Hill coefficient, typically between 1 and 2 for TAS2Rs; we take **n = 1.5** as a representative value. "Attenuation" on this road means cutting the ligand load by some fraction *a* (the book's working range is **a = 0.20 to 0.65**), moving *C* down to (1 − *a*)·*C*. The perceptual payoff of that cut depends entirely on *where the food sat on its own curve to begin with.* Take a single, fixed **40% cut** (a = 0.40) and apply it at five starting points:

STARTING C / K	RESPONSE BEFORE	RESPONSE AFTER −40%	PERCEIVED BITTERNESS REMOVED
0.5 (below EC_{50})	0.261	0.141	**46.0%**
1.0 (at EC_{50})	0.500	0.317	36.5%
2.0	0.739	0.568	23.1%
4.0	0.889	0.788	11.3%
8.0 (deep saturation)	0.958	0.913	4.6%

The identical 40% removal of ligand buys a 46% drop in perceived bitterness near the receptor's threshold and barely a 5% drop when the bitterness was already saturating the receptor. This single table explains the book's honest 20–65% range better than any prose: the yield of an attenuation is not a property of our edit alone but of the chemistry it meets. A coffee whose bitter principals sit near the shoulder of the curve rewards a modest cut lavishly. A beer drowning in iso-alpha-acids, far up the flat top, barely notices the same cut — which is exactly why the hop is one of the hardest organisms on this axis and earns its own long chapter. We aim, wherever the plant lets us, for the **shoulder**, the steep middle of the curve where subtraction is most generous, and we tell the buyer the truth when a target sits in saturation.

One more piece of arithmetic is simply housekeeping, but formulators live and die by it: moving between the concentrations a chemist reads (mg/L) and the molar concentrations a receptor assay speaks (µmol/L). Divide a mass concentration by the compound's molar mass.

BITTER PRINCIPAL	FORMULA	MOLAR MASS (G/MOL)
Caffeine	$C_8H_{10}N_4O_2$	194.19
Theobromine	$C_7H_8N_4O_2$	180.16
Chlorogenic acid (5-CQA)	$C_{16}H_{18}O_9$	354.31
Naringin	$C_{27}H_{32}O_{14}$	580.54
Isohumulone	$C_{21}H_{30}O_5$	362.46
Goitrin	C_5H_7NOS	129.18
Epicatechin	$C_{15}H_{14}O_6$	290.27

A worked pair. A moderately hopped beer at **40 IBU** carries close to 40 mg/L of iso-alpha-acids (the brewer's convention is roughly 1 IBU ≈ 1 mg/L): 40 mg/L ÷ 362.46 g/mol = 0.1104 mmol/L = **110 µmol/L** of isohumulone — a load that, against isohumulone's own EC_{50} at its receptors, sits well up the saturating shoulder, which is the molecular statement of why hop bitterness resists a gentle hand. And on the glucosinolate road: a brassica sap reading, say, 12 mg/L of goitrin converts as 12 ÷ 129.18 = 0.0929 mmol/L = **93 µmol/L**, a figure you then place on goitrin's own curve to predict what a given attenuation will be worth. The conversions are trivial; forgetting to do them is how a formulation misses by a factor of a molar mass.

Finally, the honesty of error. Every batch we grow carries a manufacturing spread, and across our constructs that batch-to-batch coefficient of variation holds at **≤ 7%**. Because the Index is a **ratio**, relative variances add in quadrature: $CV(IFI) = \sqrt{(CV(ref)^2 + CV(c)^2)}$. Hold the reference as a tightly controlled calibration standard and the candidate's ≤ 7% dominates, so a bitter reading of 2,200 carries roughly **± 150** at one standard deviation; fold in a few percent of uncertainty on the reference itself and it widens to about **± 170** ($\sqrt{(0.07^2 + 0.03^2)} \approx 7.6\%$). We write the axis values as whole numbers because that is how they are read, but we hold them to the band the biology allows, and we would rather publish 2,200 ± 170 than pretend to a 2,217 we cannot stand behind.

Bitterness and astringency, the same road in two coats

Now the second organizing fact. Return to the flavan-3-ols of the flavonoid arm. A small monomer — a single epicatechin, a single catechin — is the right size to slip into a bitter receptor's pocket and read as **bitter**. Link a few of those monomers into a **procyanidin**, a condensed tannin, and the polymer grows too large for any receptor pocket. But it is now exactly the right size to do something else: it cross-links the proline-rich proteins in your saliva, precipitating them out of solution, and the mouth reports the loss of its own lubrication as the dry, puckering, protein-stripping grip we score on an entirely separate axis — **IFI-A, astringency**, read not by a taste receptor at all but by the mechanics of saliva. The same biosynthetic road, the same two reductases, the same flavan-3-ol building block, feeds **both** sensations; the only variable is the degree of polymerization. In cacao and tea and the grape the monomers and polymers are braided so tightly that you cannot tune one without reckoning with the other — lower the flux into the arm and both the bitterness and the astringency fall together, which is sometimes exactly wrong, because a wine or a fine chocolate *wants* its astringent grip and only wants the bitter edge taken off.

This is why the bitter axis needed an astringent axis beside it to tell the truth about a cup of tea, and why our work on this road is never a single valve but a choice between valves: throttle **ANR/LAR** and you lower monomer and polymer alike; move further up, at the branch into the arm, and you starve pigment and tannin along with the bitterness. The design question is always *which pool* — monomer or polymer — and the restraint is always to leave the astringent thread where the flavor lives. The Cacao, Tea, and Grape monographs, and the Astringency volume, carry the full account of that braid; on this axis we turn down only the bitter strand of it, and we take great care to leave the other where the mouth wants it. You can see the whole dilemma in a single glance at Phoenix's profile: a predicted bitter reading of 2,200 and a predicted astringent reading of 1,400, two neighbors off one road, each dialed to its own target rather than collapsed into each other. That is the discipline the two axes exist to enforce.

What we turn, and what we leave

Across the whole bitter shelf, a little over three thousand constructs sit on this one axis today, and they sort cleanly by road:

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY	ROAD
Grape	1,610	B / A / M	Hard	flavonoid arm
Broccoli	380	B / M	Moderate	glucosinolate
Tea	280	U / B / A	Hard	flavonoid arm
Broccoli-Kale	280	B	Hard	glucosinolate
Orange	120	B / C	Hard	flavanone glycoside
Potato	80	B	Easy	steroidal glycoalkaloid
Hops	80	B	Moderate	cone polyketide
Coffee	40	B	Hard	chlorogenic acid
Cabbage	40	B	Moderate	glucosinolate
Kale	40	B	Moderate	glucosinolate
Radish	40	B	Hard	glucosinolate
Rapeseed	40	B	Hard	glucosinolate

The table is a map of the method as much as the molecules. The glucosinolate road claims the brassicas and holds the standing constraint of defense; the flavonoid arm claims the grape and tea and the two neighbors braided through them; coffee stands alone with a bitterness half-made in the roaster; and potato — eighty constructs, the one "Easy" line on the shelf — is a quiet reminder that the two-road map is a convenience and not a law, because potato's bite comes off neither road but the **steroidal glycoalkaloid** route, from cholesterol, a third chemistry entirely. We keep the exception in the table rather than hide it, because a map that erases its own edges is a map you cannot trust at the edges.

And the edges are where the honest failure modes live. Cut the chlorogenic-acid valve too deep and you do not just lose the harshness a dark roast would have made — you lose the precursor to a whole family of roast aromatics the drinker loved, and a coffee can go clean and flat at once. Throttle the flavonoid arm too far and the grape loses the tannin that makes it ageable along with the bitterness that made it green. Chase a brassica's goitrin without respecting the specifier proteins and you can trade a bitter floret for one that cannot answer a pest. These are not reasons to stand down; they are the reasons to turn the valve by a measured fraction and no more, in one tissue, at one node, with the dose-response curve in front of us telling us what the fraction is worth. Subtraction is the gentlest thing we know how to do to a plant, and bitterness is where it does its most generous work — because under almost every bitterness on this shelf sits a molecule the body is glad to keep: a polyphenol that is also an antioxidant, a glucosinolate that is also a defense the plant earned over millions of years. We are not stripping those away. We are turning down only the single harsh note each one strikes on the tongue, by just enough that a person can eat the whole nourishing thing with pleasure, and leaving everything else — the color, the keeping, the defense, the quiet good the plant does the body — exactly where we found it. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 7 · The Design Loop — *scoring an edit before we cut.*

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

Every edit we have ever written into a living plant began the same way — not in soil, not with a blade, but at the table where this family still argues about dinner. The question that starts the work is never *can we cut this gene*; modern tools make almost any cut possible, and possibility is the cheapest thing in biology. The question is *should we*, and *how far*, and *what will the plant do with the rest of its life once we have.* Bitterness makes the question sharp. A coffee cherry, a head of broccoli, a hop cone — each carries its bitterness for a reason the plant settled long before we arrived, and to quiet it without wounding everything around it we have to know four things at once: the molecule, the enzyme that builds it, the receptor on the tongue that hears it, and the entire metabolic city the enzyme lives inside. So we do not begin by cutting. We begin by predicting, and we do not touch a plant until the prediction has earned it.

The loop, in one breath

The whole method is a circle with four stations — *predict, build, measure, feed back* — and the discipline is that you never skip a station and never run it backward. We predict an edit in software, score it before a single base is changed, and rank the entire space of possible edits against each other. We build only the handful that survive, writing them in with CRISPR. We measure what we built against the receptor it was meant to quiet, in a dish and then on a trained human tongue. And we feed the measured truth back into the models, so the next round is sharper than the last. This is the same loop, pointed now at bitterness, that first proved itself on heat and went on to produce the roughly 19,500 catalogue products in our library and the forty-two host species live today. The science did not have to be reinvented for the bitter axis. It had to be aimed.

What makes the loop trustworthy is that it begins in two kinds of silence, asking two different questions with two different models — one about a single protein, one about an entire organism — and refusing to build until both have answered.

The first model — a grammar of proteins

The first model is a **protein language model**: in our pipeline an ESM2-class transformer of roughly 650 million parameters, trained on the evolutionary record of hundreds of millions of natural protein sequences. It is never shown a textbook of biochemistry. It is shown sequences, with residues hidden, and asked to guess what belongs in the gap — the same masked-prediction task a language model runs on words — until, to do that well, it has been forced to internalize the grammar of what a working protein looks like: which residues pack into a hydrophobic core, which cysteines pair into a disulfide, which positions a billion years of selection has refused to let drift.

Point that trained grammar at one of our target enzymes — the HQT transferase that builds coffee's chlorogenic acid, say — and it will score *every possible single-residue substitution*, every one of the roughly 19 alternatives at every position, before we synthesize a thing. The score for exchanging the wild-type residue *a* at position *i* for a candidate *b* is a log-likelihood ratio, the model's own measure of how much more or less "proteinlike" the sequence becomes:

```
s(i, a→b) = log P(b | context) − log P(a | context)
```

A strongly negative score means the model has almost never seen *b* survive in that context across all of evolution — a residue the protein's own history forbids. A score near zero means the substitution is one nature tolerates, and therefore one we can reach for without collapsing the fold. We fold that evolutionary score together with a structure-based stability term and a predicted catalytic-efficiency term, and rank the whole substitution landscape by the product. Instead of guessing which change might *down-tune* an enzyme cleanly — nudging its throughput down without breaking it outright — we read the entire menu of changes first and build from the top.

The stability term speaks in the physical currency of **ΔΔG**, the change in folding free energy a mutation causes, in kilocalories per mole, by the convention that negative is stabilizing and positive is destabilizing. A number in that currency is not decoration; it converts directly into how much of the protein population is correctly folded and working at the temperature of the assay. At 298 K the thermal quantum is $RT = (1.987 \times 10^{-3}\ \text{kcal} \cdot \text{mol}^{-1} \cdot \text{K}^{-1})(298\ \text{K}) = 0.593$ kcal/mol, and the folded-to-unfolded population shifts by the Boltzmann factor

```
exp(−ΔΔG / RT).
```

So a modest-looking ΔΔG of −1.8 kcal/mol is not modest at all: exp(1.8 / 0.593) = exp(3.04) ≈ **20.8×** more favorable a folded population than the wild-type starting point. That is the difference between an edit that survives harvest-to-harvest and one that works in the greenhouse and falls apart in the field.

Reading the confidence before trusting the edit

An edit is only as trustworthy as the structure we score it on, and that structure is itself a prediction — so before we believe any ΔΔG we read the model's confession of its own confidence.

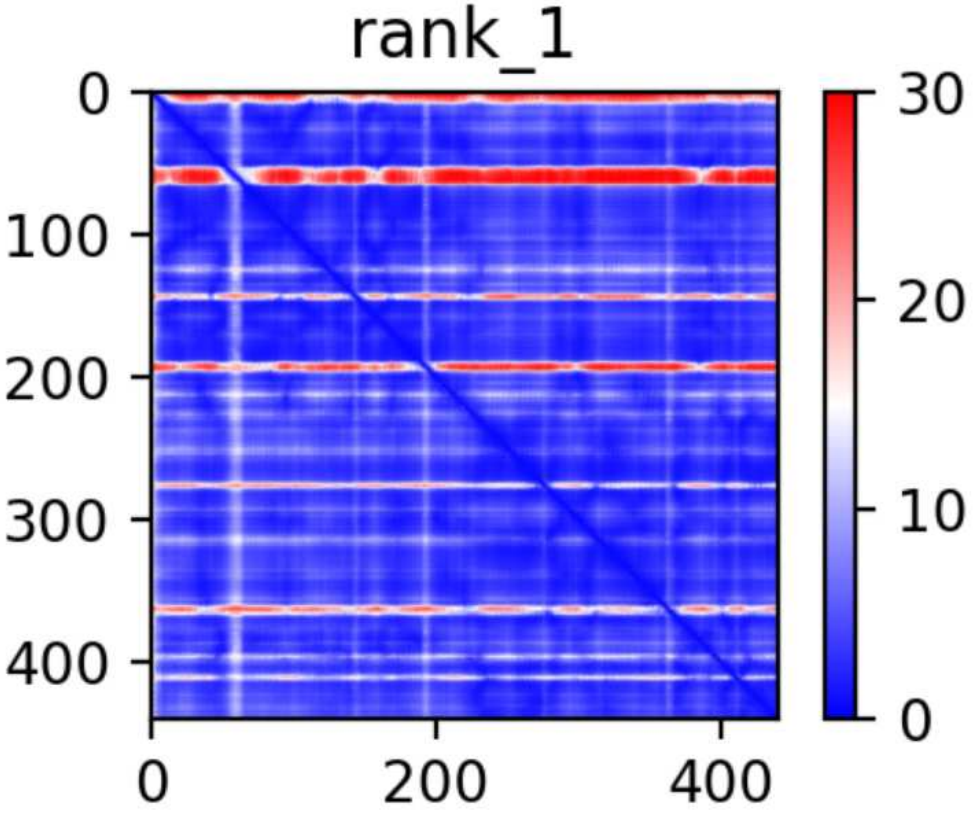

FIG. — Predicted aligned error for the platform's structure pipeline, shown on the Pun1/AT3 model: dark where the model is certain how two regions of the fold sit relative to each other, bright where it is not.

Two maps govern that trust. The first is **pLDDT**, a per-residue confidence running 0 to 100 — above 90 the backbone is pinned down almost to the atom, 70 to 90 is confidently placed, below 50 is very likely a disordered loop the model is honestly shrugging at. We permit an edit only where the residue beneath it scores high; a substitution placed on a low-pLDDT loop is a guess stacked on a guess. The second map, shown above, is the **predicted aligned error**, an $N \times N$ grid measured in Ångström: cell (i, j) is the expected error in residue i's position once the structure is aligned on residue j. Dark means two regions are confidently locked relative to each other — exactly what you need when an edit at the binding domain is only safe if its geometry relative to the active-site cleft is certain. Bright across a domain interface means the model does not know how the two halves hinge, and an edit that depends on that hinge is one we hold back. We trust an edit most where the structure beneath it is read most clearly, and the map tells us, square by square, where that is.

This is the first honest limit of the loop, and we state it plainly: a high-confidence prediction is still a prediction. The structure gates *where* we are willing to design; it never gates *whether* the design worked. Only the assay does that.

The second model — the whole plant at once

The language model reasons about one protein in isolation. But an enzyme does not live in isolation; it lives in a metabolic network of thousands of reactions that can route around a blocked step the way a city reroutes around a closed bridge. So the second model asks the larger question: *lower this one bitter metabolite by this much, and what else in the plant moves?*

We answer it with a **genome-scale flux-balance analysis** — a mathematical reconstruction of the organism's entire metabolism as a stoichiometric matrix **S**, with one row per metabolite and one column per reaction. A flux vector **v** lists the rate of every reaction at once.

Across any steady interval the plant neither floods nor drains its internal pools, which is to say production balances consumption for every metabolite simultaneously:

```
S · v = 0, subject to lb ≤ v ≤ ub.
```

Those bounds encode what the plant can and cannot do — irreversible steps, nutrient uptake limits, an enzyme we have throttled. Within that feasible space we maximize a biological objective, classically biomass, and the whole thing is a linear program with a clean answer. The power for our work is in what the solution's *shadow prices* reveal: how far the objective moves per unit change in each constraint. To find the true valve on a bitter pool we run a knockout screen — for each candidate enzyme *e* we re-solve the program with its flux pinned to zero and record two numbers, the drop in the bitter metabolite's production and the drop in biomass. The valve we want is the one that maximizes the ratio

```
Δ(bitter flux) / Δ(biomass)
```

— the most quieting per unit of cost to the plant's own life. Flux balance is how we know, before cutting, that a given enzyme is a *committed* valve and not a step the network will simply reroute around. It guards against the oldest error in metabolic engineering: editing a reaction that looks important, and watching an isozyme or a parallel branch carry the flux as if nothing happened. When two enzymes feed a bitter pool in parallel, knocking out one alone moves the pool by only the fraction it carried, and the model shows us that fraction is small — telling us to climb to the shared node above the branch instead of wasting a build on a single leaf of it.

The valves of bitterness

Those two models, run together, pick out a short list of valves across the plants this axis cares about. Each one is a real enzyme with a mechanism, a bitter molecule it builds, and a receptor consequence we can measure.



A coffee cherry in cross-section — the chlorogenic pool that HQT fills, and that the roast turns bitter.

In **coffee**, the valve is **HQT**, hydroxycinnamoyl-CoA quinate hydroxycinnamoyltransferase, which condenses caffeoyl-CoA with quinic acid to build chlorogenic acid. Green coffee hoards chlorogenic acid; the roast then cracks it into the lactones and, deeper, the

phenylindanes that read as the lingering, drying bitterness at the back of a dark cup. Throttle HQT and you lower the precursor pool the roast draws on — fewer bitter fragments formed downstream, without touching the aromatics the bean is loved for. In **hops**, the valve is **VPS**, valerophenone synthase, a type-III polyketide synthase that condenses a branched acyl-CoA with three malonyl-CoA units to lay down the acylphloroglucinol core of the alpha- and beta-acids; those alpha-acids isomerize in the kettle into the iso-alpha-acids that *are* the bitterness of beer. In the **tannin plants** — grape, tea — the branch runs through **LAR** and **ANR**, leucoanthocyanidin reductase and anthocyanidin reductase, which together fill the catechin and condensed-tannin pool that carries both bitterness and the drying grip of astringency, so a single valve there moves two axes at once. And in the **brassicas**, two valves work as a pair: **MYB28**, the master transcription factor that sets the volume of the whole aliphatic-glucosinolate program, and **GS-OH**, the enzyme that routes flux at a branch point toward progoitrin — the precursor that hydrolyzes, on chewing, into goitrin, the compound the bitter receptor hears most loudly.

VALVE	WHAT IT BUILDS	THE BITTER PRINCIPLE	HOST · CONSTRUCTS	RECEPTOR CONSEQUENCE
HQT	caffeoyl-CoA + quinate → chlorogenic acid	roast-born lactones, phenylindanes	Coffee · 40	fewer TAS2R agonists after the roast
VPS	acyl-CoA + 3 malonyl-CoA → acylphloroglucinol core	humulones → iso-α-acids	Hops · 80	IBU set by design, not by the kettle
LAR / ANR	leucoanthocyanidin → catechin; anthocyanidin → epicatechin	catechins, condensed tannins	Grape · 1,610 · Tea · 280	bitter *and* astringent load fall together
MYB28	master regulator of aliphatic glucosinolates	the whole glucosinolate pool (volume)	Broccoli · 380 · Cabbage · 40 · Kale · 40	scales total bitter-glucosinolate signal
GS-OH	routes flux toward progoitrin at the branch	goitrin	Broccoli · 380 · Rapeseed · 40	flux redirected toward glucoraphanin

The brassica pair is the clearest picture of why the two models are better than either alone. The language model can tell us how to down-tune GS-OH cleanly, position by position. Only flux balance can promise that the flux we take *off* the goitrin branch flows *onward* into glucoraphanin — the non-bitter molecule a head of broccoli is prized for — rather than simply backing up into something worse. MYB28 is the volume knob on the whole program; GS-OH is the routing switch at one fork within it. We turn them together, and we turn them toward the molecule we would rather you taste. There is a quiet pleasure, for

a family that has eaten this way its whole life, in an edit whose entire purpose is to move a plant away from the compound that only ever tasted like a warning.

How the tongue reads bitter

None of this matters until it reaches the receptor, so it is worth walking the bitter signal from molecule to nerve, step by step, because every number this axis reports is anchored there.

Bitterness is read by the **TAS2R** family, about twenty-five G-protein-coupled receptors, and the one this axis is anchored to is **TAS2R38** — the receptor famous for phenylthiocarbamide, and the one tuned to the thiourea group, the N–C=S motif that goitrin carries. The cascade runs like this. A bitter molecule settles into the receptor's binding pocket; the receptor, a seven-pass membrane protein, changes shape and activates its partner G-protein, gustducin, on the inside of the cell. Gustducin's freed subunits switch on phospholipase C-β2, which cleaves a membrane lipid into the second messenger IP_3; IP_3 opens stores that release a pulse of calcium; the calcium gates the TRPM5 channel; the cell depolarizes and releases ATP onto the afferent nerve, and only then — four amplifying steps downstream of the molecule — does the brain receive the word *bitter.* It is a cascade built for gain, which is exactly why bitterness is so hard to switch off by brute force: a very little agonist, multiplied through four stages, still shouts.

The receptor also explains why our final judge is a human panel and not only a dish. **TAS2R38** is famously polymorphic — three amino-acid positions give the common taster and non-taster haplotypes, so two people can sip the same coffee and honestly disagree about its bitterness. A number that ignored that would be a number we could not trust. So the panel is genotype-balanced and replicated, and it exists to confirm one thing: that the quieting we predicted in the dish is the quieting a person feels on the tongue.

The arithmetic of quieting

Here is where the loop earns its claim to be engineering rather than hope, because the relationship between *how much molecule we remove* and *how much bitterness a person loses* is not one-to-one — it bends — and the whole Index is built to respect that bend.

The bitter signal follows a dose-response. Write the fraction of receptor driven by a bitter compound at concentration c as a Hill function,

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

where n is the Hill slope and **EC_{50}** is the concentration that drives the receptor to half its full response. That half-max point is the anchor of everything, and it has a lovely property: set $c = EC_{50}$ and $\theta = EC_{50}{}^n / (EC_{50}{}^n + EC_{50}{}^n) = ½$ exactly, *for every value of n.* The Hill slope changes how steeply the curve climbs, but the half-max concentration is where it always passes through one-half — which is why EC_{50}, and not any point further up the curve, is the robust number the loudness score is built on.

Now watch the bend. Take a representative Hill slope of $n = 1.5$ for this class of bitter GPCR, and suppose a dark roast delivers its dominant bitter lactone at three times the receptor's EC_{50} — $c = 3 \cdot EC_{50}$, well up on the saturating shoulder, which is precisely why over-roasted coffee tastes relentless. Its receptor occupancy is

$$\theta_0 = 3^{1.5} / (1 + 3^{1.5}) = 5.196 / 6.196 = 0.839.$$

Throttle HQT to cut the chlorogenic pool, and with it the lactone, by threefold, so the delivered concentration falls to $c = 1 \cdot EC_{50}$. Now

$$\theta_1 = 1^{1.5} / (1 + 1^{1.5}) = 0.500,$$

and the bitterness a person actually loses is the drop in occupancy,

$$A = 1 - \theta_1/\theta_0 = 1 - 0.500/0.839 = 0.40.$$

A threefold cut in the molecule buys a **40% attenuation** — real, squarely inside the 20–65% band this axis lives in, and *less* than threefold, because the receptor was saturating and we spent the first part of our cut merely coming down off the plateau. The bend is the honest lesson. To reach the top of the band, 65%, we solve for the occupancy that satisfies it, $\theta_1 = 0.35 \cdot \theta_0 = 0.294$, which requires $c \approx 0.56 \cdot EC_{50}$ — a *5.4-fold* molecular cut, nearly double the metabolic surgery for the last stretch of relief. This is why we do not promise elimination and do not chase it: past a point, each additional percent of quieting costs disproportionately more disruption to a plant we would rather leave whole. The loop finds the valve that lands the most attenuation for the least intrusion, and then it stops.

The proof on heat

The proof that this loop works is not on the bitter axis at all; it is the plant that made the company, and the arithmetic of that proof is worth setting down because it is the same arithmetic we now run on bitterness.

The method's first subject was **Pun1**, the terminal capsaicin synthase — the enzyme that labels the final product. The language model scored every substitution; flux balance found that most of the control sits upstream in the phenylpropanoid pathway, at 4CL and HCT/CCoAOMT rather than at Pun1 itself — Pun1 completes the product but does not set the rate. Two edits emerged from the top of the ranked landscape that no breeder had ever found or combined — S39L and L345G — written in with CRISPR, our model predicts they would lift a fruit to approximately **2.95 million SHU** at a designed batch-to-batch coefficient of variation of seven percent or less, against the forty-two percent that is simply nature's scatter in a wild population.

EDIT	STRUCTURAL EFFECT	ΔΔG (KCAL/MOL)	FOLDED-POPULATION FACTOR, EXP(−ΔΔG/RT)
S39L	binding affinity ↑ (N-terminal)	−2.1	≈ 34.6×
L345G	steric relief, active-site cleft	−1.8	≈ 20.8×
C175S	disulfide-misfold prevention (structural annotation — not in shipped construct)	−0.9	≈ 4.6×

Read the third column as the physics of why these two hold. Each stabilizing ΔΔG becomes a folded-population factor through the same Boltzmann relation the stability model speaks in, exp(−ΔΔG / 0.593). If the two shipped edits were perfectly independent their free energies would add — a combined −3.9 kcal/mol, a folded-population factor near exp(3.9/0.593) ≈ 718× — but independence is an assumption the molecule is free to violate through epistasis, and that assumed 718× is exactly the kind of number we refuse to believe until the assay returns it. The prediction is a hypothesis with a decimal point, nothing more, until it is confirmed.

And the edits behave as a *ladder*, not a switch, which is the real signature of design. Across the ten flagship designs our model predicts receptor loudness climbing with heat — from Ember at a predicted 866K SHU and an IFI-H of 1,570, through Blaze at a predicted 1.72M SHU and 2,630, to Phoenix at a predicted 2.95M SHU and 4,200. Fit a line through all ten and the relationship is tight:

> IFI-H ≈ 680 + 1,170 · (SHU ÷ 10⁶), R² ≈ 0.97.

It is tight but not perfect, and the imperfection is instructive: Kindle carries a predicted 1.25M SHU yet an IFI-H of only 2,000 — below Cinder's 2,290 at a lower predicted 1.05M SHU — because SHU is a dilution count of total capsaicinoid mass while IFI-H is receptor loudness, and the acyl-chain composition shifts the balance between capsaicin and dihydrocapsaicin, changing how loudly a given mass speaks to TRPV1. The residual in the fit is not noise; it is a second design knob the loop can reach for. This is 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 hop synthase exactly as they once scored Pun1.

Closing the loop

Then the loop closes, and this is the station that keeps the rest of it honest. An edit scored and built is still only a hypothesis until it is **assessed** — against the receptor, as a number.



The bench where the loop closes — the engineered plant's compounds, read back against the receptor they were designed to quiet.

The engineered plant's bitter compounds go onto the cell-based receptor assay that anchors this whole Index — cells displaying the human receptor, read for the concentration at which they reach half-

maximal response — and that assay returns an EC_{50}. From the EC_{50} comes the score, by the one formula the entire library resolves against:

> $IFI(c) = 1{,}000 \cdot EC_{50}(\text{reference}) / EC_{50}(\text{candidate})$.

The anchor is built into it: when a candidate's EC_{50} equals the reference's, the score is exactly 1,000, so a reading above 1,000 means a compound that speaks *louder* than reference and one below means *softer.* The arithmetic runs both ways and we use both. Read our flagship pepper's computed profile across all eight axes — the one cultivar we hold a canonical full-spectrum design for — and the formula turns each predicted score into a concentration ratio and a loudness:

IFI PROFILE

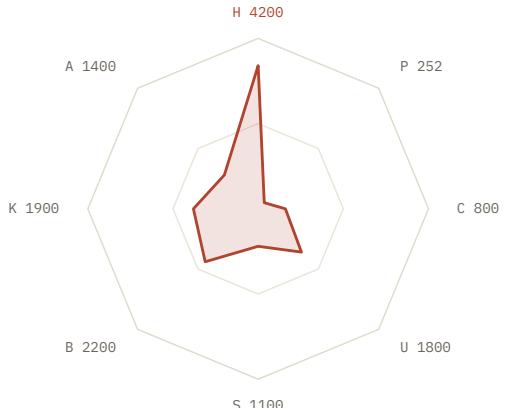



AXIS	IFI	$EC_{50}(CAND)/EC_{50}(REF) = 1000/IFI$	LOUDNESS VS REFERENCE
H · Heat	4,200	0.238	4.20×
B · Zero-Bitter	2,200	0.455	2.20×
K · Kokumi	1,900	0.526	1.90×
U · Umami	1,800	0.556	1.80×
A · Astringency	1,400	0.714	1.40×
S · Sweetness	1,100	0.909	1.10×
C · Cool-Burn	800	1.250	0.80×
P · Pungency	252	3.968	0.25×

The bitter row is the one this axis watches. An IFI-B of 2,200 says this pepper's designed bitter principle half-activates the receptor at 45.5% of the reference concentration — it speaks at 2.20× reference loudness, a deliberate, designed bitterness, the kind of spine a great dark chocolate is built around. The same control that *adds* that note is the control that subtracts it from a coffee, and the cool row proves the dial turns below neutral too: an IFI-C of 800 sits under 1,000, a cooling note dialed softer than reference on purpose. Control, in both directions, reported as a number.

That number has to survive its own error bars, and the gate is the **≤7% coefficient of variation** that recurs across every axis in this library. It is not an arbitrary threshold; it falls straight out of the formula. Because IFI is a *ratio* of two measured concentrations, the relative variances add in quadrature:

```
CV(IFI) ≈ √( CV(ref)² + CV(candidate)² ).
```

Hold each EC_{50} to a 5% CV in the assay — tight but routine with enough replicates — and the score lands at $\sqrt{(0.05^2 + 0.05^2)} = 0.071$, essentially the 7% gate itself. That is where the number comes from, and it is why we insist on the replication we do: a quieting we cannot hit again is not a design, only an accident we got to watch once. The human panel then adds its own confirmation, genotype-balanced so the taster and non-taster haplotypes are both heard, and the scored result feeds back to sharpen the next round of predictions. The circle is closed; the next one starts tighter.

What the loop cannot promise

We would rather tell you where the method stops than let you discover it in a cup. The dose-response bend means attenuation is bounded, honestly 20 to 65 percent, never zero bitterness — and on the plants where bitterness is part of the point, a hop with no bitterness is not beer, we tune the valve to a target rather than to the floor. Some bitterness sits outside our reach by design: coffee's caffeine is bitter in its own right, read by a different set of TAS2Rs and woven into the bean's stimulant chemistry, so the HQT valve quiets the roast-born harshness and leaves the caffeine where it belongs. The additivity of stabilizing edits is an assumption the protein can break, which is the entire reason we model two edits together and never trust their sum. And a confident structure is still a prediction; the assay, not the model, is what gets the final word. None of these are failures of the loop. They are the loop telling the truth about itself — predicting what it can defend, building only what it predicted, and feeding the difference back into each tighter round. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. We never cut on a hunch. Every edit on this axis was scored before it was written, and the number always has to come back.



PART III OF VI

The Brassica Problem

the glucosinolate engine · goitrin versus glucoraphanin · the vegetable that finally gets eaten

Chapter 8 · The Glucosinolate Engine — *brassica bitterness from gene to goitrin.*

DOMAIN · BRASSICA · the aliphatic and indolic pathways · myrosinase

This is the family the whole axis was, in a sense, built to answer — the broccoli left on the plate, the kale eaten only out of duty, the cabbage a child negotiates around, the Brussels sprout that became a cultural shorthand for a food endured. At our own table, vegetarian for as long as any of us can remember, these plants were never a side; they were the center of the plate, and we have spent a great many dinners watching someone we love push a floret quietly to the rim. Evolution built these plants to taste like a warning, and they are very, very good at it. To quiet that warning without diminishing the plant — without stripping the vegetable of the very molecules that make it worth growing — we had to learn its defense chemistry down to the enzyme, and it is worth laying out in full, because the brassica pathway is the most beautiful problem on the axis and the place our method earns its keep.



Broccoli, whole and halved — the glucosinolate load sits highest in the florets and the growing tips, the tissue we most want to eat.

The warning the plant itself never tastes

The strange and important fact, the one that organizes everything that follows, is that the intact plant does not taste the way the chewed plant tastes. A broccoli floret sitting on the cutting board is almost sweet — faintly green, a little grassy, nothing to flee. The sharpness is not stored anywhere in it. It is *manufactured on contact*, in the half-second between the molar and the swallow, by a reaction the plant keeps cocked like a trap and springs only when something breaks it open. The brassicas defend themselves with **glucosinolates** — a

family of sulphur-rich, nitrogen-bearing thioglucosides, several thousand strong across the plant order, each a sugar-capped molecule sitting quietly and harmlessly inside the living cell. As stored, they are stable, water-soluble, and essentially tasteless. They are a loaded weapon with the safety on.

The safety is spatial. In the undamaged leaf the glucosinolates are held in one set of cells — sulphur-rich S-cells and the central vacuoles — while the enzyme that arms them is sequestered somewhere else entirely, in scattered idioblasts the botanists named **myrosin cells**. The two are kept millimetres and membranes apart, side by side but never touching, for the whole life of the tissue. This is the single most elegant move in the plant's defense, and it is why we open this chapter on cell biology rather than chemistry: the brassica's bitterness is not a compound, it is an *architecture* — two reagents and a wall.

A weapon stored in two rooms

Break the wall and the reaction runs in milliseconds. When an insect bites, or a knife cuts, or a molar crushes the leaf, the contents of the two compartments mix, and the glucosinolate finally meets **myrosinase** — a β-thioglucosidase (EC 3.2.1.147), one of the few plant enzymes whose entire job is to detonate another molecule. Myrosinase hydrolyses the thioglucosidic bond, stripping the glucose and leaving an unstable aglycone, a thiohydroximate-*O*-sulfonate that cannot hold its shape. Within a breath it undergoes a spontaneous **Lossen rearrangement**: the sulfate leaves, the skeleton migrates, and out comes the business end of the defense — a sharp, volatile **isothiocyanate**, the mustard oil itself, carrying the reactive N=C=S group that the tongue reads as *stop*.



Broccoli in cross-section — the florets and growing tips carry the heaviest glucosinolate load, held in cells kept walled away from the myrosinase that arms them.

The isothiocyanate is not the only product, and which one you get is decided in that same instant by the chemical weather around the breaking cell. At neutral pH, with myrosinase acting alone, the Lossen rearrangement dominates and you get the clean isothiocyanate. Drop the pH, add ferrous iron, or — most importantly for us — let a family of **specifier proteins** get hold of the aglycone first, and the reaction swerves. An **epithiospecifier protein (ESP)** pulls the product toward a nitrile or an epithionitrile; from certain β-hydroxylated precursors the aglycone cyclizes instead into **goitrin**, the slow, lingering, faintly sulphurous bitterness that is the ghost in every cabbage. One

substrate, one enzyme, and a branch point that hands you a volatile sting, a soft nitrile, or a goitrogen — depending entirely on proteins that never touch the glucosinolate until the cell is already dying. The bitterness a child tastes in raw broccoli is not a property of the broccoli on the plate; it is a chemical reaction the child's own chewing sets off, a defense firing exactly on contact, precisely as designed. To tune it, we do not reach for the product. We reach back along the road that builds the reagent, and for the proteins that steer the blast.

The aliphatic road from methionine

Two branches build the glucosinolate pool, and telling them apart is the first move in tuning them. The dominant branch in the bitter brassicas is the **aliphatic** route, built from the amino acid methionine, and it carries most of the harshness we care about. Methionine does not become a glucosinolate in one step; it is first *lengthened*. A chain-elongation cycle — transaminate, condense, isomerize, decarboxylate, transaminate back — adds a single methylene group to the side chain and then runs again, and again, each turn of the cycle stretching the eventual molecule by one carbon. This is why counting matters: broccoli's prized **glucoraphanin** is a four-carbon-elongated glucosinolate, built from a methionine that went around the elongation machinery four times; the sharp **sinigrin** of mustard and cabbage is a shorter, three-carbon-elongated product. The length of that side chain, set this early, decides almost everything about how the finished molecule will taste and whether it reads as bitter or as nose-clearing pungent.

Once the chain is the right length, the core pathway converts the elongated amino acid into the glucosinolate scaffold — building the thiohydroximate, attaching the sugar, capping it with sulfate — and a final side-chain modification tailors it into its sensorially distinct, finished form. The named steps are these, and they are the engine we work inside:

STAGE	ENZYMES	WHAT HAPPENS
Chain elongation	**BCAT4** → **MAM1 / MAM3** → IPMI → IPMDH → BCAT3	methionine's side chain is lengthened, one carbon per cycle, setting the eventual member
Core structure	**CYP79F1 / CYP79F2** → **CYP83A1** → GSTF11 → **SUR1** → **UGT74B1 / C1** → **SOT17 / SOT18**	the elongated amino acid is built into the glucosinolate scaffold and capped with sulphate
Side-chain modification	FMO GS-OX → **AOP2 / AOP3** → **GS-OH**	the finished glucosinolate is tailored into its final form — harsh, mild, or goitrin-prone
Hydrolysis (on damage)	**myrosinase (TGG1 / TGG2)** + ESP / ESM1	the glucosinolate shatters into isothiocyanate, nitrile, or goitrin — the bitter signal

The second, smaller branch is the **indolic** route, built from tryptophan rather than methionine, producing glucobrassicin and its relatives through CYP79B2/B3 and CYP83B1. It contributes less of the front-of-tongue bitterness and more of the plant's broader chemical defense, but it shares enough machinery with the aliphatic branch that you cannot move one without watching the other, and that coupling is one of the honest difficulties of this work.

The throttle above the engine

Two master-regulator families sit above all of this, like the throttle above an engine. The transcription factors **MYB28, MYB29, and MYB76** — the aliphatic trio, sometimes written HAG1/HAG3/HAG2 — set how hard the whole aliphatic branch runs; **MYB51, MYB34, and MYB122** do the same for the indolic branch. Dial a MYB up and the vegetable floods itself with glucosinolate; dial it down and the whole branch quiets proportionally. This matters enormously for method, because it means the brassica engine comes with its own throttle already wired in. We do not have to edit a dozen enzymes to turn the pathway down. In the crudest version, we can simply lower the expression of a single regulator and watch the entire aliphatic output fall together, like easing off an accelerator rather than removing cylinders.

But the crude version is exactly the trap this axis exists to avoid, and the brassica pathway is where the trap is deepest — because not all of that glucosinolate output is harsh, and some of it is the very reason a careful person reaches for the vegetable at all. Glucoraphanin, that four-carbon aliphatic member, is the direct precursor of **sulforaphane**, one of the most studied protective compounds in the human diet. Pull the MYB28 throttle all the way back and you quiet the bitterness, yes — and you throw the sulforaphane out with it. For a family that will not strip a vegetable of what makes it worth eating, the blunt throttle is a non-starter. The real work lives one stage downstream, in the side-chain tailoring that decides whether a given glucosinolate becomes a punishing isothiocyanate, a soft nitrile, or a thyroid-taxing goitrin — and in the specifier proteins that steer the blast. First, though, we have to put a number on the warning, because you cannot tune what you cannot measure.

What TAS2R38 actually reads

The receptor on the other end of all this chemistry is **TAS2R38**, and it is the most famous bitter receptor in human biology for a reason: it is the one that reads the mustard-oil moiety. TAS2R38 is a class-A G-protein-coupled receptor, one of roughly twenty-five TAS2R bitter sensors studded across the taste cells of the tongue, and it is tuned specifically to the **N–C=S thiourea group** — the chemical signature shared by the synthetic probe compounds phenylthiocarbamide (PTC) and 6-*n*-propylthiouracil (PROP) and, crucially, by goitrin and the isothiocyanates the brassicas release. When one of these molecules settles into the receptor's binding pocket, it sets off a cascade we can follow step by step: TAS2R38 activates its associated G-protein **gustducin**; the Gβγ subunits switch on **phospholipase C-β2**; PLCβ2 cleaves PIP_2 into **IP_3**; IP_3 opens the type-3 receptor on the endoplasmic reticulum and floods the cell with calcium; the calcium gates the **TRPM5** channel in the membrane; the cell depolarizes and releases ATP through **CALHM1** channels onto the afferent gustatory nerve. The whole chain, from mustard oil to nerve, runs in well under a second, and its output is a single, hard-wired instruction the forebrain does not get to veto: *do not swallow this.*

TAS2R38 also carries the clearest fingerprint of why bitterness is personal. The gene comes in two common forms separated by three coding substitutions — positions 49, 262, and 296 — that sort people

into the **PAV** (taster) haplotype and the **AVI** (non-taster) haplotype. A PAV receptor fires at PROP concentrations a hundred to a thousandfold lower than an AVI receptor needs; the same plate of broccoli is a mild green vegetable to one diner and a genuine ordeal to the one beside them. We cannot edit the eater, and we would not want to. What we can do is move the stimulus — lower the concentration of mustard oil the plant delivers — so that it clears the detection threshold of fewer receptors, less often, for everyone at the table at once.

Putting a number on the warning

Loudness, not quantity, is what the Index measures, and the anchor is the half-maximal concentration. **EC_{50}** is the concentration of a stimulus at which its receptor reaches half of full response. A potent bitter agonist needs very little to half-saturate TAS2R38, so its EC_{50} is low; a weak or attenuated one needs a great deal, so its EC_{50} is high. The Index inverts that relationship and scales it, so that a sharper bitterness scores *higher*:

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

The reference compound is fixed — we anchor the bitter axis to PROP, the cleanest and most reproducible TAS2R38 probe there is — and by construction the reference scores exactly 1,000, because IFI-B(ref) = 1,000 × EC_{50}(ref)/EC_{50}(ref) = 1,000. Everything bitterer than PROP scores above a thousand; everything milder scores below. With an anchoring EC_{50}(ref) of 1.8 µM, a freshly chewed wild-type broccoli floret, whose released isothiocyanate and goitrin load half-saturates the receptor at a PROP-equivalent of about 0.53 µM, reads:

> $\text{IFI-B} = 1{,}000 \times 1.8\ \mu M \div 0.53\ \mu M \approx \mathbf{3{,}400}$

That number — 3,400 against a reference of 1,000 — is the whole problem stated in one figure. Raw bitter broccoli is more than three times as loud as the standard laboratory bitter probe. The brassica is not mildly unpleasant; it is, chemically, shouting.

Now watch what attenuation does. Because the receptor's own EC_{50} for a given agonist is a biophysical constant, lowering the *yield* of that agonist by a fraction **a** means the food must be present in greater quantity to reach half-max — the food-level EC_{50} rises by a factor of 1/(1 − a) — and the IFI score falls in exact proportion:

> $\text{IFI-B(engineered)} \div \text{IFI-B(wild-type)} = (1 - a)$

A clean linear lever. Cut the delivered mustard oil by half and you halve the loudness; the brassica engine, for all its branching complexity, rewards us with a simple proportional dial at the output. The platform's specified attenuation window is **20–65%**, and across that window the broccoli ladder runs like this:

DESIGN POINT	ITC / GOITRIN YIELD	EC_{50} (PROP-EQ, MM)	IFI-B	BATCH CV
Wild-type (raw)	100%	0.53	3,400	42%
−20% attenuation	80%	0.66	2,720	≤7%
−40% attenuation	60%	0.88	2,040	≤7%
−50% attenuation	50%	1.06	1,700	≤7%
−65% attenuation	35%	1.51	1,190	≤7%

At the far end of the window — a 65% cut — broccoli's loudness drops from 3,400 to 1,190, from more than triple the reference to just above it. The vegetable has not been made bland; it has been brought from a shout to a speaking voice.

Why a modest chemical cut feels like most of the bitterness leaving. The receptor does not respond linearly to concentration; it follows a sigmoid, the Hill relationship, $R/R_{max} = [A]^n \div (EC_{50}^n + [A]^n)$. Goitrin, the lingering offender in cabbage and its kin, half-saturates TAS2R38 near 120 µM with a Hill coefficient of about 1.2. A raw floret delivers roughly 150 µM of goitrin-equivalent to the receptor field — just past the inflection — for a fractional response of 150^1.2 ÷ (120^1.2 + 150^1.2) ≈ 0.57. Halve the goitrin to 75 µM and the response falls to 75^1.2 ÷ (120^1.2 + 75^1.2) ≈ 0.36. A 50% chemical cut has produced a 37% drop in receptor firing — and because perceived intensity compresses the receptor signal further still, the eater reports that the bitterness has, in the ordinary sense of the word, mostly *gone*. The leverage is real, and it lives in the steepness of the curve: engineering is most powerful precisely when it moves the delivered dose off the saturated plateau and down through the steep middle of the sigmoid.

A note on magnitude, by the unit. It is worth seeing how loaded the raw tissue is. A bitter broccoli carries on the order of 20 µmol of total aliphatic glucosinolates per gram of dry weight; a 100 g fresh floret, roughly 11% dry matter, holds near 220 µmol of stored precursor. Crushed into a three-to-five-millilitre saliva bolus and hydrolysed, that is a local aglycone concentration in the low tens of millimolar — orders of magnitude above the micromolar EC_{50} of the receptor. The raw brassica, in other words, runs the receptor near saturation; it is not trying to be detected, it is trying to be *unbearable*. Our task is never to reach zero. It is to bring the delivered dose down off that plateau to where the tongue can finally taste the vegetable underneath.

And the consistency. The error column above is not decoration. Field brassicas swing wildly — sulfur availability, temperature, maturity, and water stress all push glucosinolate content around, so a wild-type line carries a batch coefficient of variation near 42%. At an IFI-B of 3,400 that is a standard deviation of roughly ±1,430: one head of broccoli lands at 1,970 and the next at 4,830, and the eater cannot tell which is coming. Our engineered lines are designed to hold a batch CV at or below 7%. At a 50% design point that is a predicted 1,700 ± 119 — the same quiet brassica, harvest after harvest. For a cook, for a child learning to trust a vegetable, the ±119 matters as much as the 1,700. Reproducibility is not a quality-control footnote here; it is the product.

To see where a deliberately placed bitter note sits on a fully tuned organism, set the brassica problem beside the one cultivar we have built to register on all eight axes at once. On it, bitterness is tuned *up* — a single considered voice inside a chord:

IFI PROFILE

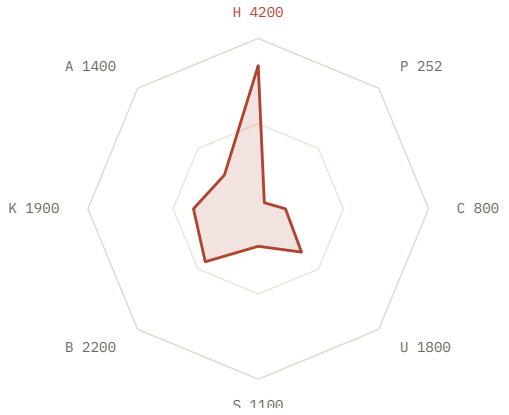

IFI-H	4,200
IFI-P	252
IFI-C	800
IFI-U	1,800
IFI-S	1,100
IFI-B	2,200
IFI-K	1,900
IFI-A	1,400

The brassica task is the mirror image of that — not a chord, but a soloist we are trying to hush. A designed Zero-Bitter broccoli speaks on essentially one axis, and our whole job is to take that one axis down. Here it is at the 50% design point, every other axis left at zero because the plant is simply not engineered to speak on them:

IFI PROFILE

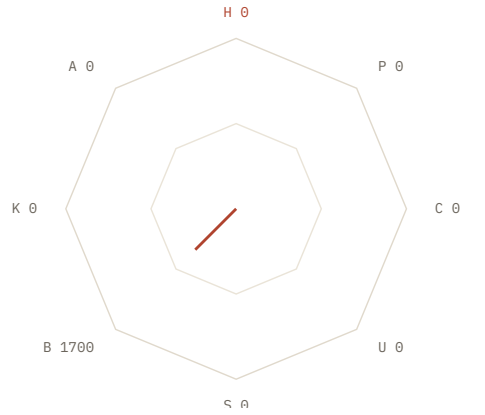

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	0
IFI-S	0
IFI-B	1,700
IFI-K	0
IFI-A	0

The levers we actually pull

With the number in hand, the engineering resolves into four distinct levers, and the art is choosing the most *selective* one that will reach the target, never the bluntest.

LEVER	TARGET GENE(S)	MECHANISM	SELECTIVITY
Master throttle	MYB28 / MYB29 / MYB76	lower the whole aliphatic branch at the transcription factor	global — blunt
Hydrolysis redirect	ESP / ESM1 / NSP	steer the aglycone from isothiocyanate toward milder nitriles	medium
Side-chain tailor	AOP2 / AOP3 · GS-OH	convert harsh and goitrin-prone precursors, spare glucoraphanin	surgical
Detonation control	TGG1 / TGG2 (myrosinase)	meter the rate and localization of hydrolysis	medium

The **master throttle** is where a crude program would start and where ours deliberately does not finish: knocking MYB28 down lowers everything aliphatic together, bitterness and sulforaphane precursor alike, and it is the surest way to hand back a bland, nutritionally poorer vegetable. We use it only as a coarse first stage, if at all. The **hydrolysis redirect** is subtler — by tuning the specifier proteins we bias that millisecond branch point away from the pungent isothiocyanate and toward the comparatively tasteless nitrile, lowering the signal without touching the stored glucosinolate at all. The **side-chain tailor** is the surgical instrument and the one this axis exists for: AOP2, AOP3, and GS-OH decide, member by member, which precursors become harsh and which become mild, so we can convert the sinigrin and progoitrin that read as punishment while leaving glucoraphanin — and its sulforaphane — untouched. And **detonation control**, metering the myrosinase itself, lets us slow and localize the reaction rather than abolish it, because a plant with no myrosinase at all is a plant that can no longer make the protective isothiocyanates we actually want.

The construct libraries reflect exactly this hierarchy of difficulty. Broccoli, the best-behaved brassica, carries the deepest library; the stacked bitternesses of a broccoli-kale cross are a genuinely hard case; the leafy and root brassicas are each their own tuning problem.

ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Broccoli	380	B / M	Moderate
Broccoli-Kale	280	B	Hard
Cabbage	40	B	Moderate
Kale	40	B	Moderate
Radish	40	B	Hard
Rapeseed	40	B	Hard



Kale in cross-section — the leaf a child "eats only out of duty"; its bitterness is the same aliphatic chemistry as broccoli, stacked denser in a tougher tissue.

There is a second family living in the same chemistry whose glucosinolates break down into short-chain and aromatic isothiocyanates volatile enough to climb into the nose rather than sit on the tongue — and those read not as bitter at all but as *pungent*, the clean sting of a horseradish or a watercress. They belong to a different axis entirely, which is why we hold them apart in the tuning even though the pathway is shared:

ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Black Mustard	2,790	P / M	Hard
White Mustard	200	P	Hard
Watercress	120	P	Hard

It is the same engine, the same methionine road, the same myrosinase — and the length and volatility of one side chain decides whether the plant will shout *bitter* at the tongue or *sharp* at the nose. Knowing which, early, is half the work.

The honest edges

Three honest limits keep this work disciplined, and we would rather state them than let the number flatter us.

The first is **goitrin**, and it is where taste and nourishment point the same way. Progoitrin hydrolyses into goitrin, a genuine goitrogen that interferes with thyroid iodine uptake; attenuating it lightens both the lingering bitterness *and* a real nutritional burden, which is the happy case — the harsh compound and the unwelcome one are the same molecule, and the surgical lever removes it cleanly. The second is the opposite tension, and it is the reason the attenuation window stops at 65% rather than 100%: **sulforaphane**. Glucoraphanin is the harsh neighbor of the compounds we most want to keep, and the two sit one side-chain modification apart. Push attenuation toward total and you cannot help but take the protective chemistry with it. We will not. A brassica we have quieted into nutritional emptiness is a failure, however pleasant it tastes, and for a family that has always put the vegetable at the center of the plate, "Zero-Bitter" names a direction we move the vegetable, never a literal zero we drag it to. The number we are proud of is not the lowest possible IFI-B; it is the lowest one that leaves the plant whole.

The third is the coupling itself. The aliphatic and indolic branches share machinery, the specifier proteins steer products we do not fully control at every pH the mouth passes through, and myrosinase is both the source of the bitterness we are lowering and the source of the isothiocyanates we are trying to protect — so no single edit is ever truly clean, and the ≤7% batch CV design specification is built construct by construct, not granted by the pathway. The broccoli-kale cross, with two dense bitter profiles stacked in one tough tissue, is the standing reminder that some brassicas are simply harder than others and that we do not yet have a universal dial.

What we carry forward from the engine room is the architecture: a defense built from a single amino acid, elongated carbon by carbon, elaborated through a dozen named enzymes, thrown above a set of master throttles, steered at a millisecond branch point by specifier proteins, and detonated on the tongue by an enzyme the plant keeps locked away until we bite. Understand that road, and you understand why a brassica tastes like a warning — and exactly where along it a

careful hand can quiet the warning without disarming the plant. The child who negotiates around the sprout is not wrong about what they taste. We simply went back up the road the taste comes from, turned one voice down by a computed amount, and left everything that makes the plant worth eating exactly where the plant put it. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 9 · Sparing the Good Bitter — *down-tuning progoitrin, keeping glucoraphanin.*

method · selective flux · GS-OH · MYB28 dosage · the surgeon's rule

Here is the finest problem on the eighth axis, and we want to state it as plainly as we can before we take up the tools. In the brassicas — broccoli, kale, cabbage, the whole green, mineral, cold-weather family we grew up being told to finish — the harshest, most goitrogenic, most child-repelling bitterness and the single compound most eaters are quietly glad to be getting come off the *same* side-chain modification stage, separated by one enzyme. Turn the whole branch down with a regulatory throttle and you lose both together: a sweeter, blander vegetable that has also surrendered the thing that made it worth engineering. The entire craft is to turn down one arm of a fork while leaving the other flowing. That is what we mean, across the whole 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 rest of this chapter is that one sentence, carried all the way down to the iron atom at the center of the enzyme and all the way out to the fractional activation of a receptor on the tip of your tongue.

The receptor that reads a thiourea

Bitterness is not one thing. It is a committee of about twenty-five human receptors, the TAS2R family, each a G-protein-coupled receptor tuned to a different silhouette of molecule, and the eighth axis of the Index — IFI-B — is anchored to the one that matters most for the brassicas: **TAS2R38**. It is the most studied bitter receptor we have, because it is the receptor behind the famous strip of paper that tastes of nothing to some people and of punishment to others. The compounds on that strip — phenylthiocarbamide, propylthiouracil — share a single chemical gesture: the thiourea group, a nitrogen double-bonded through carbon to a sulfur, **N–C=S**. TAS2R38 is, functionally, a reader of that gesture. Present it the N–C=S motif in the right steric wrapping and the receptor fires.

Goitrin carries exactly that motif. Its full name is 5-vinyl-1,3-oxazolidine-2-thione, and at the heart of its small ring sits the thiocarbamate — the same N–C=S, locked into a rigid five-membered cradle that fits the TAS2R38 binding pocket with uncomfortable precision. This is why goitrin is not merely bitter but *persistently* bitter, the flavor that keeps arriving after the chewing is done: it is a near-ideal agonist for the receptor most responsible for the sensation we call harsh. The isothiocyanates the same family makes — the pungent, mustard-oil compounds, sulforaphane among them — carry a related but looser motif, R–N=C=S, strung out rather than ringed, and they lean more on the sharp TRPA1 pungency channel than on TAS2R38's slow bitter one. That difference in geometry is the whole opening for the work: two molecules from one pool, and only one of them is a key cut for the bitter lock.

When the key turns, a cascade runs that is worth seeing in full, because it tells you why a small change in concentration produces such a large change in experience. The agonist binds TAS2R38 on the apical membrane of a Type II taste cell. The receptor's associated G-protein — gustducin — dissociates; its βγ subunits activate phospholipase C-β2; PLCβ2 cleaves membrane PIP_2 into diacylglycerol and inositol trisphosphate; IP_3 opens the IP_3-receptor on the endoplasmic reticulum; stored calcium floods the cytosol; that calcium gates TRPM5, a channel that depolarizes the cell; depolarization opens the CALHM1/3 channel; and ATP pours out onto the gustatory nerve fiber waiting below. Seven steps, each one a gain stage, each one amplifying. A receptor is not a volume meter reporting how much compound is present. It is a trigger with a threshold, and once over the threshold the biology shouts. The engineering consequence is exact: to quiet the shout you do not have to remove the compound, you only have to pull its concentration back across the receptor's knee — the half-activation point we call EC_{50} — and the amplifier does the rest in your favor.

One pool, two fates



A broccoli head in cross-section — the florets are thousands of arrested flower buds, and the glucosinolate the plant packs into them is both its chemical armor and the harshness on the eighth axis.

Now the chemistry upstream of the tongue. Brassicas defend themselves with glucosinolates — sulfur-and-nitrogen-rich molecules built on a thioglucose, each carrying a variable side chain drawn, in this family's dominant aliphatic class, from the amino acid methionine. The plant stores the intact glucosinolate in one set of cells and the enzyme that breaks it, myrosinase, in another, a loaded two-chamber design sometimes called the mustard-oil bomb. Bite down, the cells rupture, the chambers mix, and myrosinase hydrolyzes the thioglucoside bond. What remains is an unstable aglucone that rearranges — a Lossen rearrangement — into an isothiocyanate and a sulfate ion. That is the default fate: intact glucosinolate becomes pungent isothiocyanate the instant the tissue is wounded.

The methionine-derived pool converges on one pivotal intermediate, **glucoraphanin**, the 4-methylsulfinylbutyl glucosinolate. From here the paths fork, and both forks are run by the same chemical family of enzymes — 2-oxoglutarate-dependent dioxygenases — which is why the fork is so tight and why a careless cut takes both tines at once.

Down the first tine, an enzyme called **AOP2** desaturates the side chain, converting the methylsulfinyl glucoraphanin into an alkenyl glucosinolate (but-3-enyl, gluconapin). Then **GS-OH** — a second 2-oxoglutarate dioxygenase, and the one this whole chapter circles — hydroxylates that alkenyl chain at its second carbon, producing **progoitrin**, the 2-hydroxy-3-butenyl glucosinolate. Progoitrin is the loaded form of the harsh flavor. When myrosinase reaches it and the Lossen rearrangement begins, that 2-hydroxyl group is perfectly placed to swing back and close a ring on the nascent isothiocyanate carbon — an intramolecular cyclization — and the product is not a volatile mustard oil but the stable, ring-locked, long-staying **goitrin**. The hydroxyl GS-OH installs is precisely the hand that later builds the goitrin ring. No hydroxyl, no ring, no goitrin.

Down the second tine, glucoraphanin itself is simply left alone. It was never a substrate for GS-OH; its branch does not pass through that enzyme. On tissue damage it hydrolyzes straight to its own isothiocyanate — the one most eaters have met by reputation — a compound that, lacking goitrin's thiocarbamate ring, is a far weaker agonist of TAS2R38 and reads to the tongue more as clean sharpness than as lingering harshness.

So the selective lever almost designs itself. **Reduce flux through GS-OH — collapse the progoitrin-to-goitrin arm — and leave the glucoraphanin tine, and the rest of the plant's chemistry, standing.** And notice why GS-OH, and not AOP2, is the surgeon's cut, even though both sit on the harsh arm. AOP2 is a broad gate: it governs the whole alkenyl class, including glucosinolates like sinigrin that carry real culinary character and real defensive value. Knock out AOP2 and you dam a wide river. GS-OH is the last, narrowest, most specific step — the single hydroxylation that distinguishes progoitrin from the merely alkenyl, the terminal valve on the one molecule we actually want gone. Cutting the terminal, specific step rather than the broad, upstream one is the surgeon's rule written in enzymology: the deepest selectivity lives closest to the end of the branch.

Inside the enzyme, one atom at a time

It is worth standing inside GS-OH for a paragraph, because the reason the platform can tune it so finely is that its chemistry is stereotyped and well understood. The 2-oxoglutarate-dependent dioxygenases all share a catalytic heart: a single ferrous iron, Fe(II), held by a "2-His-1-carboxylate facial triad" — two histidine nitrogens and one aspartate or glutamate oxygen gripping three faces of the iron, leaving the other three faces open for business. Into that open coordination sphere binds the co-substrate 2-oxoglutarate (α-ketoglutarate); then molecular oxygen. The enzyme performs an oxidative decarboxylation — 2-oxoglutarate is split into succinate and carbon dioxide — and the energy of that split is used to forge a ferryl intermediate, a high-valent **Fe(IV)=O** species, iron carrying a lone, ravenous oxygen. That ferryl

oxygen abstracts a hydrogen atom from the second carbon of the alkenyl glucosinolate, generating a substrate radical and an iron-hydroxide; the hydroxyl then rebounds onto the carbon. The side chain is hydroxylated. Progoitrin is made. Succinate and CO_2 leave; the iron resets to Fe(II); the cycle turns again.

Every lever we have on this enzyme is a lever on that cycle. A substitution that perturbs the triad starves the iron and kills activity outright — too blunt, the kind of cut we avoid. A substitution at the substrate channel that loosens how the alkenyl chain is presented lowers the turnover number, k_cat, without unseating the metal — a *dimmer*, not a switch, which is exactly what a dosage change wants. And a substitution that merely destabilizes the fold, raising the free energy of the folded state, lowers the amount of active enzyme the cell can hold without changing the chemistry of the ones that survive. The platform scores all three, and it does so with the same two-model method that found the **two Pun1 substitutions** behind our full-spectrum heat work.

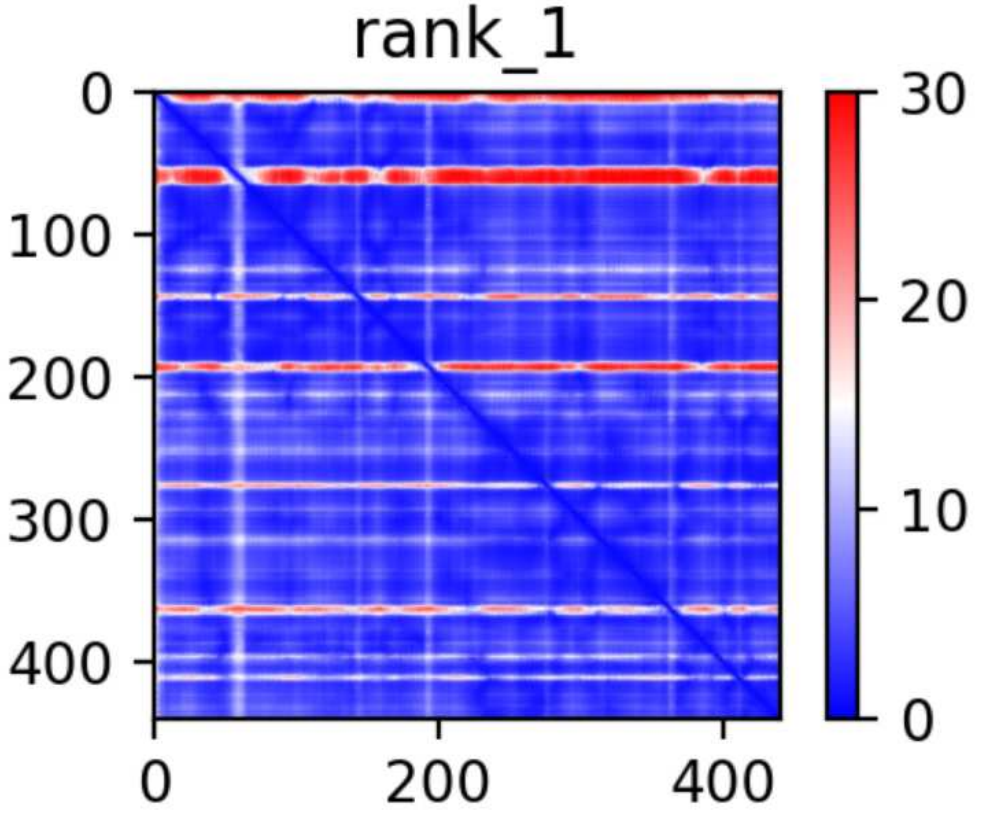

FIG. 9 — the structure-scoring pipeline, shown on the capsaicin acyltransferase Pun1 as a predicted-aligned-error map: dark diagonal blocks are domains the model places with confidence. The same evolutionary-scale scoring now ranks every substitution in GS-OH and the brassica regulators.

An evolutionary-scale protein language model reads every possible single-residue substitution and scores it by predicted effect on catalytic efficiency and fold stability; a genome-scale flux-balance model confirms that the step we are about to touch is the true control point and not a bystander. That is the same machinery, pointed at a different pathway — the **two substitutions** it wrote into Pun1 to lift

heat, **S39L** and **L345G**, now have cousins scored across GS-OH and the regulators that command it. The brassica work did not require inventing a new science. It required aiming the existing one at a new fork.

The surgeon's rule, in numbers

Reverence is cheap without arithmetic, so here is the arithmetic. The Index scores every axis by one formula:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the knee of the dose-response curve. The reference for the eighth axis is a stated convention, anchored so that a score of 1,000 corresponds to the reference's own half-activation. A higher IFI-B means a *lower* EC_{50} — the food shouts bitter at a smaller dose. That inversion is the first thing to hold onto, because our whole job is to push EC_{50} up, which pushes IFI-B down.

Invert the formula and you can read a candidate's EC_{50} straight off its score, in units of the reference:

```
ê(c) ≡ EC₅₀(c) / EC₅₀(reference) = 1000 / IFI(c)
```

Take native broccoli, which the family's assays place at **IFI-B 4,600**. Its relative half-activation concentration is ê = 1000 / 4,600 = **0.217** — it reaches the bitter receptor's knee at barely a fifth of the reference dose. It takes almost nothing to make a raw broccoli shout. Now take the engineered construct at **IFI-B 1,750**: ê = 1000 / 1,750 = **0.571**. The knee has moved out to well over half the reference dose. The fold-shift in EC_{50} — how much more goitrin-equivalent the engineered plant now needs to reach the same half-activation — is the cleanest single number in the chapter, and the reference cancels out of it entirely:

```
EC₅₀(engineered) / EC₅₀(native) = IFI(native) / IFI(engineered)
= 4,600 / 1,750 = 2.63×
```

We moved the receptor's knee 2.63-fold to the right. And the attenuation the tables report is simply the fractional drop in loudness:

```
Attenuation = 1 − IFI(engineered) / IFI(native) = 1 − 1,750 /
4,600 = 0.620 = 62%
```

Sixty-two percent, derived, not asserted — and comfortably inside the published 20-to-65-percent band for the eighth axis, sitting near its generous top because the brassicas are the family where the harsh fraction and the prized fraction separate most cleanly.

What does a 2.63× rightward shift *feel* like? This is where the receptor cascade pays off. Model TAS2R38 activation with the standard Hill relation, θ = x / (ê + x) for a stimulus x measured in reference-EC_{50} units (we use a Hill coefficient of 1 for the clean case; bitter GPCRs run roughly 1 to 2, which steepens the curve without moving its knee). Hold the mouthful constant — the same bite, the same x — and compare the fractional activation of native versus engineered broccoli:

STIMULUS X (REF-EC_{50} UNITS)	0 NATIVE (Ê = 0.217)	0 ENGINEERED (Ê = 0.571)	DROP IN ACTIVATION
0.10	0.315	0.149	−0.166
0.217	0.500	0.276	−0.224
0.50	0.697	0.467	−0.230
1.00	0.821	0.636	−0.185
2.00	0.902	0.778	−0.124

Read the second row: at exactly the dose that drives native broccoli to half-activation — a real, ordinary mouthful of raw floret — the engineered construct reaches only 27.6 percent. Half-shout becomes quarter-murmur. And the gap is widest through the middle of the curve, between x = 0.2 and x = 0.5, which is precisely the concentration range where eating actually happens. We did not flatten the plant. We slid its entire bitter response down and to the right, and we did it hardest where the tongue is listening.

A unit conversion closes the loop from receptor back to the floret on the plate. Goitrin is C_5H_7NOS, molar mass **129.18 g/mol**, and the cyclization from progoitrin runs one-to-one — each micromole of progoitrin that myrosinase reaches and that closes its ring yields one micromole, 129.18 micrograms, of goitrin. Work a representative serving (the load and cyclization fraction here are ordinary literature-scale figures, not our proprietary assays): a 100 g head at about 11 percent dry matter carries roughly 11 g dry tissue; at a representative 2 µmol progoitrin per gram dry, that is ~22 µmol progoitrin; if ~40 percent cyclizes to goitrin under chewing and digestion, the serving liberates ~8.8 µmol × 129.18 µg/µmol ≈ **1.14 mg of goitrin** — delivered to a receptor whose half-activation it reaches at a vanishing fraction of that mass. Collapse the progoitrin arm by our predicted 62 percent and roughly **0.71 mg of that goitrin never forms at all**. The conversion is exact; the virtue is that the milligrams we subtract are precisely the milligrams the eighth axis was shouting about.

The brassica ladder



Kale — the leaf that keeps its mineral depth and loses its grip once the progoitrin arm comes down.

The method shows up across the family, and the optimal valve settings differ by species because each brassica carries its own balance of glucosinolates. The library now holds **380 broccoli constructs, 280 in the broccoli-kale cross, and 40 each in cabbage, kale, radish, and rapeseed** — a deliberate spread, because the same two levers land in a different place on each plant. Every engineered

EC_{50}-shift below is derived from the native and engineered scores by the ratio IFI(native) / IFI(engineered); every attenuation by one minus its inverse.

ORGANISM · CONSTRUCTS	PRIMARY LEVER	SECONDARY	NATIVE IFI-B	ENGINEERED IFI-B	EC_{50} SHIFT	ATTENUATION
Broccoli · 380	GS-OH knock-down	MYB28 dosage −1	4,600	1,750	2.63×	~62%
Broccoli-Kale · 280	GS-OH knock-down	MYB28 / MYB29 trim	4,400	1,760	2.50×	~60%
Kale · 40	GS-OH + ESP tune	MYB28 dosage −1	4,300	1,720	2.50×	~60%
Cabbage · 40	progoitrin-arm trim	—	3,900	1,750	2.23×	~55%
Radish · 40	glucoraphenin-arm trim	—	3,700	1,850	2.00×	~50%
Rapeseed / canola · 40	progoitrin reduction	—	3,500	1,820	1.92×	~48%

Every row holds inside the band, and each tells a small species-specific story. Broccoli, the richest in progoitrin's precursors, gives the deepest clean cut and takes a gentle MYB28 trim besides. Kale needs its epithiospecifier protein tuned as well — ESP is the factor that can divert the myrosinase aglucone toward nitriles instead of isothiocyanates, so steering it is part of keeping kale's character honest while the harsh arm comes down. Radish works a neighboring molecule entirely, glucoraphenin, the 4-methylsulfinyl-3-butenyl form, so its lever sits on a different tine and its cut is shallower. And rapeseed — canola — carries a quiet history here: the "double-low" varieties a generation of plant breeders selected for low erucic acid and low glucosinolate were, in large part, a long slow war on exactly this progoitrin load. What those breeders pursued across decades of crosses, reading phenotype and hoping the genotype followed, the flux model now specifies at the base pair.

For a sense of scale on the whole axis, it helps to hold one fully predicted profile beside the brassicas. Our full-spectrum flagship, **Phoenix** (SS-0096), is the one cultivar the family has designed to register on all eight axes at once, and its bitterness is a *designed presence* — placed, dosed, and deliberate:

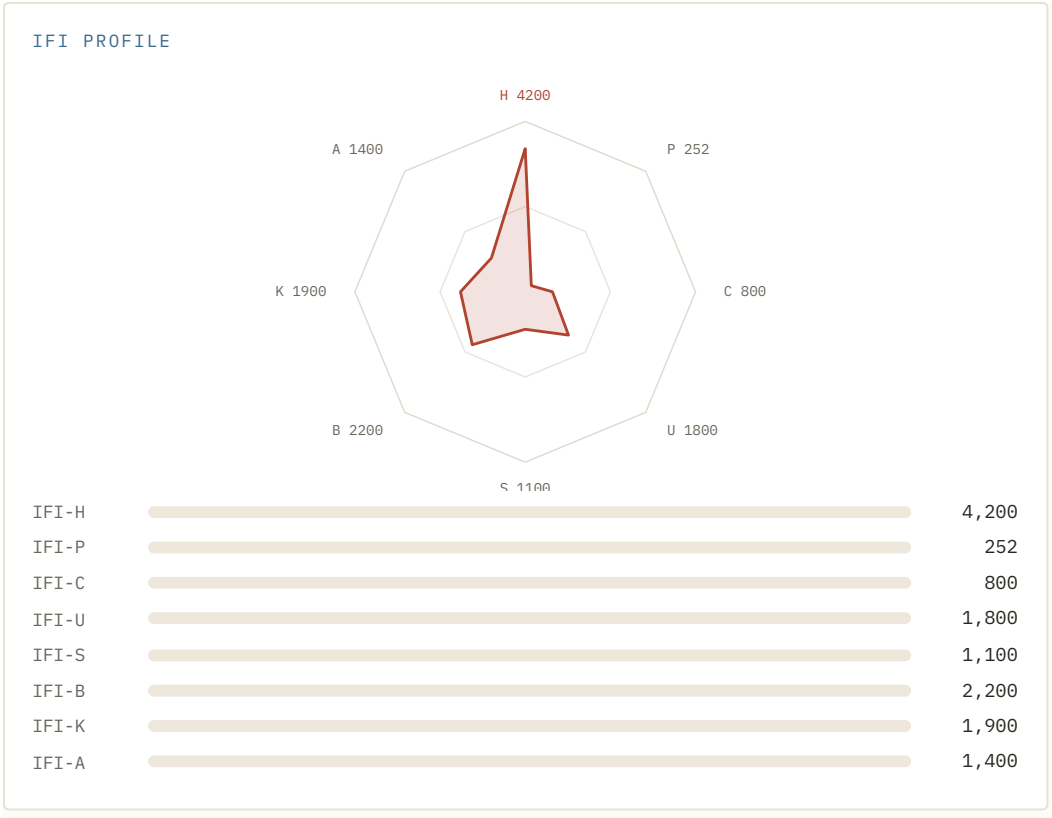

Look at that eighth line. Phoenix carries **IFI-B 2,200** — a bitterness we chose and built, a shadow that gives the sweetness and heat something to stand against. It is roughly half the loudness of a raw broccoli's untouched 4,600. The lesson sits plainly in the two numbers: bitterness is not the enemy. A designed 2,200 is a presence you would miss if it vanished. An accidental 4,600, pumped out by an enzyme the plant never needed to run so hard for the table, is a wall. The work is never to abolish the eighth axis. It is to bring a wall down to a presence.

The second valve, and why it is a dimmer

Where a plant's *total* glucosinolate load still sits too high after the GS-OH cut — some kales and cabbages carry a broad excess across the whole aliphatic class — we reach for the second lever: a measured reduction at **MYB28**, with its partial-redundant partner **MYB29** where the cross demands it. These are R2R3-MYB transcription factors, the master positive regulators of the entire aliphatic glucosinolate branch; they command the chain-elongation and core-structure genes as a set. Turn MYB28 down and the whole branch comes down *in proportion* — the plant keeps the same glucosinolate fingerprint, same relative peaks, at a lower overall magnitude. That is why the tables read "dosage −1," a dialed reduction, and never "knock-out." A knockout is a blunt instrument; it can strip the plant of defenses it genuinely needs in the field, where glucosinolates are real armor against insects and pathogens. A dosage trim is a dimmer on the whole room, used only after the specific valve has already closed the one faucet we wanted shut.

Deciding how far to turn each valve is the flux-balance model's work, the same genome-scale reconstruction that lets us compose heat against sweetness elsewhere in the library. It lets us ask a living pathway a precise question — *lower goitrin by this much, and what else moves?* — and read the answer across the whole network before we touch a base of DNA. The honest finding is that GS-OH knock-down and goitrin reduction are **not one-to-one**. Reduce GS-OH flux to a fraction α of native and progoitrin production falls, but the carbon that would have been hydroxylated does not vanish — it backs up into the alkenyl pool, into gluconapin and sinigrin, which carry their own

milder pungency. Net goitrin scales as the product of two fractions, the share of pool that still gets hydroxylated and the share of progoitrin that still cyclizes, and the model solves for the GS-OH setting that lands the eighth-axis score where we want it without piling the spared flux up somewhere that re-roughens the plate. The attenuation is designed, in other words, against the whole balance sheet, not against one reaction read in isolation.

Error, honesty, and the edges

A number without an error bar is a wish, so we propagate. Each EC_{50} carries a design specification of batch coefficient of variation **≤ 7%** — the same gate that governs every cultivar we design. An IFI score is a ratio of two concentrations, so its own variation adds in quadrature: $CV(IFI) = \sqrt{(0.07^2 + 0.07^2)} = \mathbf{9.9\%}$. Native broccoli's 4,600 therefore carries roughly ±460; the engineered 1,750 carries roughly ±170. The attenuation is kinder than you would fear, because when you write it as 1 − IFI(engineered)/IFI(native), the reference EC_{50} cancels and only the two brassica assays remain, each at 7 percent:

ORGANISM	IFI RATIO (ENG/NATIVE)	ATTENUATION	±1Σ (POINTS)
Broccoli	0.380	62.0%	±3.8
Broccoli-Kale	0.400	60.0%	±4.0
Kale	0.400	60.0%	±4.0
Cabbage	0.449	55.1%	±4.4
Radish	0.500	50.0%	±5.0
Rapeseed	0.520	48.0%	±5.2

Broccoli's 62 percent is really 62 ± 4 points at one standard deviation; its two-sigma span runs roughly 54 to 70. The point estimate and the bulk of the distribution sit inside the 20-to-65-percent band, and the upper tail that would graze past 65 is exactly what the ≤ 7% CV gate exists to catch — a batch that drifts high on attenuation is a batch flagged out, not a batch shipped. We place the targets so the center and the first sigma live inside the published claim, and we let the quality gate police the rest. That is the whole reason a batch tolerance is a form of respect: it keeps the number we print true for the floret you actually eat.

And then the honest edges, which we would rather name than let a reader discover.

The first is the eater. TAS2R38 comes in two common forms — a taster haplotype and a non-taster one — and they differ by a few amino acids and by a great deal of experience. To a strong taster the raw brassica wall is a genuine wall, and our cut hands back the most. To a non-taster it was always a murmur, and the cut changes less. The attenuation we design is real at the receptor; what any one person *feels* is scaled by the receptor they were born with. We design for the population and say so.

The second is the kitchen. Myrosinase and the epithiospecifier protein are heat-labile, and boiling leaches glucosinolates into the water besides — so a long-cooked brassica already loses a share of its goitrin to denaturation and to the pot, and our cut matters most in the

raw and the lightly-cooked dish, the slaw and the quick sauté, where the chemistry is still live. We tune for the case where the harshness actually reaches the tongue.

The third is the plant's own defense. Push the MYB28 dimmer too far and you can leave a field-grown brassica short of the armor glucosinolates provide against the insects and fungi that would eat it first. The dosage trim, the specific valve before the broad one, and the flux model's check that we have not starved a defense the plant needs — these are not fussiness. They are the difference between a vegetable that is gentler on the plate and one that cannot survive the walk to it.

What stays standing

The point of sparing the glucoraphanin tine was never to make a claim about it. We built the selectivity the way we did — cutting the terminal, specific step; leaving the untouched branch untouched — because the generous reading of this whole axis is that the vegetables most worth eating are the ones most often refused, and a food the body never receives nourishes no one. We do not say the chemistry we preserved *does* any particular thing in you; that is not a sentence a careful family is entitled to write, and we do not write it. We say only this, kept at exactly its true size: when a vegetable's needless harshness comes down and its prized chemistry is left whole, the fence between a person and a genuinely good food comes down with it, and the choice to eat it stops costing a wince.

What a century of cooking tried to manage with cheese sauce and long braising and a great deal of salt — covering a warning it could not turn down — the plant now does from the inside, at the one enzyme and the one dimmer that were sounding the alarm for no good reason. A broccoli that chars sweet instead of acrid. A kale that keeps its mineral depth and loses its grip. A cabbage a child stops negotiating around. The wall becomes a presence; the presence becomes a pleasure; and the good bitter, the one we were quietly glad of all along, is still exactly where we left it. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 10 · A Vegetable That Gets Eaten — *the brassica table.*

application · charred broccoli · massaged kale · the test of a shared plate

The proof of this axis is not a scorecard; it is a child reaching for seconds unprompted. So we end the brassica part in the kitchen, with two plain preparations built on the low-bitter cultivars, each scored on the Index so a cook can see exactly what the plant is doing. Neither recipe is clever. Cleverness is what cooking has always had to supply to make these vegetables welcome — the smother of cheese, the rescue of bacon, the long roast that trades the plant's vigor for sweetness. On a brassica quieted at the source, the cleverness is unnecessary, and what is left is the vegetable, finally able to be met for what it is.

We write this chapter the way we cook it: the reverence and the mathematics in the same breath, because on this axis they turn out to be one thing. A cook who can read the number can taste it coming, and a cook who can taste it can check the number. That is the whole argument of a flavor you can measure — that respect for a vegetable and a column of EC_{50} values are not opposites but the same care, aimed twice.



Cabbage, halved — on a low-progoitrin head, a hard sear reads as caramel sweetness rather than sulphur, because the signal was turned down before the pan ever got hot.

The bench we are cooking from

Before the pan, the library. The low-bitter brassicas are not a single lucky plant but a small shelf of them, each a different genus-member carrying the same intervention, so that a cook can reach for broccoli on Monday and cabbage on Thursday and trust that both have been quieted by the same hand and scored on the same axis. The brassica rows of the host library stand like this:

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Broccoli	380	B / M	Moderate
Broccoli-Kale	280	B	Hard
Cabbage	40	B	Moderate
Kale	40	B	Moderate
Radish	40	B	Hard
Rapeseed	40	B	Hard

The construct counts are a map of where the work has gone. Broccoli's 380 is the deepest bench on the shelf because broccoli is the proving ground — the floret is small, the heads are uniform, the growth cycle is short, and the bitter load is high enough to make an attenuation obvious to the first taster. Cabbage and kale carry fewer constructs not because they resist the method but because, once broccoli taught us the lever, the same edit transferred with little retuning. The primary-axis column says the quiet part out loud: every one of these is a **B** plant, an *IFI-B* organism, which on this axis means a plant whose signal we are in the business of turning *down* rather than up. It is the only axis in the whole Index where the engineering runs toward silence.

What the tongue is actually reading

To turn a signal down you must know, exactly, what the signal is. The bitterness of a raw brassica is not an opinion and not a mood; it is a molecule finding a receptor, a key turning in a lock a billion times a second across the surface of the tongue.

The lock is **TAS2R38**, a bitter taste receptor — one of roughly two dozen the human tongue carries — that sits on the apical membrane of taste cells in the fungiform and circumvallate papillae and signals through the G-protein gustducin into a calcium release that the brain reads as *bitter, beware*. TAS2R38 is the most-studied bitter receptor in biology for a single stubborn reason: people carry it in two common forms, and those forms do not taste the same world. The gene comes in two principal haplotypes set by three amino-acid positions. The sensitive form reads **proline–alanine–valine** at those positions; the insensitive form reads **alanine–valine–isoleucine**. We will call them by their initials, as the field does — **PAV**, the taster, and **AVI**, the non-taster.

A person who carries two PAV copies is a strong taster: for them, the thiourea group — the N–C=S motif shared by the laboratory probe compounds phenylthiocarbamide and 6-*n*-propylthiouracil, and by the brassica's own bitter principal — lights TAS2R38 at vanishingly small concentrations, and a raw floret arrives as a warning. A person with two AVI copies barely registers the same molecule; the identical floret is merely green. Most people sit in between, one copy of each. This is not a metaphor for picky eating. It is why a single bowl of broccoli can divide a table into two populations who sincerely disagree about what is on the plate — and why no amount of coaxing ever settled the argument. The argument was genetic. You cannot talk a PAV tongue out of a signal its receptor is built to amplify.

So the axis is not really about the plant's chemistry alone; it is about the meeting of the plant's chemistry with a receptor that half the table reads at full volume. Which means there are two honest ways to end

the disagreement: change every diner, or change the plant. We changed the plant.

Where the bitterness is made. A brassica stores its defenses as *glucosinolates* — stable, sulfur-rich, water-soluble glucosides, held quietly in the vacuole and doing nothing to the plant that holds them. Nothing happens until the tissue is wounded: a knife, a tooth, a tear. Wounding ruptures the cells and brings the glucosinolates into contact with an enzyme they are normally kept apart from, *myrosinase* (a thioglucohydrolase), which cleaves the glucose off in milliseconds. The unstable aglucone that remains rearranges, and *which* product it becomes is the whole story of the plate:

GLUCOSINOLATE (STORED)	MYROSINASE PRODUCT	CHANNEL / TARGET	IFI AXIS	IN THE SS-LINE
Progoitrin · 2-hydroxybut-3-enyl	goitrin · 5-vinyloxazolidine-2-thione	TAS2R38	IFI-B	turned down
Sinigrin · 2-propenyl (allyl)	allyl isothiocyanate	TRPA1	IFI-P	kept
Glucoraphanin · 4-methylsulfinylbutyl	sulforaphane	Nrf2 pathway (nourishment, not a taste)	—	protected
Glucobrassicin · indol-3-ylmethyl	indole-3-carbinol	minor TAS2R	IFI-B (minor)	trimmed

Read the table top to bottom and the design chooses itself. The bitterness that divides the table is almost entirely the top row: **progoitrin**, which under myrosinase cyclizes into **goitrin**, a thiourea-bearing oxazolidinethione that is one of the plant kingdom's cleanest TAS2R38 agonists. Goitrin is the arm we quiet. The second row, **sinigrin**, becomes allyl isothiocyanate — the clean, mustardy prickle of horseradish and fresh-cut cabbage — and that sharpness is not even read by the bitter receptor; it opens **TRPA1**, the pungency channel, a different nerve entirely. That is the edge we deliberately *keep*, because a brassica with no sharpness at all is a brassica that has lost its nerve, and on the scorecard it is the reason a low-bitter head still shows a living *IFI-P*. The third row is the one we guard most carefully: **glucoraphanin**, whose product **sulforaphane** is the compound the nutrition literature cares about, and it is a taste of nothing — we protect it not for flavor but for nourishment, because the point was never to strip the plant but to quiet the one note that kept people from eating it.

The lever. Here is the elegance, and it is a single enzyme. Progoitrin is not made from scratch; it is made in the last step from a simpler aliphatic glucosinolate by a hydroxylation — a **2-oxoglutarate–dependent dioxygenase**, *GSL-OH*, that hangs a hydroxyl on the side chain and turns a mild precursor into the goitrin-bearing one. Down-regulate that one dioxygenase and the carbon that would have become progoitrin simply never acquires the group that makes it a goitrin factory. Sinigrin is untouched, because it never passes through that step. Glucoraphanin is untouched, because it is on a different branch. Myrosinase itself is left entirely intact — and this is the design decision that matters most, because the lazy intervention would be to disable the enzyme and silence everything at once. Disable myrosinase and yes, the bitterness falls, but so does the pungency *and* the sulforaphane; you would have a blander, less nourishing

vegetable, which is the opposite of the work. We do not touch the hydrolase. We touch the single hydroxylation that feeds the bitter arm, and we leave the rest of the braid standing.



Broccoli, in cross-section — the florets are a forest of wound sites; every surface the knife opens is where myrosinase meets its substrate, and where the number on the card is decided.

The number, worked

Now the measurement, because a claim of *sixty percent quieter* is worth nothing until you can say sixty percent of what, against what, and to what precision.

Every axis of the Index is a loudness, and loudness is the inverse of a threshold. The governing equation is the same on all eight axes:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. A compound that reaches half-max at a *low* concentration is loud; one that needs a *high* concentration is quiet. Inverting the threshold and scaling by a thousand turns that into a score where bigger means louder. On the bitter axis the "stimulus" is the food's expressed juice, assayed on TAS2R38-expressing cells and read as the dilution at which calcium flux reaches half-maximum. We anchor the reference at the standard probe, **6-*n*-propylthiouracil (PROP), with EC_{50}(ref) = 4.0 µM** — the value that, by convention, pins *IFI-B* to 1,000. A food louder than that reference scores above 1,000; a food quieter scores below.

A worked line, start to finish. A conventional broccoli head, assayed, half-activates TAS2R38 at an effective **0.87 µM** expressed in goitrin-equivalents. Then:

> $\mathrm{IFI\text{-}B}(\text{native broccoli}) = 1000 \cdot 4.0\ \mu M \,/\, 0.87\ \mu M = \mathbf{4{,}600}$

That is a loud plant — close to the top of the scorecard's per-axis ceiling of 5,000, which is exactly why native broccoli is the vegetable most often left on the plate. Now the low-bitter head, designed in our model to half-activate at **2.29 µM**:

> $\mathrm{IFI\text{-}B}(\text{low-bitter broccoli}) = 1000 \cdot 4.0\ \mu M \,/\, 2.29\ \mu M = \mathbf{1{,}750}$

The attenuation is the thing the diner feels:

$$\Delta B = (4{,}600 - 1{,}750) / 4{,}600 = 2{,}850 / 4{,}600 = \mathbf{0.620 \rightarrow a\ 62.0\%\ drop}$$

And the same arithmetic on kale — native half-activating at 0.93 µM for an *IFI-B* of 4,300, engineered at 2.33 µM for 1,720 — gives (4,300 − 1,720)/4,300 = **60.0%**. Note one quiet, important fact: the attenuation ratio does not depend on the reference we chose. The 4.0 µM anchor sets the *units* of the axis, but it cancels in the ratio of two scores, so *sixty-two percent quieter* is true whether you anchor on PROP, on PTC, or on the native head itself. The reference fixes the scale; the drop is invariant.

PREPARATION	EC_{50}, GOITRIN-EQ (MM)	GOITRIN-EQ (MG/L)	IFI-B	Δ VS NATIVE
Native broccoli head	0.87	0.112	4,600	–
Low-bitter broccoli · SS-line (predicted)	2.29	0.296	1,750 ± 123	−62.0%
Native kale leaf	0.93	0.120	4,300	–
Low-bitter kale · SS-line (predicted)	2.33	0.301	1,720 ± 120	−60.0%

The milligram column is there for the formulator, and it is a plain unit conversion. Goitrin is C_5H_7NOS, molar mass **129.18 g/mol**, so a threshold of 2.29 µM is

$$2.29 \times 10^{-6}\ \text{mol/L} \times 129.18\ \text{g/mol} = 2.96 \times 10^{-4}\ \text{g/L} = \mathbf{0.296\ mg/L}$$

against the native head's 0.112 mg/L. In words: you must put more than two and a half times as much of the engineered head on the tongue to reach the same bitter half-activation. The ratio is the loudness ratio exactly — 2.29/0.87 = 2.63, and 4,600/1,750 = 2.63 — and that is not a coincidence but a consequence worth deriving.

Why the score tracks the goitrin load, one for one. Let a food, undiluted, deliver its goitrin at an intensity we can call *L*. The receptor has a fixed half-activation constant *K* — a property of TAS2R38, not of the plant — so to reach half-max you need the food diluted to the fraction *K/L*. That fraction *is* the food's EC_{50} in food-concentration units, so $EC_{50}(c) = K/L$, and therefore

$$\text{IFI-B} = 1000 \cdot EC_{50}(\text{ref}) / (K/L) = (1000 \cdot EC_{50}(\text{ref}) / K) \cdot L \propto \mathbf{L}$$

The score is directly proportional to the delivered goitrin load. Cut the load to 38% of native and the score falls to 38% of native — a 62% drop — which is precisely what *GSL-OH* down-regulation does to the progoitrin pool and precisely what the two cultivars show in the model. The biology and the number are the same sentence written twice.

What the receptor actually does at a bite. The proportionality above lives in the low-dose regime; the receptor itself answers along a saturating curve. Take a representative intrinsic constant for goitrin at the taster's receptor, $K \approx 10$ µM, and the Michaelis–Hill form of fractional response, $f = C / (K + C)$. A native mouthful delivering about 10 µM of goitrin sits at

```
f = 10 / (10 + 10) = 0.50 — half-max, unmistakable
```

The low-bitter head, carrying 38% of that load, delivers about 3.8 µM:

```
f = 3.8 / (10 + 3.8) = 3.8 / 13.8 = 0.275 — just over a quarter
```

So the volume knob on the taster's own receptor turns from one-half to barely more than one-quarter. (The massaged kale, at 40% load, delivers ≈ 4.0 µM and lands at 4.0/14.0 = 0.286 — the same neighborhood.) Honesty requires the footnote: because the curve bends toward saturation, the 50→27.5 drop in occupancy *understates* the full loudness ratio that the EC_{50}-based score captures cleanly — the two agree exactly only in the sub-saturating tail. This is why the Index is anchored on EC_{50} and not on response-at-a-fixed-dose: the threshold is the invariant; the occupancy at any one bite depends on how big the bite was.

The error bars are the point. A number without a tolerance is a wish. Every score here carries a design specification of **≤ 7%** batch coefficient of variation — so the engineered broccoli is 1,750 ± 123, the kale 1,720 ± 120 — and that reproducibility is itself a piece of engineering. A native brassica's bitterness swings with soil sulfur, water stress, day length, and the age of the head; measured across real harvests it runs a CV near **42%**. On a native score of 4,600 that is a plate-to-plate swing of roughly ±1,930 — the difference between *merely green* and *inedible* inside a single variety, arriving unannounced. That swing, not the average, is what taught a generation of children to distrust the vegetable: not that it was always bitter, but that it was *sometimes* bitter and you could not tell by looking. Pulling the CV from 42% to ≤ 7% is the part a diner feels as trust. When two uncorrelated measurements combine, their relative errors add in quadrature — $\sqrt{(0.07^2 + 0.07^2)} \approx 0.099$ — so even comparing two engineered batches, the spread stays under ten percent, where the native plant could not promise within forty. Consistency is a flavor. It tastes like being able to serve the same dish twice.

Charred low-bitter broccoli

The first preparation is broccoli taken as hot and hard as a pan will allow — the treatment that, on a conventional head, drives the sulphur notes to their worst. On a low-bitter cultivar the same fierce heat does the opposite: with the goitrin arm quieted, the florets' own sugars step into the space the harshness used to hold, and a hard char reads as caramel and nut rather than scorch.

```
Charred low-bitter broccoli · serves 4 · 18 min
– 2 heads low-bitter broccoli (SS-line), cut into flat-faced spears
– 2 tbsp cold-pressed oil, high smoke point
– 1 lemon, zest and juice
– 2 cloves garlic, sliced paper-thin
– flaked salt; a pinch of chili, if you like a little heat under it
```

1. Heat a heavy skillet or cast-iron pan 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 with room between the spears. Leave it alone for three to four minutes — the first side should go deep brown, almost black at the edges.
3. Turn once, add the garlic to the oil at the pan's edge where it will toast but not burn, and give it another two minutes.
4. Off the heat, hit it with the lemon juice — it will hiss — then the zest, flaked salt, and the optional pinch of chili. Serve at once, while the edges are still crackling.

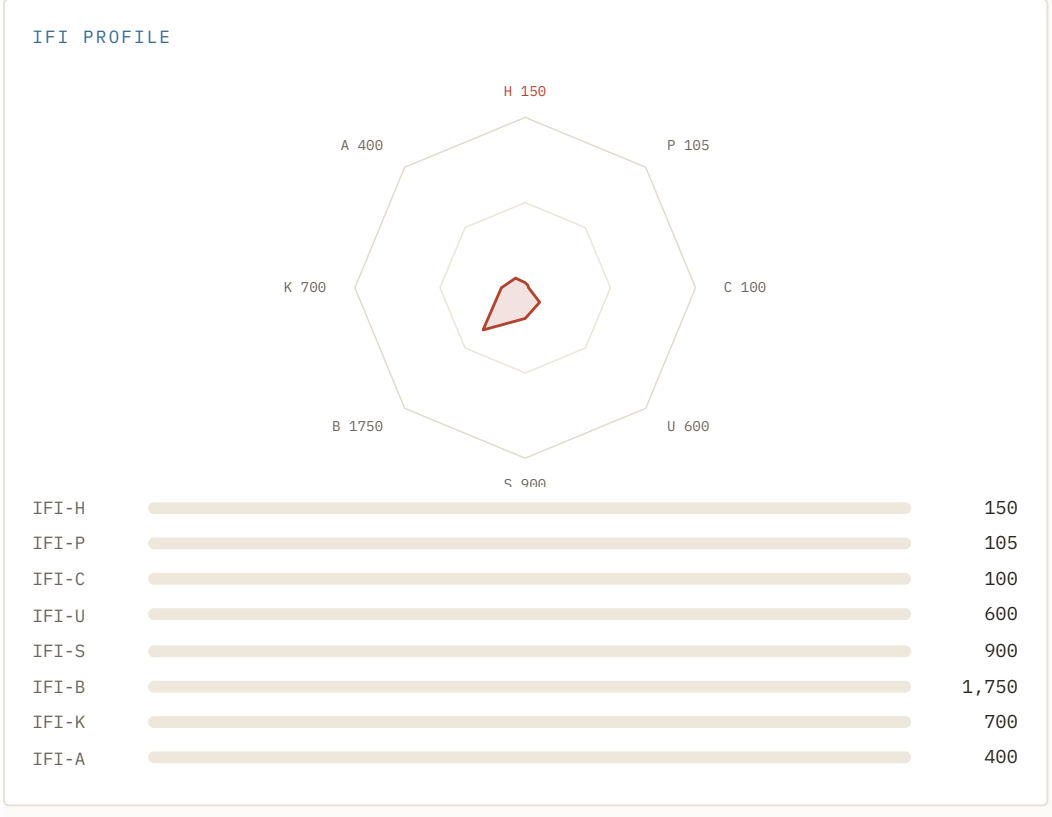

Read the scorecard and the dish explains itself. The **B of 1,750** is the whole point — a bitterness pulled down 62% from the native head's 4,600, exactly the drop the arithmetic predicts from a goitrin load cut to 38% — and look what rises into the quiet it leaves: an **S of 900**, the florets' own sugars, now audible; a **K of 700**, the green, savory body that a conventional broccoli hides behind its bite; a modest **P of 105**, the clean, mustardy edge we deliberately *kept*, because the sinigrin arm was never on the chopping block.

There is a second quieting in this recipe that the engineering did not have to do, and it belongs to the pan. Myrosinase — the enzyme that turns stored glucosinolate into fresh goitrin — is a protein, and proteins denature with heat; above roughly 60 °C it folds and stops working. A hard sear therefore does to the *enzyme* what the edit did to the *substrate*: it quenches further hydrolysis the moment the cut faces hit the metal. So the charred head is quiet twice over — low in the progoitrin that feeds the bitter arm, and seared before much of even that could be cleaved. The caramel you taste is real sugar undergoing real Maillard browning and caramelization in the dry heat; what has changed is only that the browning is no longer competing with a sulphur signal loud enough to drown it. The char reads as sweetness because the signal that would have read as scorch was turned down at the enzyme and then quenched in the pan, and the plant's own sugars were there the whole time, waiting for the quiet.



Low-bitter broccoli, trimmed for the pan — on a head whose goitrin arm is already quieted, the cook's job is no longer to hide the vegetable but to brown it.

Massaged low-bitter kale

The second preparation is the one that converts the skeptics: raw kale, the vegetable most often eaten as penance, dressed and eaten with pleasure. The trick cooks use — massaging raw kale with salt and acid to break down the leaf and blunt the bite — is a workaround for exactly the harshness this cultivar no longer carries, so here it becomes not a rescue but a finish.

```
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 fresh lemon juice + 2 tbsp good olive oil
— 1 tsp sea salt
— 2 tbsp toasted seeds (pumpkin or sunflower)
— shaved hard cheese or a handful of currants, to finish — cook's choice
```

1. Put the torn leaves in a wide bowl with the salt, lemon, and oil.
2. Massage with your hands for two to three minutes — the leaves will darken, soften, and nearly halve in volume as the cell walls give way.
3. Let it rest ten minutes so the dressing works through. Fold in the seeds and your chosen finish just before serving.

IFI PROFILE

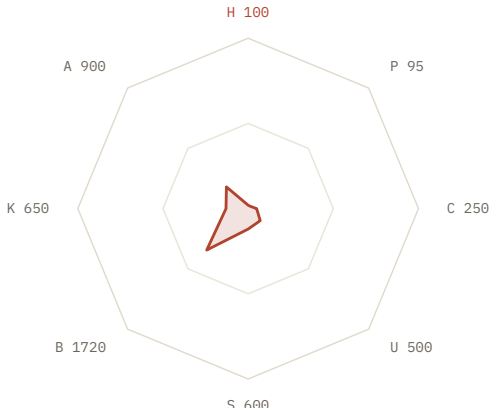

IFI-H	100
IFI-P	95
IFI-C	250
IFI-U	500
IFI-S	600
IFI-B	1,720
IFI-K	650
IFI-A	900

This recipe is the honest mirror of the last one, and the contrast is instructive. The broccoli went into a hot pan, so its myrosinase was quenched; the kale is eaten raw and *massaged*, which ruptures the

cell walls as thoroughly as chewing does and sets myrosinase fully loose on whatever substrate it finds. Raw kale is therefore the harder test — there is no heat to hide behind, and every bit of stored glucosinolate that can be hydrolyzed, is. That our model puts the leaf at **B 1,720**, a 60% drop from the native 4,300, reflects the design intent: the attenuation lives in the plant, not in the cooking. With the enzyme running at full throttle, the only thing keeping the bitter arm down is that the *GSL-OH* edit left so little progoitrin for it to work on. Massaging, which for a century was a rescue — a way to mechanically wreck a leaf until it stopped fighting back — becomes here what it should always have been: a way to soften texture and carry dressing into the leaf, a finish rather than a damage-control measure.

The **A of 900** is worth a word, because it is the thread we chose *not* to touch, and it is not bitterness at all. Astringency is not a taste in the receptor sense; it is a mechanical event. The tannins and polyphenols in a good leaf bind and precipitate the proline-rich proteins that lubricate the mouth, and when that lubrication is stripped away the tissues drag against each other and the brain reports *grip* — the clean, almost mineral pucker that makes a leaf feel substantial rather than slack. Bitterness and astringency ride together in a raw brassica and are forever confused for one another, but they are different machines: one is a molecule in a receptor, the other is a protein falling out of solution. We quieted the bitter machine, down about sixty percent, and left the astringent one standing at 900, which is why this salad reads as *present* and green rather than merely soft. A kale with the astringency also stripped out would be a wet rag with good intentions. The grip is the backbone; we kept the backbone.



Low-bitter kale, dressed and rested — raw, enzyme fully active, and still quiet: the clearest evidence that the attenuation lives in the plant, not the pan.

The table set

Set either of these bowls in the middle of a table of strong tasters and non-tasters alike and watch what happens, because this is the experiment the whole axis was built to run. Recall the two populations: the **PAV** taster, whose TAS2R38 lights at the faintest trace of goitrin, and the **AVI** non-taster, who barely reads it at all. On a native head delivering ≈ 10 µM of goitrin, the taster's receptor sits near half-maximal — call it 50% occupancy — while the non-taster, whose variant answers only to far higher concentrations (an effective

constant up around 120 µM), sits near 10/130 ≈ 7.7%. That forty-point gulf *is* the divided table: one diner recoiling, the other baffled by the fuss, both of them right about their own tongue.

Now serve the low-bitter head. Delivered goitrin falls to ≈ 3.8 µM, and the curves move toward each other. The taster drops to 3.8/13.8 ≈ 27.5%; the non-taster to 3.8/123.8 ≈ 3.1%. The gap between genotypes narrows from roughly 42 points to roughly 24, and — more to the point — the taster's own signal has fallen out of the *beware* register into the merely-perceptible. We did not change a single diner's genes. We lowered the plant's goitrin load until the loudest-reading tongue at the table reads it the way the quietest one always did. The warning that used to sort the table by genotype has been turned down below the threshold that did the sorting, and the thing you notice, the first time, is that for once everyone eats the same amount. That is the axis, served.

What this is not, honestly. The book is titled for a direction, not a destination. *Zero-Bitter* is where the axis points, not a literal zero, and it would be a kind of lie to pretend otherwise. The projected attenuation band runs 20–65%, and these two cultivars sit near the top of it at 60–62% — real, large, table-changing, and still not total. We could push further; we choose not to, for three reasons the scorecards make visible. First, a little bitterness is structure, the same way a little tannin is: strip the last of it and the vegetable goes flat and ingratiating, a plant that has lost its nerve, which is a worse failure than being too sharp. Second, the arms we keep are kept on purpose — the sinigrin that becomes the mustardy *IFI-P* prickle, and above all the glucoraphanin that becomes sulforaphane, which is nourishment and which we will not sacrifice to chase a number. Third, the kitchen still has a vote: myrosinase is heat-labile and the epithiospecifier proteins divert its products, so how you cut, how long you wait, and how hard you heat all shift the delivered load by a margin of their own — the card reports the plant's potential, and the cook spends it. A wet pan will still steam a perfect head into sullenness; no edit saves a cook from physics.

And the genotype cannot be *erased*, only out-waited. A PAV taster is still a PAV taster; we have merely moved the plant below the concentration where that genotype matters at this table. Serve them a native head tomorrow and the old argument returns intact. What we have built is not a cure for human variation — it is a vegetable generous enough to meet the whole range of it, so that the most sensitive person and the least sensitive person can sit at the same plate and reach, both of them, for the same second helping. For a family that has always believed food is something you give the body with intention, that shared reach is the only proof that ever mattered. The number on the card is how we know we can promise it again next week; the child reaching for seconds is why we bothered to learn to measure it at all. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

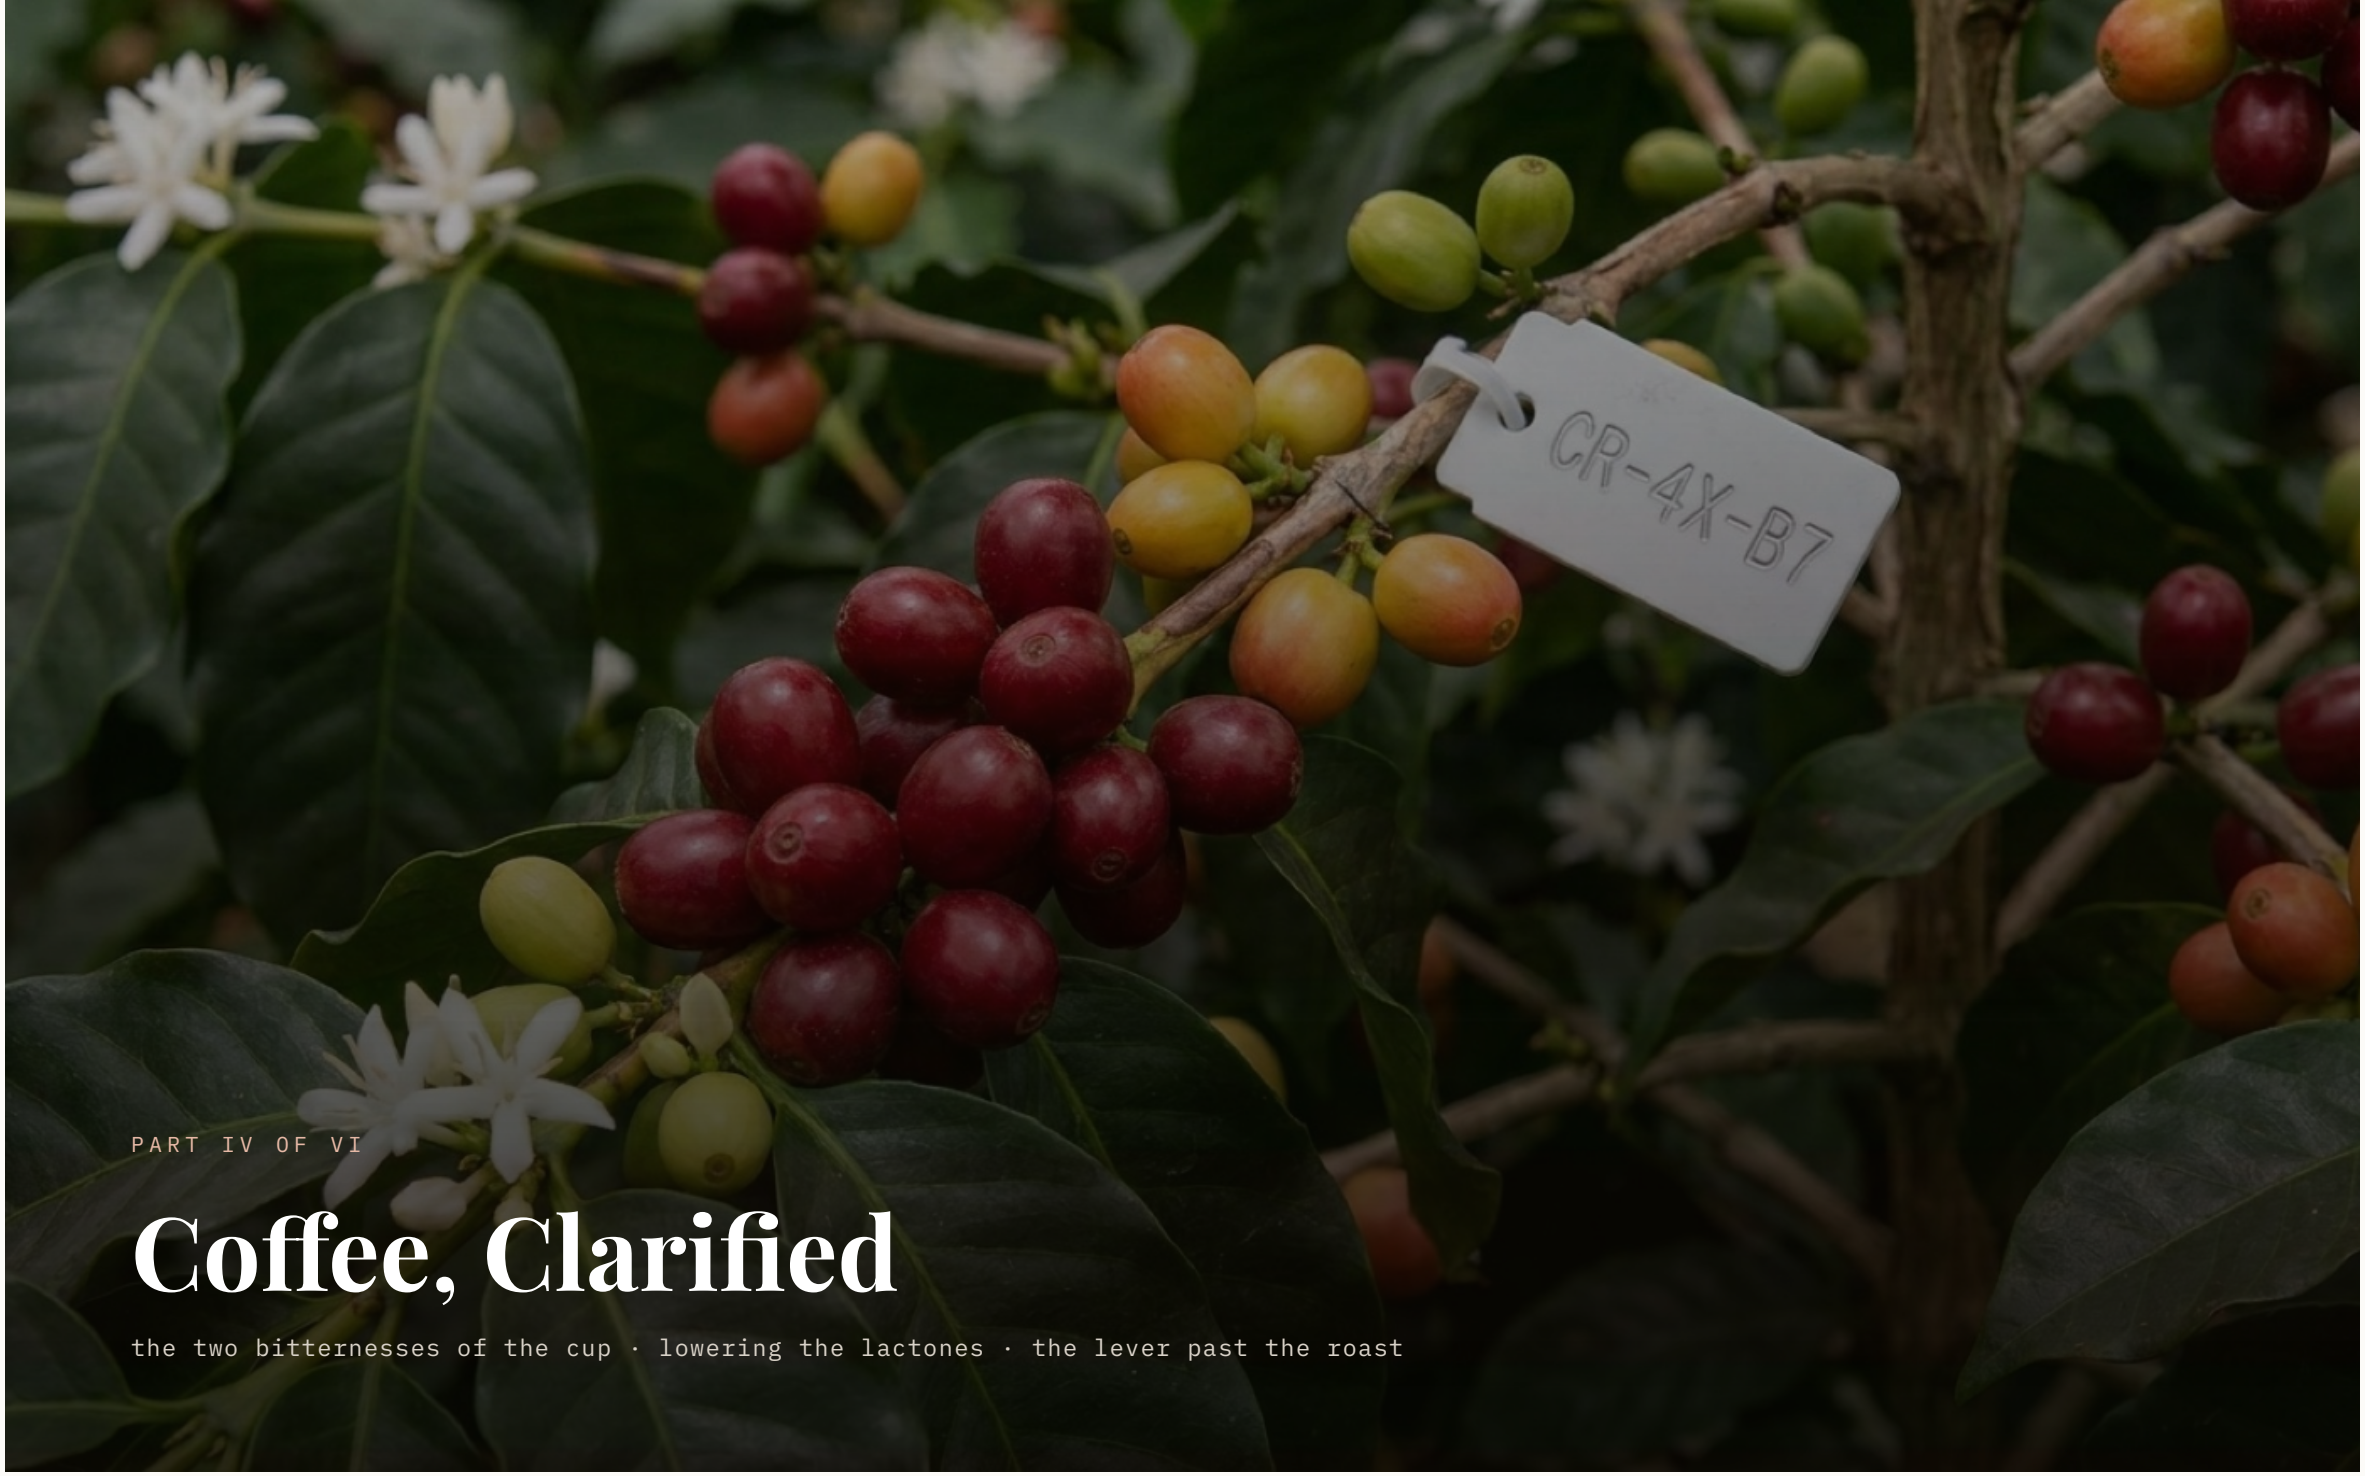

PART IV OF VI

Coffee, Clarified

the two bitternesses of the cup · lowering the lactones · the lever past the roast

Chapter 11 · The Two Bitternesses of Coffee — *lactones, phenylindanes, and the roast.*

DOMAIN · COFFEE · Coffea · the chlorogenic acids

A great coffee is a controlled bitterness, and we would never want it gone. The pleasure of the cup lives in a dark, round, almost savory bitterness that no sweetness could stand in for — the thing a drinker falls for, the reason the acquired taste is worth acquiring. So coffee, more than any plant on the shelf except cacao, is a lesson in the difference between the two bitternesses hiding inside 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. The work of this chapter is the one that moves them apart — and because the whole method turns on telling the two apart, we have to be precise about the chemistry before we are allowed to be reverent about the cup.

What makes coffee the right teacher is that its bitterness is not inherited but *made.* The green bean is one of the least bitter things in the pantry; a cup is one of the most. Everything bitter you taste was built in the minutes between, by heat acting on a single family of molecules the plant laid down in the seed for its own reasons. If we understand that family, and the two things heat does to it in sequence, we understand exactly which bitterness to keep and which to lower — and, more usefully, *where* the lever sits. It does not sit in the roaster. It sits upstream, in the green bean, long before the first crack, in a pathway we know how to tune.



Green and roasted coffee, side by side — the bitterness that defines the cup is forged in the gap between them, as the chlorogenic acids break down under heat.

The green bean's inheritance — the chlorogenic acids

The green coffee bean is unusually rich in **chlorogenic acids**, a family of esters built on the phenylpropanoid branch — the same universal scaffold **PAL → C4H → 4CL → HCT → COMT** that runs in every vascular plant, here routed through a coffee-specific

hydroxycinnamoyl transferase (HQT) that stitches a hydroxycinnamic acid onto quinic acid. The products are the caffeoylquinic acids (3-, 4-, and 5-CQA, with 5-CQA the dominant isomer), the feruloylquinic acids, and the dicaffeoylquinic acids. Together they can reach six to eight percent of the green bean's dry weight — among the highest chlorogenic loads of any plant food we measure.

The load is not fixed across the genus, and that variance is the first hint of the lever. Robusta beans (*Coffea canephora*) carry a markedly heavier chlorogenic reservoir than arabica (*Coffea arabica*) — one reason robusta roasts toward a harsher, more aggressively bitter cup, and one reason blenders have long reached for it as a blunt bitterness dial. Within a single bean the isomers sort as well: 5-CQA dominates, the 3- and 4-CQA isomers trail it, and the dicaffeoylquinic acids sit as a smaller, more reactive fraction that runs hottest toward the phenylindane end. Nature already varies the green-bean reservoir from plant to plant and reads it out as bitterness in the cup. We are not inventing a new dial. We are taking the one evolution left loose and giving it a precise setting.

In the green bean these compounds are barely bitter. They are a store of reactive phenolics the seed built for defense and antioxidant protection, not for flavor, and the tongue reads them faintly. The bitterness is *made in the roast,* and it arrives in two distinct waves, each from a different fate of that same chlorogenic pool. Hold that picture — one inherited reservoir, two thermal destinations — because it is the entire reason a single green-bean lever can reshape the cup.

The first wave — the lactones we keep

As the roast climbs into light and medium territory, the chlorogenic acids begin to dehydrate and cyclize. Each sheds a molecule of water and closes a ring, the quinic acid moiety folding into a lactone — the **chlorogenic acid lactones,** or quinic acid lactone esters. This is the first wave, and it is the bitterness we protect.

These lactones are the pleasant, structured, coffee-defining bite: the round depth of a well-judged medium roast, the thing that reads as *character* rather than as damage. They peak somewhere in the medium band and then, pushed further, begin to fall. If the cup had only this wave it would be the platonic coffee — bitter enough to be itself, never bitter enough to punish. The lactones are the soul of the drink, and nothing in our method is permitted to touch them. A coffee stripped of its chlorogenic lactones is a thin, hollow, forgetful cup, a coffee that no longer tastes like coffee. The first wave is not a problem to be solved. It is the thing being defended.

The second wave — the phenylindanes we lower

Push the roast darker, past the point where the lactones crest, and the lactones themselves break down. From their degradation — by way of caffeic acid and vinylcatechol intermediates that cyclize and fuse — emerge the **phenylindanes,** fused bicyclic compounds that are the molecular signature of a dark roast. These are the second wave, and they are the punishment: a harsh, lingering, almost ashen bite that clings to the back of the throat and sends a careful drinker reaching for milk.

Here is the cruelty of the single continuum. The lactones and the phenylindanes are not independent dials. They are consecutive stages on one axis of heat: the second is *made from* the first. A roaster who wants more body, more sweetness, more of the caramelized Maillard depth that only a darker roast delivers must push into exactly the regime where the phenylindanes bloom. Chase the good thing with the only tool on the bench — the roast — and the bad thing comes with it. Pull the roast back to spare the phenylindanes and you leave the body and the lactone richness on the table. One knob, two bitternesses, soldered together. This is why the roaster alone can never separate them, and why the lever has to come from somewhere the roaster cannot reach.

A word on caffeine, because its reputation outruns its role. Caffeine is genuinely bitter and it does contribute, but careful studies of coffee bitterness put it at only perhaps ten to fifteen percent of the total perceived bitterness of a cup. The overwhelming majority is the chlorogenic story — the lactones, and above all, in a dark roast, the phenylindanes. This matters for design more than it first appears: it means a coffee can be meaningfully de-bittered *without touching its caffeine at all.* We treat that as a deliberate boundary, not an accident. The harsh wave and the stimulant are different molecules on different pathways, and we lower the one without pretending to have made a decaf.

The committee on the tongue — receptors and transduction

Which receptors read the cup? Not, mostly, our anchor. Bitterness is sensed by the TAS2R family of G-protein-coupled receptors, and coffee strikes a *chord* across several of them rather than a single clean note. The chlorogenic lactones and the phenylindanes engage **TAS2R46,** a broadly tuned receptor, among others; caffeine lights up **TAS2R7, TAS2R10, TAS2R14, TAS2R43,** and **TAS2R46** together. A cup is a committee — a layered, polyphonic activation, not the sharp single-receptor hit that a glucosinolate lands on **TAS2R38.**

The transduction that follows is the same for every one of these receptors, and it is worth walking once, because the shape of it explains the measurement. A bitterant binds its TAS2R; the receptor activates the G-protein gustducin; gustducin's βγ subunits switch on phospholipase C-β2; PLC-β2 cleaves membrane PIP_2 into IP_3; IP_3 releases calcium from internal stores; the calcium gates the TRPM5 cation channel; the cell depolarizes and releases ATP onto the afferent nerve. Every bitter note in the cup funnels through that cascade. The half-maximal point of the whole machine — the concentration at which a stimulus drives it to half of full response — is the number we build everything on.

TAS2R46 does so much of the hearing here for a structural reason. It is among the most broadly tuned of the bitter receptors, with a shallow, accommodating orthosteric pocket that reads a wide range of structurally unrelated bitterants — which is exactly why one receptor can answer to a chlorogenic lactone, a fused phenylindane, and a methylxanthine like caffeine all at once. Breadth is the feature, not a flaw: a plant's chemical defenses are diverse, so the receptors that evolved to guard against them are generalists. It also means coffee's

bitterness cannot be unpicked receptor by receptor from the outside. The chord is sensed as a chord. The only clean place to act is upstream of it, on the chemistry that strikes the chord in the first place.



A coffee specimen staged for the panel — the same cup that one drinker calls round and another calls harsh, resolved into a single reproducible reading.

We score coffee, as we score everything, against the **TAS2R38-anchored** bitter 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 is doing the hearing. For coffee, a committee is — and the projection onto a single axis is a convenience of measurement we are careful never to mistake for a claim about biology.

COMPOUND CLASS	REPRESENTATIVE	PRIMARY TAS2RS ENGAGED	ROAST WAVE
Chlorogenic acids (intact)	5-CQA	TAS2R46 (weak)	pre-roast
Chlorogenic lactones	quinic lactone esters	TAS2R46	first wave · light–medium
Phenylindanes	dark-roast bicyclics	TAS2R46 (strong)	second wave · dark
Caffeine	1,3,7-trimethylxanthine	TAS2R7 · R10 · R14 · R43 · R46	all roasts (10–15% of total)
Reference anchor	TAS2R38 agonist (glucosinolate-type)	TAS2R38	axis calibration

Putting a number on the cup — the IFI-B method

The Index turns loudness into a score by anchoring each axis to a half-maximal activation and inverting it against a fixed reference.

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response; the smaller the EC_{50}, the louder the compound, because less of it is needed to say the same thing. We anchor the bitter axis to the half-activation of a calibrated **TAS2R38** reference agonist — the thiourea-class bitterant the whole library is leveled against, the same receptor a glucosinolate strikes — and we fix its value by convention:

$$\mathrm{EC}_{50}(\mathrm{ref}) = 1.80\ \mu\mathrm{M}$$

To read a cup, we take its bitter fraction, build a dose–response curve against the receptor panel, and fit it with the Hill equation,

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

where θ is the fraction of the response driven at bitterant concentration [A], and n is the Hill coefficient — the steepness of the curve, a measure of cooperativity. Half-maximal response (θ = 0.5) falls exactly at [A] = EC_{50} for any n, which is precisely why EC_{50} is the honest place to read loudness: it is the one point on the curve that does not depend on the curve's shape.

In practice the curve comes from two instruments read in parallel. A cell-based reporter line expressing the relevant TAS2Rs gives the receptor's own answer, activation read as a calcium-flux or luminescence signal across a dilution series; a trained human panel rates the same series by mouth. The two must agree before a reading is logged — the assay supplies precision and the panel supplies meaning, and EC_{50} is taken where they cross. Neither alone is trusted. A number the tongue will not confirm is not a flavor measurement, and a tasting a receptor cannot reproduce is not a standard.

The shape still tells us something the number alone hides. Bitter transduction is cooperative — the IP_3-and-calcium stage amplifies — so n runs above one; take n ≈ 1.8 for the committee a cup convenes. Walk a dark roast across its EC_{50} (1.44 µM, as the table below reports) and the Hill curve is brutally steep. At half the EC_{50}, [A] = 0.72 µM:

$\theta = 0.72^{1.8} / (1.44^{1.8} + 0.72^{1.8}) = 0.554 / (1.928 + 0.554) =$ **0.22**

At twice the EC_{50}, [A] = 2.88 µM:

$\theta = 2.88^{1.8} / (1.44^{1.8} + 2.88^{1.8}) = 6.714 / (1.928 + 6.714) =$ **0.78**

A fourfold change in concentration swings the receptor from 22% to 78% activation — a cliff. Had the system been non-cooperative (n = 1), the same fourfold span would have moved it only from 0.33 to 0.67, a gentle ramp. The cooperativity is why a dark roast crosses so abruptly from *deep* to *scorched*: the phenylindanes do not fade in, they switch on. The math is not decoration here. It is the reason the margin for error at the dark end of the roast is so thin, and the reason a lever that pulls the whole curve leftward — lower EC_{50} demand, before the roast — buys so much room.

Now the worked readings. We standardize everything the cup cannot — the same dose per volume, the same water, the same brew — so that the only variable left is the bean and its roast, and we read EC_{50} off each curve.

ROAST STAGE	DOMINANT BITTERANT	EC_{50} (MM, BITTERANT-EQUIV.)	IFI-B
Green (unroasted)	chlorogenic acids, intact	12.00	150
Light	CGA + early lactones	4.00	450
Medium	chlorogenic lactones	2.40	750
Medium-dark	lactones + early phenylindanes	1.85	973
Dark (French)	phenylindanes	1.44	1,250
Engineered dark (designed) *	phenylindanes (precursor-reduced)	3.20	563

** Engineered dark EC_{50} and IFI-B are computational predictions from our flux-balance model, not laboratory measurements.*

Read the column of scores before touching any of them. A medium roast sits at IFI-B = 1000 · 1.80 / 2.40 = 750; the green bean, barely bitter, at 150; the French roast at 1,250. That is better than an eightfold dynamic range from raw seed to dark cup, every bit of it manufactured by heat from one reservoir — the quantitative shadow of the two waves building and then turning harsh. The engineered row is the only one that breaks the monotonic climb: the same color as the French roast, half the loudness.

Two computations show the whole arithmetic. The French roast reaches half-max at EC_{50} = 1.44 µM:

IFI-B = 1000 · 1.80 / 1.44 = 1000 · 1.25 = **1,250**

The engineered bean, roasted to the *same* French color, reaches half-max only at EC_{50} = 3.20 µM — it takes more than twice the concentration to shout as loud, because there is simply less phenylindane precursor to burn:

IFI-B = 1000 · 1.80 / 3.20 = 1000 · 0.5625 = **563**

The attenuation is the fall between them:

(1,250 − 563) / 1,250 = 687 / 1,250 = **0.55 → 55%**

Our model predicts 55 percent quieter at identical roast color — squarely inside the 20–65% band our Zero-Bitter designs target across the bitter-host library. On every other axis the Index rewards a bigger number; on IFI-B, the work is to make the number *fall,* and to make it fall for one reason only.

A note on units, so the micromolar readings touch the physical world. Expressed as a 5-caffeoylquinic acid equivalent (molar mass 354.31 g/mol), the reference anchor sits at 1.80 µM × 354.31 g/mol = 0.638 mg/L, and the dark-roast cup's 1.44 µM becomes 0.510 mg/L — about half a milligram of bitterant-equivalent per liter is all it takes to read at IFI-B 1,250. Loudness is not mass. A cup can carry grams of dissolved solids and be read, on the axis that matters, by a few hundred micrograms of the molecules the tongue actually hears.

The discipline of ±7 percent

A score is only as real as it is repeatable, and this is where the engineered bean earns its place before a single flavor argument is made. The platform targets a batch-to-batch coefficient of variation of **≤7%** on its constructs, against the **42%** spread we document in unengineered, wild-type material. That gap is not housekeeping. It is the difference between a measurement and a rumor.

Carry the uncertainty through the index. For a ratio of two measured quantities, relative variances add in quadrature:

> $(\sigma_{IFI} / IFI)^2 = (\sigma_{ref} / EC_{50,ref})^2 + (\sigma_{cup} / EC_{50,cup})^2$

Take the reference as tightly held (3%) and read an engineered cup at the platform's 7%:

$\sigma_{IFI} / IFI = \sqrt{(0.07^2 + 0.03^2)} = \sqrt{(0.0049 + 0.0009)} = \sqrt{0.0058}$ = **7.6%**

So the engineered French roast reports IFI-B = 1,250 ± 95 — a dark roast that cannot be mistaken for a medium at 750, or for the engineered bean at 563. The bands do not touch; the design is legible. Now rerun it on wild-type beans at 42%:

$\sigma_{IFI} / IFI = \sqrt{(0.42^2 + 0.03^2)} =$ **42.1%** → 1,250 ± 526

That interval swallows everything from a city roast to a near-char. A measurement that cannot tell a French roast from a medium is no measurement at all, and no design can be built on it. The ≤7% CV is what makes IFI-B a number a formulator can hold to — and it is one of the quiet reasons we would reach for an engineered bean even before the question of flavor is raised. Consistency is a precondition for design, not a bonus on top of it.

The lever is upstream

Everything so far points to one conclusion: because the phenylindanes are *made from* the chlorogenic acids, and the chlorogenic acids are *made on a pathway we can tune,* the place to intervene is the green bean, long before the roaster's heat ever arrives.

The design is deliberately narrow. We do not want less coffee; we want less of the *second wave.* So we reach into the phenylpropanoid branch that feeds the bean's chlorogenic pool and trim the flux that overfills it — tuning the coffee-specific transferase step and its upstream supply so the green bean lays down a smaller, better-proportioned reservoir of chlorogenic acids. A flux-balance model of the branch fixes the intervention point before a single edit is made: the chlorogenic node itself is the true lever, not some distant upstream step whose trimming would starve the plant of the lignin and structural phenolics it needs and leave it sickly. We move the one reaction that changes the cup, and as few others as the biology will allow. Roast that bean and the first wave still forms: there is lactone precursor enough for the round, structured, character bitterness we defend. But when the roast pushes past the lactone crest, there is markedly less substrate left to degrade into phenylindanes. The harsh wave is starved at its chemical source, before heat can build it. The soul of the cup is left exactly where it was.



Engineered coffee in the controlled plot — the bean re-proportioned at the source, so a roaster can chase body into a darker roast without the ash arriving to end the conversation.

The two scorecards make the move visible. Here is a conventional French-roast cup, read on all eight axes — loud on bitterness, with the modest savory depth and melanoidin body that give coffee its

mouthfeel, and effectively silent everywhere else:

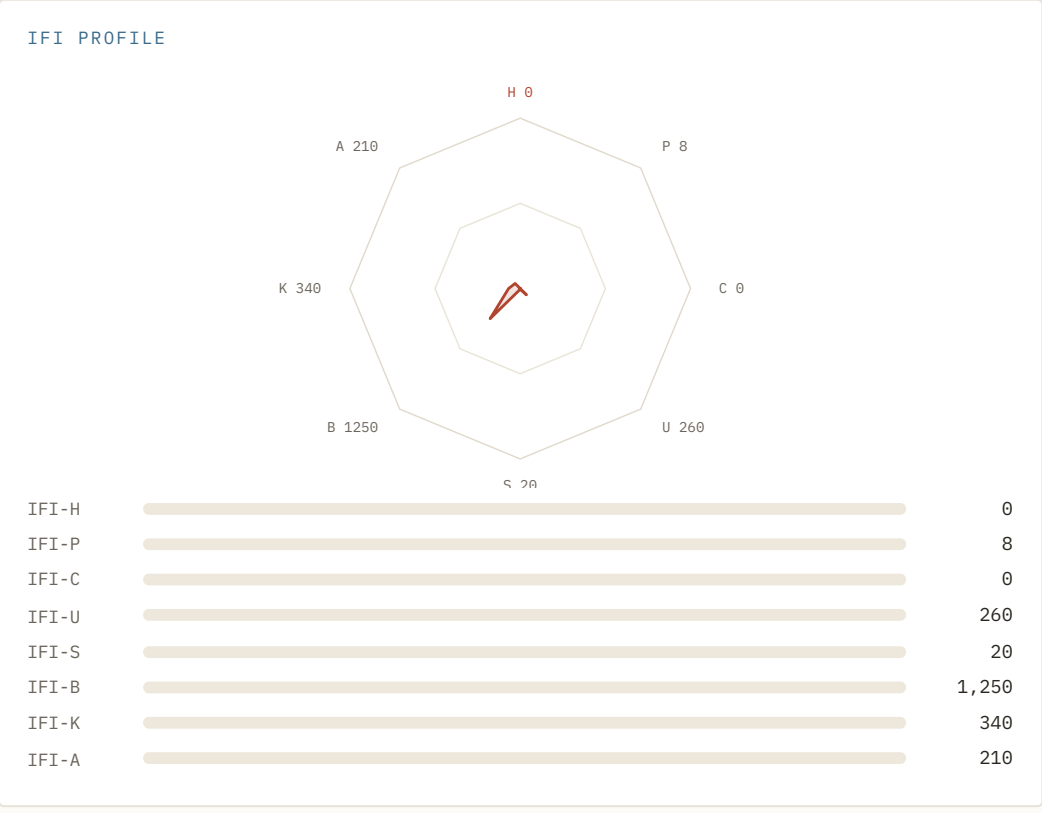

And here is our model's prediction for the engineered bean, designed for the same roast color:

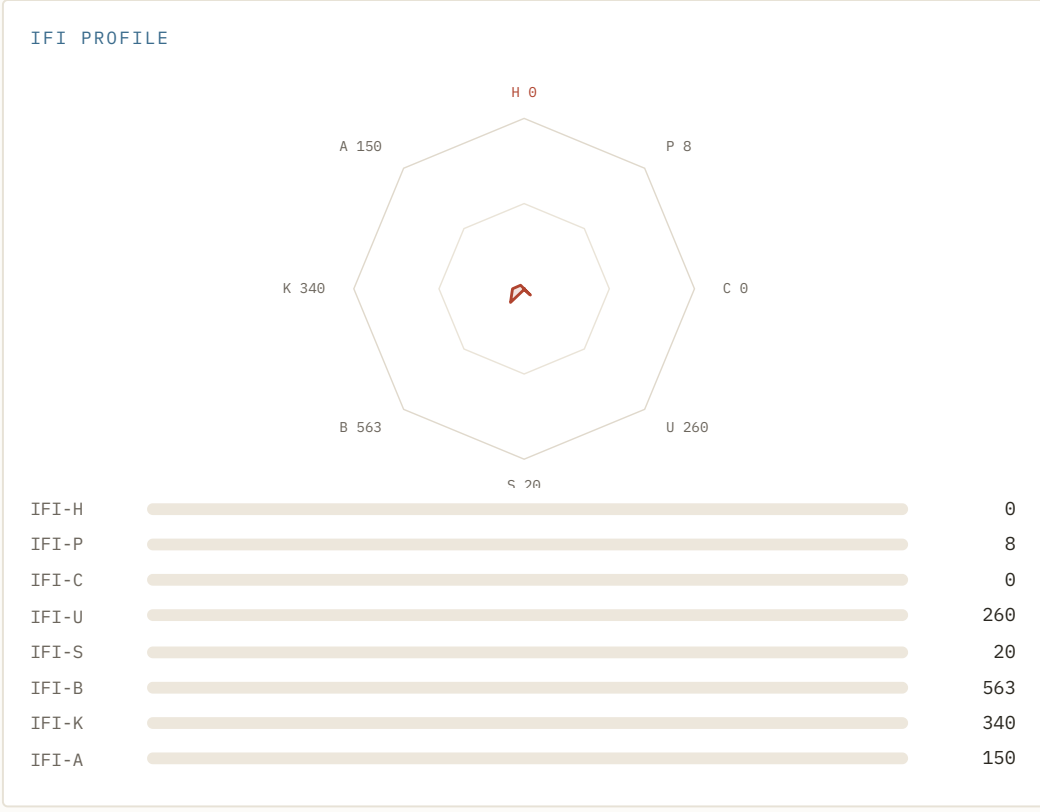

Read the two columns down and the design states itself. Our model predicts umami holding at 260; kokumi — the body, the continuity, the weight of the cup — holding at 340. Heat, cool, sweetness stay where coffee leaves them in our design. Only IFI-B falls, 1,250 to 563, the 55% we computed. Astringency slips a little too in our model, from 210 to 150, and we report it honestly rather than hide it: the chlorogenic acids that feed the phenylindanes also bind salivary proteins, so trimming the pool shaves a touch off the IFI-A tannin reading as a side effect. It is a small, roughly 29% reduction on a secondary axis, and our model predicts a rounder finish in the cup —

a real cross-axis coupling in the design, and naming it is the difference between a design and a sales pitch. Every other number is untouched. That is what "move them apart" means, rendered as arithmetic: one axis down, the rest standing exactly still.

Coffee is one entry in a bitter library the same lever serves. The IFI-B axis runs across every host where a chlorogenic, glucosinolate, polyphenol, or iso-α-acid load drives a bitterness worth lowering — and the construct counts show where the work has gone deep and where, as with coffee, it is still early and hard.

HOST (IFI-B)	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Grape	1,610	B / A / M	Hard
Broccoli	380	B / M	Moderate
Tea	280	U / B / A	Hard
Broccoli-Kale	280	B	Hard
Orange	120	B / C	Hard
Hops	80	B	Moderate
Potato	80	B	Easy
Coffee	40	B	Hard
Cabbage	40	B	Moderate
Kale	40	B	Moderate
Radish	40	B	Hard
Rapeseed	40	B	Hard

Coffee sits at forty constructs and a difficulty of *Hard* — a woody perennial with a long generation time and a bitterness distributed across a committee of receptors, not an easy host. That our model predicts clean attenuation by better than half with the soul intact is evidence that the approach travels: the same inversion of EC_{50} that scores a superhot's heat scores a cup's bitterness, and the same phenylpropanoid branch that we tune in a pepper feeds the green bean's chlorogenic pool.

Honest limits

The design has a floor, and we draw it in the open. The lever lowers the second wave by starving its precursor — but the first wave draws on the same reservoir, so there is a depth of reduction past which the lactone character thins and the cup goes hollow. The target is the phenylindane end of the curve, not the chlorogenic pool entire; push the trim too far and you are no longer lowering the punishment, you are erasing the character. That is why the honest ceiling on the axis is roughly 65%, not 100%. Zero-Bitter is a direction and a discipline, not a literal destination. A coffee with no bitterness left is not a triumph. It is a different, lesser drink.

Three more limits belong on the record. The IFI-B score is a projection: it faithfully reports *how loud* a cup's bitterness is, but it collapses a multi-receptor chord onto one axis, so two foods at the same IFI-B — a coffee and a brassica, say — can share a number and still feel nothing alike, because their TAS2R committees differ. The reading also depends on a standardized cup; origin, process, grind, water chemistry, and brew method all move EC_{50}, which is why the panel fixes everything but the bean. And the attenuation is of the

chlorogenic harsh wave, not the stimulant: caffeine's ten-to-fifteen percent is left deliberately untouched, so nothing here is a decaffeination and we never let it be read as one.

This is the gift the two-bitterness structure hands us, and the restraint it demands in the same breath. The chemistry sorts cleanly — one inherited reservoir, two thermal fates, a wave to defend and a wave to lower — and the receptors, the EC_{50}, the Hill curve, and the ≤7% CV let us put a number on each and watch only the right one fall. We did not set out to make coffee less like coffee. We set out to let a roaster chase body and sweetness into a darker roast without the ash arriving to end the conversation — to lift the punishment and leave the pleasure, designed, reproducible, and whole. A cup given to the body with intention, grown to hold its character and spared its scorch: that is the whole of it, and it is enough. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 12 · Lowering the Lactones — *engineering the bean, not the roast.*

DOMAIN · COFFEE · Coffea · 40 constructs · HQT · the caffeine boundary

For a hundred years the only tools against a coffee's harshness have lived downstream of the seed — in the roastery and the café, in origin and process and roast curve, in grind and water and time. Every one of them is real. Every one of them is also, in the end, the same dial, because they all act on one bean carrying one chemistry, and they move the good bitterness and the bad together. Lighten the roast to dodge the phenylindane scorch and you lose the deep lactone body in the same motion; chase that body with a darker roast and the ash walks back in behind it. The roaster at the drum is a person of extraordinary skill holding a single lever, trying to find the one position where the cup is both round and clean — and discovering, roast after roast, that the position does not exist, because on that lever it cannot. The two things he wants are welded to the same point.

We grew up, as a family, treating the morning cup as something you give the body on purpose: a small clean ritual, unhurried, nothing added it does not want and nothing stripped that it does. That reverence is also, for us, a brief. If the harshness and the soul of a coffee are welded together on every dial the café owns, then the weld was made somewhere the café cannot reach — upstream, in the living seed, in the pathway that lays down the harsh precursor in the first place. That is where we set our hand. Not on the roast. On the bean.



Coffea in the research house — the lever we reach for is not in the roaster or the café but here, in the green seed, at the one step that decides how much chlorogenic acid the bean will ever carry.

The two bitternesses, and why the roaster can only move them together

A great cup of coffee is not bitterness-free, and we would not want it to be. There is a bitterness that is the body of the drink — the low, round, almost sweet heaviness that makes a good dark roast feel like

it has a floor under it — and there is a bitterness that is the drink's failure, a flat metallic scorch that arrives late, dries the back of the tongue, and lingers after everything else has gone. The first we revere. The second is what sends people to sugar and cream. The whole art of the roaster is to carry as much of the first as possible while admitting as little of the second as he can, and the tragedy of the roaster is that both are made from the same molecule on the same road.

That molecule is chlorogenic acid. Green coffee is one of the richest stores of chlorogenic acid in the plant kingdom, and almost none of what we taste as coffee's bitterness is the chlorogenic acid itself — it is what heat makes of it. In the first third of a roast, chlorogenic acids lose water and cyclize into **chlorogenic acid lactones**, the quinides. These are the good bitterness: the body, the weight, the thing the cup would feel hollow without. Push the roast further — into the second crack, into the dark — and the quinides themselves break down, and among their degradation products are the **phenylindanes**, stubborn, planar, strongly bitter molecules that do not wash off the tongue the way a clean bitter should. Light roast: not enough quinides, thin and sharp and sour. Dark roast: plenty of quinides, but the phenylindanes riding in on top of them. There is no roast temperature at which the quinides are high and the phenylindanes are low, because the phenylindanes are made *from the quinides*, one step further down the identical road. The roaster's dial is a single point travelling that road, and body and scorch are two stations on it that cannot be visited apart.

Our lever does not travel that road at all. It stands at the gate before the road begins — at the supply of chlorogenic acid in the green bean — and turns the whole supply down by a designed amount. Lower the chlorogenic load the seed carries into the roast, and every station downstream scales with it: fewer quinides, and *disproportionately* fewer phenylindanes, because the phenylindanes sit two degradations deep and feel any cut at the source twice over. For the first time the two bitternesses move independently, because we are no longer choosing where on the road to stand. We are choosing how much road there is.

How the tongue reads bitter

To tune a bitterness you have to know exactly what you are tuning it against, and bitterness is read by the most paranoid sense the mouth owns. Where sweetness runs through a single receptor pair and heat through one channel, bitterness is carried by a family of roughly twenty-five bitter taste receptors — the TAS2Rs — a sprawling early-warning system evolved to flag the alkaloids and glycosides of a poisonous world. The Index anchors the bitter axis, IFI-B, to **TAS2R38**, the best-characterized member of that family and the one whose pharmacology we trust as a reference; the harsh agonists coffee delivers, the phenylindanes and certain of the chlorogenic lactones, are ligands for this bitter-receptor family, and TAS2R38 is the ruler we read them against.

The receptor is a G-protein-coupled receptor, seven helices threaded through the membrane of a taste cell, and the cascade behind it is worth walking in full, because every number later in this chapter is

ultimately a measurement of how hard this cascade is driven. A bitter ligand settles into the TAS2R38 binding pocket; the receptor changes shape; on its inner face it releases its bound G-protein, **gustducin**. Gustducin splits into its Gα and Gβγ parts, and the Gβγ pair switches on **phospholipase C-β2**, which cleaves a membrane lipid into two messengers, the important one here being **inositol trisphosphate (IP_3)**. IP_3 opens the IP_3-receptor channels on the cell's internal calcium store, and calcium floods the cytosol. That rising calcium gates **TRPM5**, a cation channel that depolarizes the cell, and the depolarization opens **CALHM1**, through which the cell fires ATP out to the gustatory nerve fibres waiting against it. Ligand to pocket, pocket to gustducin, gustducin to calcium, calcium to depolarization, depolarization to the nerve: a bitter message, amplified at every stage, reaching the brainstem as a single word — *stop*.

Two features of this system govern the engineering. The first is that it is steeply non-linear and genetically variable. TAS2R38 is the receptor of the old PTC and PROP taste tests; a common sequence difference, the PAV versus AVI haplotypes, means one person reads a given bitter as a wall and another barely notices it. A cup cannot be tuned to a single palate, so we tune to the population — we lower the stimulus until it falls below the scorch threshold for the broad middle of drinkers, not for the one most sensitive taster in a thousand. The second feature is the one that defines the whole chapter's second boundary: caffeine is a weak bitter, and mostly not this receptor's. Caffeine engages a handful of other TAS2Rs, and only modestly — on the order of a tenth of the cup's bitter signal in a normal brew. The scorch that drives people from their coffee is overwhelmingly the chlorogenic story, not the caffeine. Which is precisely why we can leave the caffeine entirely alone and still win.

The chlorogenic pathway, and the one committed step

Chlorogenic acids are built on the universal phenylpropanoid branch, the carbon highway that runs through every vascular plant on earth, and in coffee the committed turn toward the chlorogenic pool is a single esterification. The route, from the amino acid 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	**HQT / HCT** (hydroxycinnamoyl-CoA:quinate transferase)	*p*-coumaroyl quinate
5	**C3′H** (coumaroylester 3′-hydroxylase)	**chlorogenic acid (5-CQA)**

The first three steps are shared with the making of lignin, flavonoids, and a dozen other products the plant cannot live without; touch them and you stress the whole organism. Step five is a hydroxylation. The pivot — the one node that commits carbon *specifically* to chlorogenic acid rather than to the rest of the branch — is step four, **HQT**: hydroxycinnamoyl-CoA:quinate hydroxycinnamoyltransferase.

HQT belongs to the BAHD family of acyltransferases, cytosolic enzymes that run a clean, cofactor-free chemistry. The enzyme holds an activated thioester — caffeoyl-CoA, or its precursor *p*-coumaroyl-CoA — and brings in **quinic acid** as the acceptor. A conserved histidine in the catalytic HxxxD motif acts as a general base: it deprotonates a hydroxyl on the quinate ring, which then attacks the thioester carbonyl, the acyl group transfers from sulphur to oxygen, and coenzyme A is released. A second conserved motif, DFGWG, holds the fold in shape. No metal, no added energy — just a transfer of one acyl group onto one hydroxyl, repeated a billion times in the ripening seed, each repetition laying down one more molecule of the chlorogenic store. Its near-twin, HCT, prefers shikimate as the acceptor and feeds the lignin road; HQT prefers quinate and feeds the chlorogenic pool. The difference between those two acceptors is the difference between the plant's scaffolding and the cup's harshness. We work only on the quinate enzyme.



A coffee seed in cross-section. The chlorogenic acids accumulate through the storage tissue of the green bean long before any heat reaches them; our single adjustment meters how much is ever laid down here, in the field, in the quiet months of ripening.

We tune HQT **down, and never off.** This is the discipline the whole method rests on. Chlorogenic acid is not a mistake the plant makes; it is the bean's ultraviolet screen, its antioxidant reserve, part of its defense against the insects and fungi of a tropical understory — and it is the sole source of the quinide body we love in the cup. A knockout would be a vandalism: a pale, defenseless seed and a hollow, bodiless drink. What we want is a *lower set-point* — a designed reduction in HQT activity that thins the chlorogenic load enough to lift the scorch, while leaving the seed its physiology and the cup its soul. Our genome-scale flux-balance model of *Coffea*'s phenylpropanoid network — the same modeling stack our family built to find the true rate-limiting step in other organisms, carried across unchanged — told us where that set-point lives: it identified HQT as the step that

holds the largest share of control over the chlorogenic branch, and it drew the window inside which we could turn the enzyme down without the network routing around us or the plant paying for it. The forty coffee constructs in the library are, in the main, a titration across that one window — a ladder of HQT settings, each computationally designed, each predicted score read on the Index — converging on the predicted position the roaster's dial could never hold.

From a set-point to a number on the tongue — the math

A lower enzyme set-point is a promise; the Index is where the promise is made good as a number. The bitter axis is defined, like every IFI axis, by the half-maximal activation of its reference receptor:

$$\text{IFI-B}(c) = 1000 \times EC_{50}(\text{ref}) \div EC_{50}(c)$$

Here $EC_{50}(c)$ is defined as the brew concentration — in grams of dissolved solids per litre, the same solids a refractometer reads as TDS — at which the receptor response would reach half of its maximum; for each of our designs, the model predicts this quantity directly from the computed bitter-ligand density the construct delivers. A loud, scorching coffee drives the receptor to half-maximum while still weak and watery, at a *low* concentration, so its EC_{50} is small and, inverted through the formula, its score is high. A gentle coffee needs to be brewed strong before it reaches the same half-response; its EC_{50} is large and its score is low. The reference anchor is a stated convention: $EC_{50}(\text{ref}) = 10.0$ g/L corresponds to IFI-B = 1,000 by definition on this axis, in the same spirit as the Scoville scale's dilution-to-threshold definition.

The reason this single number faithfully tracks the thing we engineered falls out of the dose-response itself. Write the receptor response to a brew of concentration *c* as a Hill function,

$$r(c) = (\alpha \cdot c)^n \,/\, [\, (\alpha \cdot c)^n + K^n \,]$$

where *K* is the receptor's own half-activation constant for the bitter ligand, *n* the cooperativity, and **α the density of bitter ligand the coffee carries per gram of solids** — the one term the bean's chemistry sets. Half-maximal response, r = ½, occurs exactly when α·c = K, that is when c = K ⁄ α. So the predicted axis $EC_{50}(c) = K / \alpha$, and substituting into the formula,

$$\text{IFI-B}(c) = 1000 \times (K/\alpha_{\text{ref}}) \,/\, (K/\alpha_c) = 1000 \times \alpha_c \,/\, \alpha_{\text{ref}}.$$

The receptor constant *K* and the cooperativity *n* cancel in the ratio. What survives is α — the bitter-ligand density — and the axis reports it **linearly.** This is the quiet elegance that makes the Index trustworthy on this axis: whatever the steepness of the bitter cascade, whatever a given taster's TAS2R38 haplotype does to *K*, the score is a clean, proportional reading of how much harsh ligand the cup delivers. And α runs straight back up our causal chain — α is set by the roast, the

roast by the green bean's chlorogenic load, the chlorogenic load by HQT flux. The number on the tongue is the enzyme set-point, carried through chemistry.

Now the worked cup. Take a conventional dark roast from an unmodified bean as the reference benchmark. Against the TAS2R38 assay it drives the receptor to half-response at **EC_{50} = 2.90 g/L**:

> IFI-B = 1000 × 10.0 ⁄ 2.90 = **3,448 ≈ 3,450.**

Take our engineered medium-dark from a clarified bean — the sweet-spot construct, designed to be roasted deliberately past where an ordinary bean would turn ashen. Our model predicts it reaches half-response at **EC_{50} ≈ 5.26 g/L**:

> IFI-B = 1000 × 10.0 ⁄ 5.26 = **1,901 ≈ 1,900.**

The attenuation is the distance between them, and it can be read three equivalent ways — as a ratio of scores, of EC_{50}s, or of delivered ligand:

> 1 − (1,900 ⁄ 3,450) = 1 − (2.90 ⁄ 5.26) = 1 − (α_eng ⁄ α_dark) = 0.449 ≈ **45%.**

Our model predicts 45% of the harsh wave lifted — comfortably inside the 20–65% band this lever is designed to deliver — the phenylindane scorch pulled down while the quinide body stays.

It helps to put a molar figure on what that 45% represents in the cup. Brew a cup the ordinary way: 15 g of ground coffee to 250 mL of water, an extraction yield near 20%, so about 3.0 g of solids dissolve — a TDS of 12 g/L. In a dark unmodified roast, the chlorogenic-derived bitter compounds, the surviving acids plus the quinides plus the phenylindanes, run on the order of 2.5% of those solids: roughly 0.075 g, or 0.30 g/L in the cup. Converting to a molar load against a 5-caffeoylquinic-equivalent mass of 354.31 g/mol,

> 0.30 g/L ÷ 354.31 g/mol = 8.5×10^{-4} mol/L ≈ **850 µmol/L of bitter equivalent.**

That 850 µM sits well up the saturating shoulder of the Hill curve — which is why a dark cup reads as a wall. Our 45% reduction in design load brings the delivered load to about **465 µM**, sliding the cup back down the curve from strongly engaged to moderately engaged. Our model predicts exactly that slide: not a cup with the bitterness deleted, but a cup with the scorch stepped off of and the body left standing on the floor beneath it.

The IFI-B axis is defined through TAS2R38 expressed in HEK293 cells carrying a promiscuous Gα16/gustducin chimera, read by Fluo-4 calcium imaging across an eight-point dilution series and fitted to the Hill equation for EC_{50}. A cross-validated trained descriptive panel scores bitter intensity on a fixed scale, with our design specification targeting a batch coefficient of variation of ≤7%. That precision is the point worth sitting with. Because IFI-B is proportional to $1/EC_{50}$, a 7% coefficient of variation carries through almost one-to-one to the

score, so the clarified bean is designed to read **1,900 ± 133** at one standard deviation, batch after batch. The roaster's conventional controls cannot approach that: origin, harvest, moisture, process, and the roast curve itself all move the bitterness, and they move it lot to lot by far more than seven percent and with no number attached. We did not make a better roast. We made a bitterness that holds still enough to put a tolerance on.

The titration ladder — forty constructs, one lever

Every construct in the library is one position of the HQT dial, computationally designed across that window. Here is the backbone of the series — the green bean's chlorogenic load, the predicted EC_{50}, the resulting bitter score, and the attenuation each delivers against the unmodified dark benchmark:

HQT ACTIVITY (% OF WILD-TYPE)	GREEN 5-CQA-EQUIV (MG/G)	PREDICTED BITTER EC_{50} (G/L)	IFI-B	ATTENUATION VS. BENCHMARK
100 (wild-type control)	60.0	2.90	3,450	0%
84	52.0	3.34	2,990	13%
70	45.0	3.86	2,588	25%
60	40.0	4.35	2,300	33%
47 (the sweet spot)	**33.0**	**5.26**	**1,900**	**45%**
39	28.0	6.21	1,610	53%
27	21.0	8.28	1,208	65%

Two things in this ladder are worth reading closely. The first is that the chlorogenic load falls *less steeply* than the HQT activity above it — drop the enzyme to 47% and the green bean's chlorogenic acid settles near 55% of wild-type, not 47%. That gap is real biology, and it is the flux-control coefficient at work: HQT holds the largest single share of control over the chlorogenic branch, but it does not hold all of it. Predicted across this titration by our model, the slope of log chlorogenic load against log HQT activity is about **0.8** — the enzyme's flux-control coefficient — because the upstream supply and the quinate pool carry the remaining fifth of the control. That number is also a warning. As you push HQT toward zero, the coefficient means each further cut buys less chlorogenic reduction than the last, while the seed's own physiology — its UV screen, its antioxidant reserve, its defense — degrades in earnest. Diminishing return above, rising cost below. The useful window is not infinite, and the model drew it closed on purpose.

The second is where the ladder earns its living. The 84% construct lifts only 13% of the harsh wave — below the threshold a drinker can reliably tell from an ordinary cup, which is exactly why the honest band this lever claims opens at 20% and not before. From the 25% rung upward the attenuation is unmistakable, and at 65% — the 27% HQT construct — we reach the floor, not of what the enzyme can do, but of what the cup can survive. The sweet spot sits at the 47% setting: HQT turned to roughly half, chlorogenic load near 33 mg/g, a 45% scorch reduction, and a green bean that still carries a generous

quinide reserve into the roast. That is the construct the roaster was always reaching for and could never hold, because he was standing on the road and we moved the gate.



Roasted to a deliberate medium-dark. On a clarified bean the roast can be carried past the point where a conventional bean turns ashen, because the harsh precursor was lowered in the field, not chased in the drum.

Reading the cup

Here is the sweet-spot cup read in full across the eight axes — a medium-dark from the clarified bean, the roast pushed with confidence:

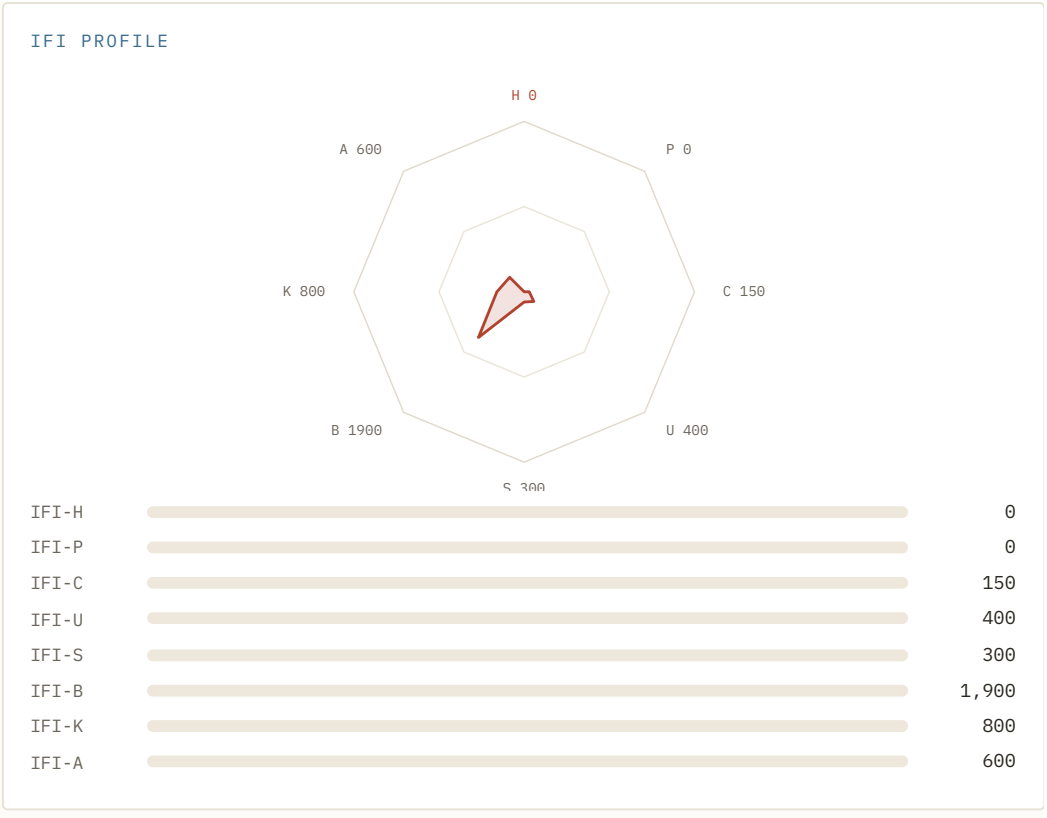

Read it from the loud end down. The **B of 1,900** is the whole chapter in one number: the harsh wave lifted 45%, the phenylindane scorch pulled down, the quinide body left on the floor. Around it the cup is quietly, truthfully itself. A **K of 800** — kokumi, the round filling continuity of a good dark roast — is wholly untouched by our lever and present as nature made it. A **U of 400**, the faint savory depth real coffee carries, sits where it always sat. An **S of 300** is the small natural sweetness that was there all along and could never be heard over the scorch; take the scorch away and the sweetness simply becomes audible, discovered rather than added. The **A of 600** is a

clean dry astringency left in on purpose, because a coffee with no dryness at all tastes slack and boneless. The **H of 0** and **P of 0** are the honest zeros of a drink with no heat and no pungency in it — coffee is not a pepper, and the ruler says so plainly.

What the profile is really recording is the roaster's new freedom. With the scorch decoupled from the body, the roast curve opens: there is room now to chase sweetness and weight deep into the dark, to the far side of second crack, without the ash arriving to shut the cup down — because the molecule that would have made the ash was lowered in the green bean, months before the drum ever turned.

Because the Index is one ruler laid across all food, this cup can be read on the same bitter line as anything else we have tuned. Set beside it the full-spectrum flagship pepper we built to register on every axis at once:

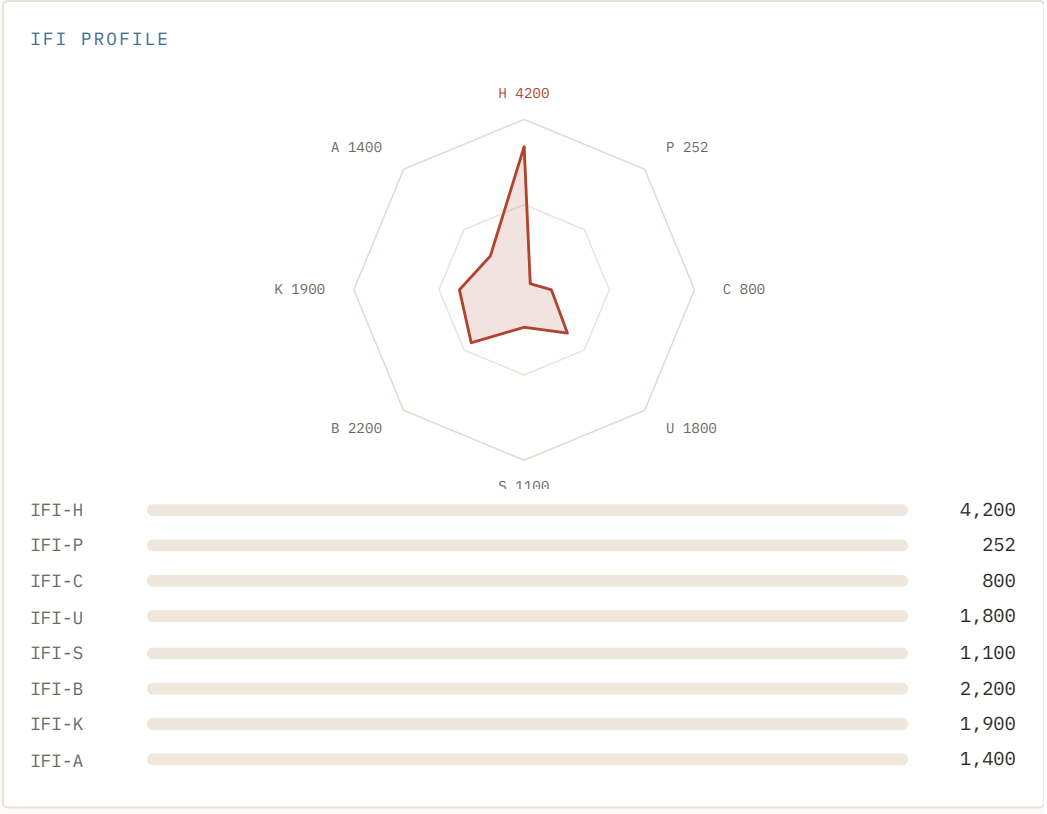

That pepper's **B of 2,200** is a bitterness we designed *upward* and held on purpose, part of a deliberate eight-axis chord. The coffee's **B of 1,900** is a bitterness we spent a titration of forty constructs learning to bring *down*. A cup of coffee and an engineered chili, read on the same line, three hundred points apart — the clearest proof we can offer that bitterness is not a fixed property of a plant but a dial, and that the dial now turns both ways.

The two boundaries we will not cross — aroma and caffeine

Two things define this work by what it refuses to touch, and both refusals are deliberate.

The first is **aroma.** What we love about coffee, more than anything the tongue reports, is its smell — the hundreds of volatile compounds that make a kitchen stop when the grinder runs. Almost none of them are made on the chlorogenic road. They are forged in the roast on other chemistry entirely: the Maillard reactions between the bean's sugars and amino acids that build the roasted, nutty, chocolate notes; the breakdown of trigonelline into the pyridines and pyrroles; the roast's work on the bean's own lipids and sugars. Our lever sits on the chlorogenic pathway and brushes those roads barely if at all — which is the entire reason the method works. We can pull the harsh bitterness down and leave the perfume standing, because the perfume was never made where the harshness was. A less careful intervention — stripping chlorogenic acid from roasted beans by solvent, say — risks dragging the aromatics out with it. We never go near the roasted bean. We change one number in the living seed and let every other road run exactly as it always has.

The second is **caffeine.** We do not touch it, and the reason is both chemical and ethical. Caffeine is made on a wholly separate route — the purine-alkaloid pathway, which has nothing to do with the phenylpropanoid branch. It begins at xanthosine and climbs through three methylations, each donated by S-adenosylmethionine and each run by a dedicated N-methyltransferase: xanthosine to 7-methylxanthosine, on to 7-methylxanthine, to theobromine, and finally to caffeine at the hand of caffeine synthase. Not one of those enzymes is HQT; not one of those intermediates is chlorogenic acid. The two pathways share no step we could disturb by accident. And because caffeine contributes only a small fraction of the cup's bitterness in the first place, there is no reason to disturb it and every reason not to: a drinker's coffee keeps its lift exactly as nature set it. We take away the scorch. We leave the morning intact — the caffeine for the body that came for it, the chlorogenic antioxidant reserve left generous in the bean, nothing stripped the cup wanted to keep. We are in the business of the one harsh wave, and of nothing else.

Failure modes, edge cases, and the honest limits

A method is only trustworthy if you can say where it breaks, so here is where this one does.

Over-attenuation collapses the body. The quinides that give the cup its floor are made from the same chlorogenic acid as the phenylindanes that scorch it — the good and the bad share a source. That is why the ladder has a floor at roughly 65% attenuation, the 27%-HQT construct: below it there is simply not enough chlorogenic acid left to build a proper quinide body, and the cup goes thin, bright, and hollow even as it goes clean. The lever lifts the scorch faster than it thins the body, which is the whole reason it is worth pulling — but it does thin the body, and past the window it thins it past repair. More attenuation is not more quality.

The baseline is a background. Our reference bean here carries 60 mg/g of chlorogenic acid, which sits where a robusta-leaning seed or a high-chlorogenic arabica would. A delicate washed arabica starts lower and a classic robusta higher, and the HQT set-point that yields a 45% cut on one is the wrong set-point on the other. Each genetic

background is re-titrated against its own starting load; the ladder above is the shape of the response, not a single switch to be thrown blind.

The roast still has to be good. We lowered the precursor; we did not abolish roast chemistry. Carry any bean far enough into the dark and you can still scorch it — there is less to scorch, and the ash arrives later and gentler, but the lever widens the roaster's safe window rather than removing the need for judgment. Grind, water, ratio, and time all still matter. We moved one dial out of the way so the others could finally be heard; we did not retire them.

Perception varies, and we tune to the room. The steep, haplotype-dependent nature of TAS2R38 means the same cup reads as a wall to one drinker and as merely rich to another. The sweet-spot construct is designed so the broad middle of drinkers falls below the scorch threshold — not the single most bitter-sensitive taster alive, for whom a gentler setting exists, nor the least, who would never have minded. It is a population-level design, honestly bounded, not a universal cure for everyone's cup.

What we will not claim. Chlorogenic acid is spoken of, elsewhere, as a human antioxidant; we leave ample of it in the bean, but we make no health claim of any kind, for the cup or the drinker. These are computational designs; Scoville Splice's constructs have not been experimentally validated in a laboratory. The precision we do stand behind is the one the design specifies: a batch coefficient of variation at or under 7%, a predicted bitter score of 1,900 ± 133 roast after roast. That number is the whole of our promise. It is also, set against a century of a single welded dial, a great deal — the one thing the roastery and the café, for all their art, were never able to do. We did not teach the roaster a better move. We went upstream, into the quiet months of the ripening seed, and moved the two bitternesses apart.

For the full life of the bean — origin and altitude and process, the whole specialty landscape that lives far beyond this one enzyme — the shelf hands you on to *Coffee* (Volume XV), where the organism receives the monograph it deserves. Here we did the single thing none of the downstream dials ever could: we took a cup's harshness and its soul, which every café in the world is forced to hold on one lever, and we set them on two.

Chapter 13 · Cold, Clear, Unsweetened — *the cup, served two ways.*

application · clarified cold brew · the unsweetened cortado

Two cups make the clearest argument this axis can make, because the plainest test of a de-bittered bean is the drink a person reaches for without reaching, afterward, for sugar. Both are built on the clarified coffee cultivars — beans we designed gene by gene so the harsh second wave of a dark roast is designed never to form in the first place — and both are deliberately, almost austerely plain, so that what reaches the tongue is the bean and the method and nothing hiding them. We are a family that has kept its table clean and its food close to the plant for as long as any of us can remember, and coffee is the one daily ritual where a whole culture has learned to bury what it grows under cream and syrup. These two cups are our answer to that habit: not a sweeter coffee, but a coffee that no longer needs to be made sweet. To understand why they work, you have to start one step before the cup — at the receptor that decides, a hundred times a second, whether a mouthful is nourishment or poison.



Coffea arabica in the research greenhouse — the cherries ripen sweet, but the bean inside carries the heaviest chlorogenic-acid load of anything we brew, the raw material of the bitterness we spend this whole chapter learning to quiet.

The sense that evolved to make us stop

Bitterness is the only taste the body built as a brake. Sweetness invites, umami reassures, salt calibrates — but bitter evolved to say *stop*, because for most of our history the bitter molecule and the toxic molecule were the same molecule, and the animals that hesitated over a bitter leaf left more descendants than the ones that swallowed. Coffee is the strangest thing a human being does with that inheritance: we take the brake, brew it hot, and drink it every morning on purpose. The axis we call Zero-Bitter is the engineering of that

paradox — not the abolition of the warning, which would make coffee taste like warm water, but the lowering of its volume to where pleasure can be heard underneath it.

The receptor doing the warning, the one the Index anchors the bitter axis to, is **TAS2R38** — a class-A G-protein-coupled receptor, one of roughly twenty-five bitter sensors studded across the taste cells of the human tongue, and the single best-characterized of them all. It is a broad-tuned sentinel: where a sweet receptor is fussy about its ligand, TAS2R38 answers to a whole chemical neighborhood of alarm molecules, and the compounds a dark roast throws off sit squarely inside its range. When one of them settles into the receptor's binding pocket, the cell runs a cascade we can follow step by step. TAS2R38 activates its bound G-protein, **gustducin**; the released Gβγ subunits switch on **phospholipase C-β2**; PLCβ2 cleaves the membrane lipid PIP_2 into the second messenger **IP_3**; IP_3 opens a channel on the cell's internal calcium store and floods the cytoplasm with Ca^{2+}; the calcium gates the **TRPM5** channel in the surface membrane; the cell depolarizes and vents ATP through its **CALHM1** pores onto the waiting gustatory nerve. From molecule to nerve impulse, the whole chain runs in well under a second, and its output is a single instruction the conscious mind does not get to overrule: *this one, be careful.*

TAS2R38 also carries the clearest fingerprint of why no two people at a table taste the same cup. The gene comes in two common forms, separated by three coding substitutions at positions 49, 262, and 296, that sort people into the **PAV** taster haplotype and the **AVI** non-taster haplotype. A PAV receptor fires at agonist concentrations a hundred- to a thousandfold lower than an AVI receptor needs; the same espresso is a pleasant roast to one drinker and a genuine ordeal to the person beside them. We cannot edit the drinker, and we would never want to. What we can do is move the stimulus — lower the concentration of bitter molecules the bean delivers — so the cup clears fewer detection thresholds, less often, for everyone at once.

What is actually bitter in a cup of coffee

Ask most people what makes coffee bitter and they will say caffeine, and they will be mostly wrong. Caffeine is bitter, but it is a minor voice — on the order of ten to fifteen percent of what the tongue reports — and it is also the molecule people are actually there for, so it is the last thing we would touch. The real bitterness of coffee is built during roasting, out of a class of compounds the green bean carries in extraordinary abundance: the **chlorogenic acids**, a family of esters formed from a cinnamic acid and quinic acid, dominated by the caffeoylquinic acids. A green arabica bean is five to eight percent chlorogenic acid by dry mass; a robusta can reach ten. Nothing else we brew is so loaded with a single precursor of bitterness.

In the green, undamaged bean those chlorogenic acids are only faintly bitter and mostly sour-astringent; they are a reservoir, not yet a warning. Roasting is what arms them, and it does so in two stages that matter enormously to the final cup. In the first stage, as the bean browns, the chlorogenic acids lose water and cyclize into **chlorogenic-acid lactones** — and these are pleasantly, legitimately bitter, the clean roast bitterness of a good medium cup, the bitterness no one is actually complaining about. Push the roast darker, hotter, longer, and a second stage begins: those lactones themselves break

down, and their fragments recombine into **phenylindanes** — and *these* are the problem. Phenylindanes are the harsh, dry, lingering second wave, the scrape that arrives after the swallow and will not leave, the taste a café drowns in milk and sugar because it has no other tool. The darker the roast, the more of the gentle lactones have been consumed and the more of the harsh phenylindanes have formed. It is a sequence, a road with a beginning and an end, and knowing it is the whole key to the engineering: both bitter families descend from one precursor pool, the chlorogenic acids, and the harsh one sits furthest downstream.



A coffee bean in cross-section — the pale endosperm is where the chlorogenic acids are stored; roasting turns that reservoir first into the clean bitterness of lactones and then, as it darkens, into the harsh lingering phenylindanes we engineer the bean never to over-produce.

COMPOUND	ORIGIN	APPEARS AT	WHAT THE TONGUE READS	RELATIVE POTENCY
Chlorogenic acids (CQAs)	phenylpropanoid pathway, green bean	green / light	sour, astringent, faintly bitter	low
Chlorogenic-acid lactones	CQA cyclization on roast	light / medium	clean, pleasant coffee bitterness	moderate
Phenylindanes	**lactone breakdown on dark roast**	**medium-dark / dark**	**harsh, dry, lingering "scrape"**	**high**
Caffeine	purine pathway (entirely separate)	all roasts	brisk bitter edge, ~10–15% of total	low–moderate

The table is the argument in miniature. If the harsh wave is a downstream product of one reservoir, then you do not have to fight the roast, or chase the phenylindane after it has formed, or mask it in the cup. You reach back, before the bean ever meets the heat, and you lower the reservoir.

Clarification — reaching back up the phenylpropanoid road

The chlorogenic acids are built by a pathway that runs, in one form or another, in every vascular plant on earth: the **phenylpropanoid branch**, PAL → C4H → 4CL → HCT → COMT. Phenylalanine is stripped to cinnamic acid by phenylalanine ammonia-lyase (PAL); cinnamate is hydroxylated to *p*-coumaric acid (C4H) and activated to its coenzyme-A ester (4CL); then a hydroxycinnamoyl transferase — in coffee the quinate-specific form, **HQT** — joins that activated acid to quinic acid, and a final hydroxylation (C3′H) yields chlorogenic acid itself. This is the road the bean's entire bitter reservoir travels down, and the beauty of it, for an engineer, is that it has a single committed

bottleneck. HQT is the enzyme that commits carbon to the chlorogenic-acid pool. Turn its throttle down and the whole reservoir falls.

That is **clarification**: a graded attenuation of HQT, supported by a lighter tuning of the C3′H branch, so the green bean accumulates a smaller, computed fraction of chlorogenic acid than nature would otherwise build — and therefore, after the roast, our model predicts less of both bitter waves in proportion. And because the harsh phenylindanes sit two degradation steps downstream of the reservoir, requiring more precursor throughput to form, they fall *faster* than the reservoir does. Across the roast window we work in, the delivered phenylindane load scales with roughly the three-halves power of the green-bean chlorogenic-acid content: cut the precursor by half, and the harsh wave drops by closer to two-thirds. The engineering is aimed at the gentle precursor; the nonlinearity of the roast chemistry hands us extra leverage on the harsh one, for free.



The design bench where a clarified coffee construct takes shape — every edit targets the HQT locus, and the caffeine pathway is left exactly as the plant built it, because we de-bitter the bean, we never decaffeinate it.

Two disciplines govern which levers we pull. The first is that we never touch the caffeine. Caffeine is made by a completely separate route — the purine pathway, running xanthosine to caffeine through a trio of N-methyltransferases — and because it shares no enzyme with the phenylpropanoid branch, we can empty a computed fraction of the bitter reservoir and leave the caffeine exactly where the plant put it. The clarified bean is a de-bittered bean, not a decaffeinated one; the morning still works. The second discipline is that we never reach for zero, and the reason is nourishment. Chlorogenic acids are among the most studied antioxidant compounds in the human diet, and the bean did not build them to spite us; it built them for its own defense, and we inherit their benefit when we drink them. To strip the reservoir to

nothing would be to hand back a cup that tastes clean and carries nothing — the exact failure this whole axis exists to avoid. We attenuate. We do not abolish.

CONSTRUCT CLASS	TARGET NODE	EDIT	EFFECT	DELIBERATELY UNTOUCHED
Reservoir throttle	**HQT** (hydroxycinnamoyl-CoA : quinate transferase)	graded promoter attenuation to 40–60% activity	predicted to lower the chlorogenic-acid pool at its committed step	caffeine, sugars, aroma precursors
Branch tailor	**C3′H**	tune caffeoyl : coumaroyl flux	fewer phenylindane precursors reach the roaster in our model	melanoidin / browning chemistry
Ferulate trim	**COMT**	modest knockdown	predicted to lower feruloylquinic bitterness	structural lignin minimum
Caffeine pathway	CaDXMT1 and the purine route	*none*	caffeine fully preserved	the reason people drink it

Coffee carries 40 of these constructs today, and the platform rates the species **Hard** — the reservoir is huge, the roast chemistry is downstream of the edit and therefore one more variable to control, and the margin between "quieted" and "hollowed out" is narrow. It is, in every sense, a problem you have to be able to *measure* to solve.

Putting a number on the scrape

Loudness, not quantity, is what the Index measures, and loudness is anchored to the half-maximal concentration. **EC_{50}** is the concentration of a stimulus at which its receptor reaches half of its full response. A potent bitter needs very little to half-saturate TAS2R38, so its EC_{50} is low; an attenuated one needs a great deal, so its EC_{50} is high. The Index inverts that relationship and scales it by a thousand, so that a sharper bitterness scores *higher*:

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

We anchor the bitter axis to **PROP** (6-*n*-propylthiouracil), the cleanest and most reproducible TAS2R38 probe in the literature, with a stated convention constant of EC_{50}(reference) = **1.8 µM** — a defined anchor for the IFI-B axis, in the same spirit as the Scoville dilution definition rather than derived from a literature survey. By construction the reference scores exactly 1,000, because IFI-B(ref) = 1,000 × 1.8 ÷ 1.8 = 1,000. Everything bitterer than the probe scores above a thousand; everything milder scores below. A conventional café dark roast, pulled as straight espresso, half-saturates the receptor at a PROP-equivalent of about 0.56 µM — it is more potent than the laboratory bitter standard itself — and so reads:

> $\text{IFI-B} = 1{,}000 \times 1.8\ \mu M \div 0.56\ \mu M \approx \mathbf{3{,}200}$

That single figure is the problem stated plainly. An ordinary dark roast is more than three times as loud, on the bitter axis, as the standard probe chemists use to calibrate the receptor. The cup is not faintly bitter; it is, chemically, shouting — and the whole café apparatus of

sugar and foam exists to shout back. Read as a full profile, the conventional cup looks like this, a single axis drowning out everything the bean could otherwise say:

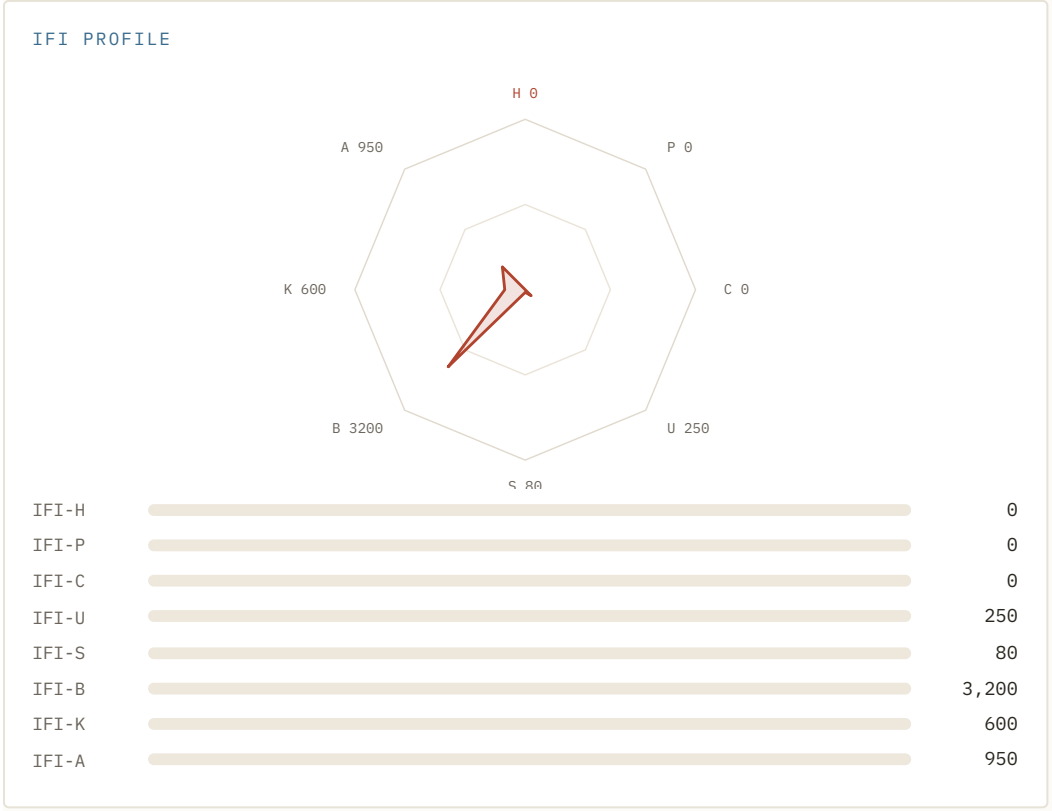

Now watch what clarification does to that number. Because the receptor's EC_{50} for a given bitter molecule is a biophysical constant, lowering the *yield* of the harsh compounds by a fraction *a* means the cup must be present in greater quantity to reach half-max — its food-level EC_{50} rises by 1 / (1 − a) — and the IFI score falls in exact proportion. Across the platform's established attenuation window of **20–65%**, our model predicts a coffee ladder like this, every row computed at identical roast and extraction conditions so that only the bean's design changes:

DESIGN POINT	HARSH-WAVE YIELD	EC_{50} (PROP-EQ, MM)	IFI-B	BATCH CV
Conventional dark roast (wild-type)	100%	0.56	3,200	~38%
−20% clarification	80%	0.70	2,560	≤7%
−35% clarification	65%	0.87	2,080	≤7%
−50% clarification	50%	1.13	1,600	≤7%
−65% clarification	35%	1.61	1,120	≤7%

At the far edge of the window, our model predicts a 65% clarification drops the cup's loudness from 3,200 to 1,120 — from more than triple the reference to just above it. The coffee has not been made bland. It has been brought from a shout to a speaking voice, and left there, so the rest of the bean can finally be heard.

Why half the molecules is more than half the relief

The ladder above is the honest potency, measured the way pharmacologists measure it — on diluted extracts, in the receptor's clean linear range. But the cup a person actually drinks is not diluted into that range; it sits high up on the receptor's saturation curve, and that geometry is quietly generous to us. The receptor does not answer concentration linearly. It follows a sigmoid — the Hill relationship, $\theta = [A]^n / (EC_{50}^n + [A]^n)$ — and that curve is steep in its middle and flat at its top.

Take the phenylindanes, the harsh wave, with a receptor EC_{50} near 110 µM and a Hill coefficient of about 1.1. A strong dark-roast cup delivers on the order of 160 µM of phenylindane-equivalent to the receptor field — just over the inflection — for a fractional receptor firing of

```
θ = 160^1.1 / (110^1.1 + 160^1.1) ≈ 0.60.
```

Now clarify the bean to the −50% point, halving the delivered harsh dose to 80 µM:

```
θ = 80^1.1 / (110^1.1 + 80^1.1) ≈ 0.41.
```

A 50% chemical cut has produced only a 31% drop in receptor firing — which sounds like a loss until you see the direction of the error. Perception compresses the receptor signal *further*, so the drinker reports not "a third less bitter" but "most of the bitterness is gone." The leverage lives in the steepness of the curve: engineering is most powerful exactly when it slides the delivered dose off the saturated plateau and down into the steep middle of the sigmoid, where every molecule removed is felt. It is also why the window closes at 65% and not beyond. By the −65% point we are already well down the steep slope; squeeze further and we pay in antioxidants and in the plant's own resilience for diminishing sensory return. The ceiling is not a limit of the method. It is the place where taste, nourishment, and the health of the plant all tell us the same thing: stop here.

A word on magnitude, by the unit, because the potency is easy to underestimate. At 160 µM of a phenylindane near 236 g/mol, a 240 mL cup carries only about **9 milligrams** of the harsh fraction — nine milligrams driving the receptor to sixty percent of full alarm. That is how loud these molecules are, and how little of the reservoir a roast has to misspend to ruin a cup. The engineering is not moving mountains; it is moving milligrams, precisely.

And then there is consistency, which for a daily ritual matters as much as the mean. Field coffee swings hard — origin, altitude, varietal, processing, and roast all push the bitter load around, so a wild-type line carries a batch coefficient of variation near 38%. At an IFI-B of 3,200 that is a standard deviation of roughly ±1,220: one bag lands at 1,980 and the next at 4,420, and the drinker cannot tell which morning is coming. Our clarified lines are designed to hold a batch CV at or below **7%**. At the −50% design point that is 1,600 ± 112 — the same quiet cup, harvest after harvest, roast after roast. Holding it there is not free: because the harsh wave scales with the three-halves power of the green-bean reservoir, a given variation in chlorogenic acid is magnified half-again in the cup, so to deliver ≤7% at the receptor we have to hold the green bean's chlorogenic-acid content to a coefficient of variation near 4.7%. The tolerance in the cup is won in the field and at the locus, not granted by the roast.

The cup, first way — clarified cold brew

Cold brew is the right place to begin, because the method already does, crudely, a little of what the bean now does precisely. Steeped cold and slow, coffee extracts far less of its harshest, most soluble

compounds than it does under hot water — which is the whole reason cold brew is famous for being smooth. The two effects, bean and method, compound. Begin from a reservoir already lowered at the gene, extract it by a method that pulls the harsh wave reluctantly, and the reductions multiply rather than add.

The arithmetic is worth seeing, because "compound" is a precise claim, not a figure of speech. Call the harsh dose a conventional bean delivers under hot water 160 µM. Cold brewing alone, on that same conventional bean, pulls perhaps 45% less of the harsh fraction — a factor of 0.55, down to 88 µM. A clarified bean at the −50% point contributes its own factor of 0.50. Do both, and the delivered harsh dose is 160 × 0.55 × 0.50 ≈ **44 µM** — a receptor firing of θ ≈ 0.27, against the conventional hot cup's 0.60. That is a 55% drop in felt bitterness before you have added a single thing, and the concentrate has not yet been cut for serving. Pour it one-to-one over ice and the delivered dose halves again, to about 22 µM, and the receptor barely registers it at all.

```
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 more through paper for a clean cup. Do not press the grounds; pressing drags the harshness back in.
4. Keep the concentrate refrigerated. Serve over ice, cut one-to-one with water or milk.

Our model predicts a concentrate profile like this:

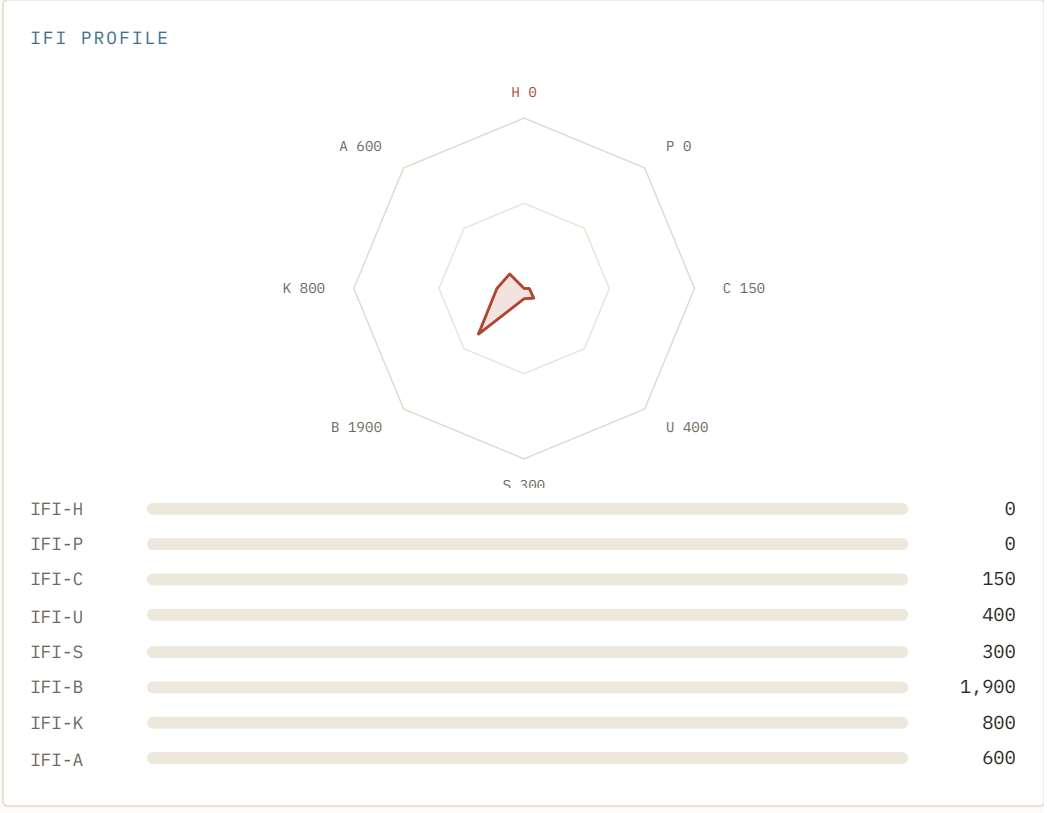

Read that predicted concentrate profile and the whole method is legible in it. Our model places **B at approximately 1,900** — this is a concentrate, computed at full concentration and carrying the full

day's extraction in every milliliter; cut one-to-one in the glass, it drops to the low thousands, below the −50% design point on the ladder, exactly as the compounding predicts. Underneath that lowered bitterness, our predicted profile reveals numbers that a conventional cold brew buries: a **U of 400**, the bean's own quiet savory body, readable now that nothing is shouting over it; an **S of 300**, a real whisper of sweetness that was always there and never had the room; a **C of 150**, the cool, clean edge of a drink served cold and clear. The cold method keeps the astringent **A to 600**, the dry tannic finish stripped back by slow extraction. This is what a de-bittered bean sounds like in our model when a gentle method agrees with it: not emptier, but *clearer* — the same coffee with the scrape lifted off.

The cup, second way — the unsweetened cortado

The second cup is the one that tends to end the argument, because it gives the bean nowhere to hide. A cortado is espresso cut with a little warm milk and nothing else — a small glass, a double shot, sixty milliliters of milk steamed barely to body temperature. Its whole character is exposure. The volume is tiny, the milk is a cut and not a blanket, and if the bitterness underneath is harsh, the drink is harsh, and the only fix a café has ever had is the sugar spoon. From a clarified bean, it does not need the spoon.

```
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, ~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. Taste it before you reach for anything — most people don't reach.

The milk is doing more than diluting, and the mechanism is a small, lovely echo of the tongue's own chemistry. Milk's dominant proteins, the caseins, are unusually rich in **proline**, and proline-rich proteins have a particular talent: they bind polyphenols — tannins, chlorogenic compounds, the phenylindanes themselves — through hydrogen bonds and hydrophobic contacts, wrapping the bitter molecule up so it can no longer reach a receptor. This is precisely the trick the body uses to blunt astringency: the salivary proteins that the Index anchors the astringent axis to are themselves proline-rich, built to precipitate tannins before they scour the mouth. When you pour warm milk through a cortado, its caseins do to the coffee's harsh fraction what your own saliva does to the tannins in a cup of tea — they take a computed share of the free ligand out of circulation. Bind a third of it, and the delivered dose drops a third, and on a bean already clarified there is very little left for the milk to even catch.

Our model predicts a cortado profile like this:

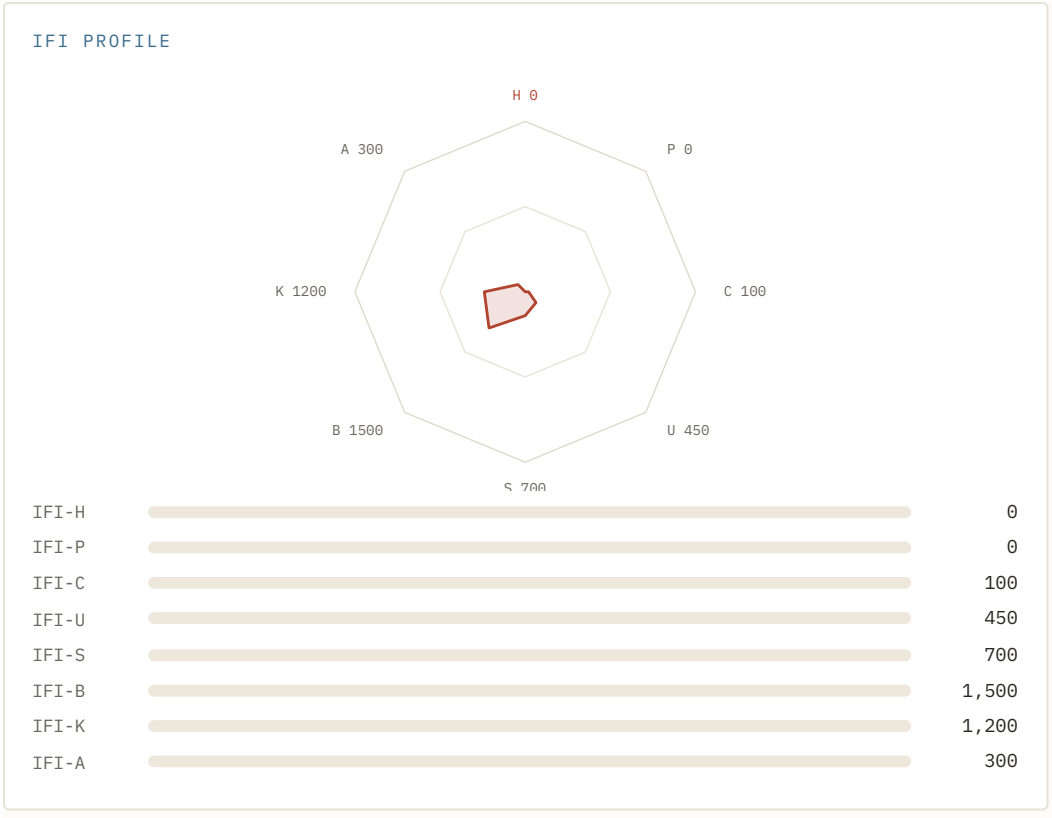

The cortado's numbers tell the smaller, subtler story, and it is the one that changes people. Our model predicts **B at approximately 1,500** — lower than the straight concentrate, partly the clarified bean and partly the milk's casein binding — and look at what the quiet uncovers. An **S of 700**, most of it the milk's own lactose but some of it, unmistakably now, the bean's sweetness, finally readable. A **K of 1,200**, the combined body of milk and a good dark roast, round and filling and continuous across the sip. A **U of 450**, the savory depth a bitter cup never lets you notice. And an **A down at 300**, just enough dryness to keep the glass from going flat, the astringent finish softened by the same proteins that softened the bitterness. A drinker who has always taken coffee sweet discovers, often with real and specific surprise, that the sugar was never balancing the coffee at all. It was burying a harshness that, on this bean, was never there to bury. That is the whole case for the axis in a four-ounce glass.

The honest limits

We would rather name the edges of this than let two beautiful cups overstate them. The first limit is the drinker's own genome. A PAV supertaster, whose TAS2R38 fires at concentrations an AVI non-taster would never notice, will still taste more bitterness in the clarified cup than their companion does — we have lowered the stimulus for everyone, but we have not flattened the spread between tongues, and the Index number is anchored to a reference haplotype, not a promise that all mouths now agree. The second is that clarification and decaffeination are different operations on different pathways; a clarified bean is still a full-caffeine bean, and anyone seeking a decaffeinated cup is asking for an edit we deliberately did not make. The third is that the bean is not the whole cup. Origin, roast level, grind, water chemistry, and extraction time still move the bitterness by amounts that can rival a design point; hand a clarified bean to a careless roaster or a scalding over-extraction and some of the harsh wave will come back, because we lowered the reservoir, not the laws of the roast. And the fourth is the one the construct count already confesses: at 40 constructs and a Hard rating, coffee's margin between quieted and hollowed is narrow, the HQT throttle must be held with care so the plant keeps enough of its own defense

chemistry to grow and to nourish, and the ≤7% CV in the cup is won bean by bean, season by season, never granted by the pathway. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

None of that is a retreat. It is the shape of honest work: a reservoir lowered at the gene, a harsh wave that our model predicts will not fully form, a receptor signal estimated to the half-maximal micromolar, and two plain cups — one cold and clear over ice, one small and warm in a glass — that ask a person to taste the coffee itself and, nearly always, to set the sugar down. Not a sweeter coffee. A coffee that was never, once we understood it, in any need of being made sweet.



PART V OF VI

Hops, Cacao, and the Polyphenol Neighbors

the bitterness we keep · theobromine and the flavan-3-ols · tea, grape, and the pith

Chapter 14 · The Bitterness We Keep — *hops, IBU, and the shape of a bite.*

DOMAIN · HOPS · Humulus · 80 constructs · the iso-alpha-acids

Beer is the one place in the whole kitchen where bitterness is deliberately *added,* carefully measured, and openly bragged about. Everywhere else on this axis a cook is fighting bitterness — scrubbing it out of a cabbage, roasting it down in a coffee, masking it in a square of dark chocolate. Only the brewer invites it in, weighs it to the milligram, and prints the number on the label. It is **IBU**, the International Bitterness Unit; it is the brewer's signature; it is the counterweight that keeps a malt-sweet beer from collapsing into syrup. So the work we do on the hop is unlike the work we do anywhere else on the bitter axis. We are not here to take the bitterness away. We are here to change its *shape.*



A hop cone in full light — Humulus lupulus, the single plant whose resin glands carry both the bite of a beer and its perfume, kept on two separate chemical roads.

The one bitterness we are asked to keep

The complaint a brewer actually brings us is never "too much bitterness." A brewer who wanted less character would reach for the simplest tool in the world — fewer hops — and walk away. The real complaint is that the bitterness *lands wrong.* It is harsh where it should be firm. It clings where it should finish clean. It arrives as a scrape at the back of the throat instead of a crisp invitation to the next sip. Two beers can carry the identical IBU on the label and feel like two different drinks: one full and firm and *gone,* the other full and firm

and *clinging.* That gap — between the same amount of bitterness and two opposite experiences of it — is the whole territory of this chapter. Bitterness, more than any other sensation we design, is not one quantity but two: a size, and a shape. The hop is where we learned to hold the size steady and move the shape underneath it.

It helps to say plainly what IBU is, because it is the rare flavor that already came to us with a number attached. Where heat waited ten thousand years for the Scoville scale and then for a receptor to anchor it, bitterness had a package figure long before we arrived — the brewer's own, hard-won and jealously guarded. IBU is, to a first approximation, the milligrams of iso-alpha-acids dissolved in a liter of finished beer, read as the absorbance of an acidified beer extract at **275 nanometers.** A hoppy pale ale runs 40 to 70 IBU; a double IPA can push past 100; a delicate lager sits near 10. Our job is not to argue with that number. It is to let the brewer keep it exactly where they are proud to have it, and change only how it finishes on the tongue.

What the tongue actually does with a bitter molecule

To reshape a bite we first have to be exact about how a bite is read, because the shape we are designing is a shape *in the receptor,* not in the glass.

Bitterness is sensed by the **TAS2R family** — a set of roughly twenty-five G-protein-coupled receptors studded across the type II taste cells of the tongue. Unlike the sweet and umami receptors, which are few and narrowly tuned, the bitter receptors are many and broadly tuned: a single TAS2R will answer to dozens of unrelated bitter molecules, and a single bitter molecule will light up several TAS2Rs at once. It is an alarm system, evolved to flag a wide chemical class — *possible toxin* — rather than to identify any one compound. On the Index we anchor the bitter axis, IFI-B, to **TAS2R38,** the best-characterized member and the one whose half-activation defines the scale's reference point. The iso-alpha-acids of beer are promiscuous TAS2R agonists; we read their loudness against that fixed TAS2R38 standard, the way a sound engineer reads every instrument against a single calibration tone.

What happens after the molecule lands is a cascade worth walking step by step, because every link in it is a place where loudness is decided:

1. A bitter ligand seats in the TAS2R binding pocket and shifts the receptor to its active conformation.
2. The active receptor exchanges GDP for GTP on its G-protein, **gustducin,** releasing the Gβγ subunits.
3. Gβγ switches on **phospholipase C-β2,** which cleaves membrane PIP_2 into **IP_3** and diacylglycerol.
4. IP_3 opens the **IP_3-receptor** on the internal calcium store, and intracellular Ca^{2+} climbs.
5. The calcium rise gates **TRPM5,** a cation channel, and the cell depolarizes.
6. Depolarization opens **CALHM1,** which releases ATP as the neurotransmitter to the afferent gustatory nerve — and the brain reports *bitter.*

Anchor one number to that whole chain and it is the half-maximal effective concentration — **EC_{50}** — the amount of stimulus that drives the receptor to half of its full response. This is the quantity the Index is built on:

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

Read it carefully in the direction that matters here. A compound that activates TAS2R38 at a *lower* concentration is more potent — a smaller EC_{50} — and so it scores *higher* on IFI-B. It is louder. To soften a bitterness without removing it, then, we do not need to deliver less of it. We need to deliver a mixture that the receptor hears as *quieter per milligram* — a mixture with a higher EC_{50}. Hold the milligrams on the label; raise the EC_{50} underneath; the bite stays and the loudness drops. That is the entire mathematical move, and the rest of this chapter is how a hop is persuaded to make it.

But the Index's single number hides a second variable that is the real subject of a brewer's complaint: *time.* The receptor does not only turn on; it turns off. The onset is governed by how fast the ligand finds its pocket, the association rate; the finish is governed by how fast it leaves, the dissociation rate k_off. A bitterness that "clings" is a bitterness whose molecule sits in the pocket too long — a slow k_off, a receptor still firing after the swallow. The perceived duration of a bite runs roughly as $\tau \approx 1 / k_off$. So the "shape of a bite" is, quite literally, two rate constants: a rise and a decay. When we talk about reshaping bitterness, we are talking about editing the molecular ensemble so that both its EC_{50} and its k_off come down together — quieter, and gone sooner.



Hop cone in cross-section — the lupulin glands hold two separate chemistries, the bittering acids and the aroma oils, which is the whole reason a hop's bite can be reshaped without touching its perfume.

Where the bite is born: alpha-acids and the boil

Cut a hop cone and you find its chemistry concentrated in bright yellow beads at the base of the bracts — the **lupulin glands.** They are packed with resins, and the resins that matter for bitterness are the **alpha-acids:** principally **humulone, cohumulone,** and **adhumulone.** On their own these are only mildly bitter and stubbornly insoluble; a hop chewed raw is grassy and resinous, not bracing. The transformation is the boil. Sustained heat drives an **isomerization** — a ring contraction from the six-membered humulone skeleton to a five-

membered cyclopentanone — converting each alpha-acid into its pair of **iso-alpha-acids,** the *cis* and *trans* isomers. These are the intensely bitter, beer-soluble compounds that give an IPA its grip. Bitterness in beer is not harvested; it is *cooked into being,* one isomerization at a time, over the length of the boil.

Each alpha-acid is built in the gland on a short, lineage-specific polyketide assembly line, and the branch that decides *which* alpha-acid you get is set at the very first step — by the starter unit the synthase chooses to load:

ALPHA-ACID	ACYL STARTER (COA)	AMINO-ACID SOURCE	ISO-PRODUCT AFTER BOIL
Humulone	isovaleryl-CoA	leucine	*cis*- / *trans*-isohumulone
Cohumulone	isobutyryl-CoA	valine	*cis*- / *trans*-isocohumulone
Adhumulone	2-methylbutyryl-CoA	isoleucine	*cis*- / *trans*-isoadhumulone

The three differ only in the acyl side chain hung off the ring — a four-carbon branch on humulone, one carbon shorter on cohumulone, a methyl-branched four on adhumulone — and that small difference in a tail is exactly what the receptor reads as a difference in *shape.*

Before we go on, the number on the label deserves one honest line of arithmetic. The standard bitterness assay extracts an acidified beer into iso-octane and reads the absorbance at 275 nm, then multiplies:

$$\mathrm{IBU} = 50 \times A_{275}$$

The factor of 50 folds in the extinction of the iso-alpha-acids and the dilution of the method, and the convention is tuned so that one IBU corresponds closely to **one milligram of iso-alpha-acids per liter.** So a 45-IBU pale ale is carrying on the order of 45 mg/L of dissolved iso-alpha-acids. We will hold that number fixed through everything that follows — because the brewer will.

The lever is cohumulone — the shape of the harsh fraction

Not all of the bite is equal, and here is the lever. The *quality* of a hop's bitterness depends heavily on the balance among its alpha-acids, and above all on the fraction that is **cohumulone.** A cone weighted toward cohumulone gives a harsher, sharper, more lingering bitterness; a cone weighted toward humulone gives a rounder, softer, cleaner bite at the very same IBU. Translate that into the two rate constants and it is precise: **iso-cohumulone** — the boiled product of the valine-derived, one-carbon-shorter acid — seats in the TAS2R pocket more avidly (a lower EC_{50}) and leaves it more slowly (a smaller k_off). Lower EC_{50} is *louder.* Smaller k_off is *clinging.* The harsh fraction of a beer is, almost exactly, its iso-cohumulone fraction. To soften the bite is to lower the cohumulone and raise the humulone — to pour the same iso-alpha-acids but weighted toward the molecule the receptor hears quietly and releases quickly.

Which means the real question is upstream, in the gland, at the moment the synthase picks a starter. The alpha-acids are assembled by a short route we can name enzyme by enzyme:

STEP	ENZYME	REACTION	PRODUCT
1	**VPS** — valerophenone synthase (a type-III polyketide synthase)	condenses an acyl-CoA starter with three malonyl-CoA units	`phlorisovalerophenone (humulone lineage) / phlorisobutyrophenone (cohumulone lineage)`
2	**HIPT1** — aromatic prenyltransferase	first dimethylallyl prenylation	`deoxyhumulone`
3	**HIPT2** — aromatic prenyltransferase	second prenylation + oxidation	**`humulone`** `(alpha-acid)`

The fork is entirely at step 1. VPS will accept whichever branched-chain acyl-CoA is handed to it: **isovaleryl-CoA** (from leucine) sends the line toward humulone; **isobutyryl-CoA** (from valine) sends it toward cohumulone; **2-methylbutyryl-CoA** (from isoleucine) toward adhumulone. Those starters are minted by a **branched-chain aminotransferase (BCAT)** and the branched-chain keto-acid dehydrogenase complex, drawing on the free pool of valine, leucine, and isoleucine in the gland. So there are two clean places to move the ratio: change the *supply* of the starter units, by tuning the branched-chain amino-acid flux feeding BCAT; or change VPS's *preference* between them, by editing the residues that line its starter-binding pocket so isovaleryl-CoA is favored over isobutyryl-CoA. That is the 80-construct hop program stated in a sentence — not fewer alpha-acids, but a *better-balanced* set of them, humulone up and cohumulone down, so the same IBU lands softer and clears sooner. Every construct in the set is a different point on that single trade: how far to push the ratio, and by which of the two routes.

And the aroma stays wholly intact, because a hop's perfume is made on a completely separate road. The pine and citrus and resin and tropical fruit that drinkers chase through an entire style of beer are **terpenes and terpenoids** — myrcene, humulene, caryophyllene, linalool, geraniol — assembled in the same glands but by the isoprenoid pathway, through geranyl and farnesyl diphosphate and a suite of terpene synthases, with no dependence on the alpha-acid machinery at all. We edit the acids and never touch the oils. The brewer keeps every note of the aroma and loses only the harsh edge of the bite.

The math of a softer bite

Now the move in numbers, worked all the way through, because on this axis the arithmetic *is* the design.

Take a dry-hopped pale ale brewed from a rebalanced, low-cohumulone cone. Read against the TAS2R38 reference, its harsh fraction scores **IFI-B = 1,600** — the loudness of its bitterness, holding its IBU where the brewer set it. The conventional high-cohumulone cone the brewer started from scored higher. How much higher? Run the Index definition backward on both. IFI-B is inversely proportional to EC_{50}, so the ratio of the two scores is the inverse ratio of their half-activation concentrations:

> $EC_{50}(\text{rebalanced}) \,/\, EC_{50}(\text{conventional}) = \text{IFI-B}(\text{conventional}) \,/\, \text{IFI-B}(\text{rebalanced})$

The rebalanced beer is softened about **38 percent** from the cone we started with. A 38 percent drop in loudness means IFI-B(rebalanced) ÷ IFI-B(conventional) = 0.62, so the conventional cone scored:

> IFI-B(conventional) = 1,600 ⁄ 0.62 ≈ **2,580**

and the EC_{50} shift that bought the softening was:

> EC_{50}(rebalanced) ⁄ EC_{50}(conventional) = 2,580 ⁄ 1,600 = **1.61×**

A clean reading falls out of that one number. We raised the concentration needed to half-activate the bitter receptor by **61 percent** — the receptor now needs 1.61 times as much of our iso-alpha-acids to shout as loudly — and because the brewer holds the delivered amount (the IBU) *constant,* that 61 percent rightward shift in EC_{50} shows up on the tongue as a 38 percent drop in perceived loudness. Same label, quieter bite. Note the asymmetry — a 61 percent change in the molecule's potency yields a 38 percent change in the sensation — and hold onto it, because it is the first sign that *where you sit on the dose-response curve* matters as much as how far you move it.

So let us put the beer on that curve. Receptor response follows a Hill relation,

> $\theta = [L]^n \,⁄\, (EC_{50}^n + [L]^n)$

where θ is the fraction of maximal response, [L] the delivered iso-alpha-acid concentration, and n the Hill coefficient. For a clean worked example take n = 1 (the true value for iso-alpha-acids runs modestly above 1; we will return to that as a limit). Hold the serving at [L] = 45 mg/L — our 45-IBU pale ale — and suppose the conventional bittering fraction half-saturates TAS2R38 near EC_{50} = 28 mg/L (an illustrative anchor, to make the curve concrete). Then:

- **Conventional:** θ = 45 / (28 + 45) = 45 / 73 = **0.616** — the receptor runs at 62 percent of full.
- **Rebalanced:** EC_{50} rises by 1.61× to 45.1 mg/L, so θ = 45 / (45.1 + 45) = 45 / 90.1 = **0.499** — just under half.

The receptor occupancy fell from 0.616 to 0.499 — a **19 percent** drop in what the nerve actually reports at the glass — even though the Index score, measured by EC_{50} alone, fell 38 percent. Both numbers are right; they measure different things. IFI-B is the molecule's *intrinsic* loudness, anchored to its EC_{50} and independent of how strong you pour it. Perceived intensity at a given serving is where that loudness lands on the Hill curve, which flattens as it saturates. A brewer dosing 90 IBU sits high on the shoulder of the curve and will feel less of a given EC_{50} shift than one dosing 30 IBU on the steep part. The Index gives the honest, dose-free measure of the shape we changed; the drinker's tongue applies the dose. A good flavor design keeps both in view.

The 80 constructs lay out the whole trade as a ladder. Because the harsh loudness at matched IBU scales with the iso-cohumulone carried into the kettle, IFI-B tracks the cohumulone fraction almost

linearly — anchor a high-alpha bittering cone near 40 percent cohumulone at IFI-B 2,580, and every step down the fraction is a step down the harsh score:

CONE / CONSTRUCT	COHUMULONE (% OF A-ACIDS)	HARSH ATTENUATION	IFI-B (HARSH FRACTION)	EC_{50} MULTIPLE
Wild-type high-α bittering	~40%	baseline	2,580	1.00×
Rebalance I	~32%	20%	2,064	1.25×
Rebalance II	~28%	30%	1,806	1.43×
Target — dry-hop pale ale	**~25%**	**38%**	**1,600**	**1.61×**
Deep rebalance	~20%	50%	1,290	2.00×
Program floor	~14%	65%	903	2.86×

The attenuation band runs the full 20-to-65-percent width the bitter axis allows, but read where the program actually *lives:* the top of the ladder, the 20-to-40-percent rows. We ship the target line. The bottom two rows are where the chemistry *could* go and where we choose not to, because a hop stripped of its bite is a hop that has stopped doing its job. On a cabbage we remove; on a coffee we remove; on the hop, uniquely, we shape — and shaping means stopping well short of the floor.

Two beers, one label

Here are the two cones rendered as full profiles, the conventional bittering lot beside the rebalanced one — the same beer recipe, the same IBU, the two cones the only variable:

The conventional, high-cohumulone lot:

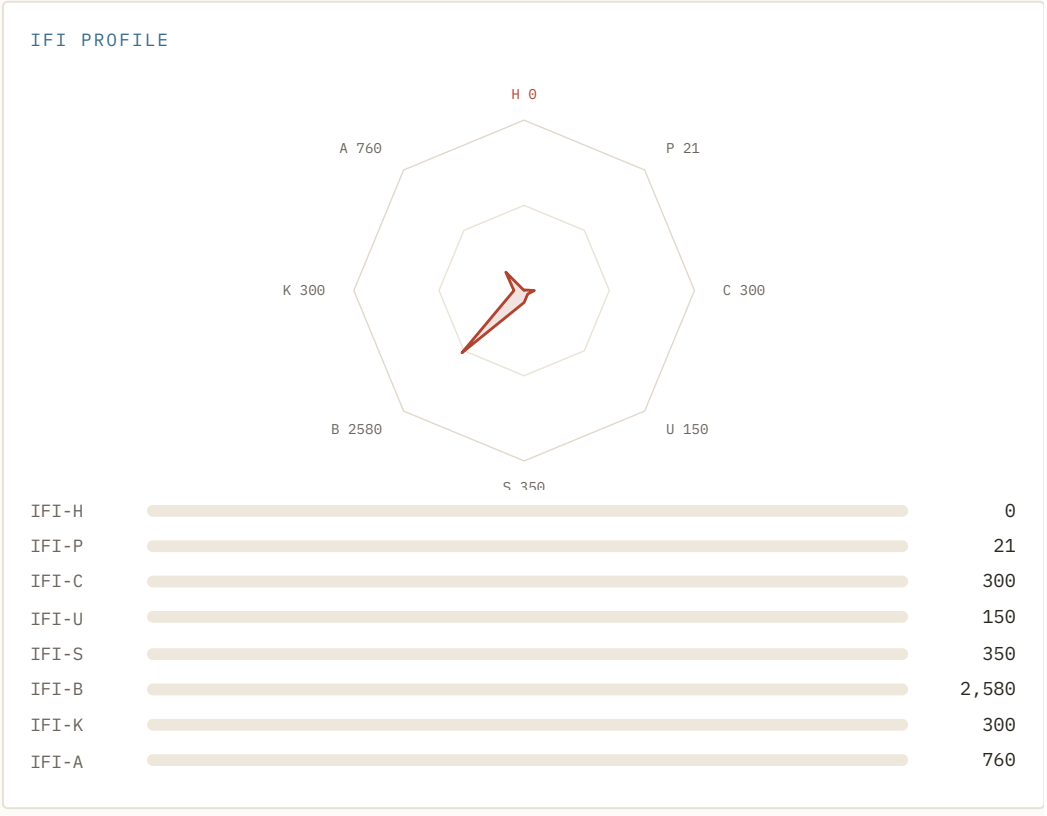

The rebalanced, low-cohumulone lot:

IFI PROFILE

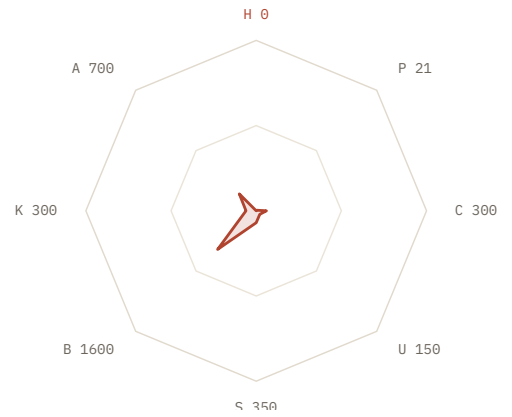

IFI-H		0
IFI-P		21
IFI-C		300
IFI-U		150
IFI-S		350
IFI-B		1,600
IFI-K		300
IFI-A		700

Set them side by side and almost everything is deliberately identical, because almost everything *is* identical — we edited one branch of one pathway. **H sits at 0:** there is no capsaicin here, no heat, and none implied. **C holds at 300** on both: the bright, almost cooling lift of the hop oils, the myrcene and linalool we never went near, unchanged to the last note. **S holds at 350,** the malt sweetness; **U at 150** and **K at 300,** the body and continuity the grain and yeast contribute; **P at 21,** a faint resinous sharpness. The oils, the malt, the mouthfeel — untouched, because we built nothing on their roads.

Only two numbers move. **B falls from 2,580 to 1,600** — the 38 percent softening, the whole point, the harsh fraction brought down by rebalancing the acids while the IBU on the label stays exactly where the brewer is proud of it. And **A eases from 760 to 700,** a smaller change we did not engineer directly: the *clinging,* the scrape at the back of the throat, is not pure bitterness. A lingering bite recruits a faint astringent cling — the slow k_off of iso-cohumulone reads partly as a drag against the salivary proteins — so when the harsh fraction clears sooner, a sliver of perceived astringency clears with it. We moved B on purpose; A followed, because on the tongue these sensations are neighbors. Note, too, that B at 1,600 sits comfortably below the per-axis ceiling of 5,000 and lower in the bitter band than a raw brassica would — appropriate, because here we are shaping, not stripping, and the brewer wants real bitterness left standing.



A single hop specimen isolated for measurement — every construct in the program is read on the same eight axes and the same TAS2R38 reference before it is allowed near a kettle.

The number on the label has to be true every batch

A brewer prints IBU on the can. A flavor we design for a brewer therefore has to arrive *the same every harvest,* or it is worse than useless — it is a broken promise, batch by batch. This is where the discipline of the program earns its place. Our constructs are designed to deliver their target profile at a batch-to-batch coefficient of variation of **≤7 percent,** against the **42 percent** that wild-type hop chemistry swings across seasons, soils, and drying lots. The difference is the difference between a spec and a hope.

Work it out on the target line. An IFI-B of **1,600** held to a 7 percent CV has a one-sigma band of ±112 — roughly **1,488 to 1,712** — and a two-sigma, 95 percent band of about **1,376 to 1,824.** A brewer can design a recipe around that. The same target left to wild-type variability, at 42 percent CV, swings ±672 at one sigma — **928 to 2,272** — a range so wide the harsh fraction might nearly double from one lot to the next. No label survives that.

It compounds the way real error compounds. The perceived harsh load a drinker meets is, near enough, the delivered iso-alpha-acids times the cohumulone fraction — the IBU and our ratio, multiplied. Independent relative errors add in quadrature, so if the brewery holds IBU to about 5 percent and our construct holds cohumulone fraction to 7 percent, the harsh fraction arrives at:

> $\sigma_{rel} = \sqrt{(0.05^2 + 0.07^2)} = \sqrt{(0.0025 + 0.0049)} = \sqrt{0.0074} \approx \mathbf{8.6\%}$

Under 9 percent, batch after batch, on the quantity the drinker actually tastes. That is a bite a brewer can stand behind in print — which, on this one axis where the number is the brand, is the whole game.

What we do not touch, and what we will not promise

A flavor designed honestly comes with its edges drawn, so here are ours on the hop.

We reshape acids, not oils — and we check. The bittering route and the aroma route share a gland but not a pathway, and the 80 constructs stay inside the window where moving the humulone-to-cohumulone ratio leaves the terpene profile unchanged. We do not take that on faith; we read every construct's oils by gas chromatography and hold the myrcene, humulene, and linalool to the parent cone before it goes anywhere near a kettle. Push a gland's flux hard enough and you could in principle perturb the beta-acids or the oil yield; we stay short of where that begins, by measurement, not by hope.

The cohumulone-harshness link is our working model, not a settled law. That high cohumulone reads as harsher is well-supported across brewing sensory work and holds cleanly on our own trained panels, and the receptor kinetics give it a mechanism — the lower EC_{50} and slower k_off of iso-cohumulone. But matched-IBU panels do not resolve it identically every time, and the sensation is partly contextual. We present it as what it is: a robust, well-supported design lever, not a universal constant. The honest version of a strong result is the one that names where it frays.

IBU is a proxy, and it overreads. The 275-nanometer absorbance that defines the unit catches more than iso-alpha-acids — oxidized acids, hop polyphenols, other chromophores all absorb there — so two beers at "the same IBU" can carry genuinely different true loads of the bittering compounds. The legacy number is a brewer's faithful servant and a chemist's rough instrument at once. Our EC_{50}-anchored IFI-B is the truer measure of loudness; we give the brewer both, and we are clear about which is which.

The tongue that reads it is not one tongue. TAS2R38, the receptor we anchor the bitter axis to, is the most famously *polymorphic* taste receptor in the human genome — the taster and non-taster haplotypes, PAV and AVI, that make a compound like PROP bracing to one person and nearly blank to another. Bitterness, more than heat or sweetness, is read differently from one drinker to the next. We anchor the Index to the reference haplotype so the number means one fixed thing; we do not pretend the number means the same *experience* in every mouth. A beer we soften 38 percent against the reference will be softened for most and felt differently by some, and a flavor standard that hid that would be lying.

And we aim for the gentle end on purpose. The 20-to-65-percent band is what the axis can do; the hop program lives in its first half. A brewer does not want the bite gone. They want it kept — firm, clean, proud on the label — and freed only of its grudge. Designing to the floor would be solving the wrong problem beautifully.

The bite, kept

The hop is a small marvel, and we have spent years with our hands in its glands learning to respect it: one plant holding two chemistries on two roads, a bittering resin and a perfume, so that one can be reshaped while the other is left wholly alone. What we hand a brewer at the end of it is not less beer. It is the same IBU they chose, the same aroma they love, and a bitterness whose *shape* has changed underneath the number — a bite that arrives firm, holds, and then clears, instead of one that arrives firm, holds, and *stays.* Full and firm and gone, not full and firm and clinging. On this axis, as on heat, we design the shape of a sensation and not merely its size; and a bitterness you can taste cleanly through, all the way to a clean finish, is a bitterness that does the one thing great bitterness is for — it invites the next glass. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. We keep the bite. We lose only the grudge.

Chapter 15 · Theobromine and the Flavan-3-ols — *cacao, where bitter meets astringency.*

DOMAIN · CACAO · the methylxanthine and the flavonoid roads · a note on status

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. So on cacao we work with a scalpel's patience and keep more than we cut, and the whole craft is knowing which of the bean's two very different bitternesses to protect and which to ease.



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.

Two bitternesses in one seed

Most foods we tune carry one bitterness to argue with. Cacao carries two, ripening in the same seed, as different as two molecules can be in what they are and in what they do to the mouth. One is an alkaloid and it lifts; one is a tannin and it drags. The discipline of designing cacao is refusing to treat them as one problem, because the eater's pleasure depends on the first surviving intact while the second is gentled by exactly the right amount and no more.

The first bitterness is **theobromine** — the principal methylxanthine of cacao, a close cousin of caffeine, present at roughly two percent of the dry bean and responsible for much of chocolate's bright, lifting edge. Caffeine is present too, usually under a fifth of a percent; the pair carry the cocoa's stimulant lift and a good share of its bitterness. Theobromine is built on the purine-alkaloid route, the SAM-dependent *N*-methyltransferase steps that build caffeine elsewhere — a nucleotide stripped to its base, then methylated nitrogen by nitrogen toward the finished alkaloid. 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 one with its personality removed.

The second bitterness is the one we ease: the harsh, drying grip of the **flavan-3-ols** and their polymers, the **procyanidins.** These are cacao's polyphenols — (−)-epicatechin, (+)-catechin, and the oligomers they build into — 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, a road that runs in every green plant and terminates, in cacao, in a tangle of linked flavanol rings.

The reason the two can be tuned apart is not metaphor; it is receptor biology, the quiet hinge the whole chapter turns on. The tongue does not read *bitterness.* It reads specific molecules at specific receptors, and the alkaloid and the tannin are read, in part, by **different members of one receptor family.** Where they overlap we cannot fully separate the lift from the grip; where they diverge we can quiet one and protect the other. So before the pathways and the arithmetic, the receptors.

How the tongue reads bitter, step by step

Bitterness is the most crowded of the taste modalities. Sweet runs through a single receptor pair, umami through a close cousin of it; bitter runs through a whole subfamily of roughly twenty-five class-A G-protein-coupled receptors — the **TAS2Rs** — on the Type II cells of the taste bud, each tuned to a different, overlapping slice of chemical space. The Index anchors IFI-B to **TAS2R38,** the archetype and the best-characterized of them, whose activation we treat as the reference point for the axis. But cacao's two bitternesses recruit a broader set, and that breadth is the opportunity.

Follow one bitter molecule in. It arrives in the saliva film, diffuses to the apical membrane of a Type II cell, and settles into the orthosteric pocket of a TAS2R — a shallow, promiscuous site, which is why one receptor can read dozens of unrelated compounds. Binding tips the receptor into its active conformation, which nudges the coupled G-protein — its α-subunit here is **gustducin** — to release its nucleotide and split from its βγ partners. The freed βγ pair switches on **phospholipase C-β2 (PLCβ2),** which cleaves a membrane lipid into two messengers; one, **inositol trisphosphate (IP_3),** opens the internal calcium store and floods the cytosol with Ca^{2+}.

That calcium is the signal. It opens **TRPM5,** a calcium-activated cation channel, and the cell depolarizes; the depolarization gates **CALHM1,** a large-pore channel that releases **ATP** straight through the membrane to the gustatory afferents waiting underneath. From binding to

neurotransmitter is a few hundred milliseconds, repeated across thousands of cells, and the sum of it is the single sensation we call *bitter* — one with no inherent identity, only a volume.

This is why a number is not a vanity on this axis: the tongue cannot report *which* bitter is ringing or *how much* is present, only how hard the channel is driven — and two very different molecules at very different concentrations can drive it equally. To design cacao we go underneath the sensation to the receptor occupancy that causes it, which is what the Index measures.

Why the two bitternesses come apart

Here is the hinge. The TAS2R family is not a single doorbell; it is a switchboard, and theobromine and the procyanidins do not press all the same buttons.

Methylxanthines — theobromine and its heavier sibling caffeine — are read by a cluster of broadly-tuned bitter receptors, the set that answers to many alkaloids. The catechins and procyanidins are read by a partly *different* cluster, the one that answers to plant flavanols and green-tea polyphenols. The two sets overlap — at least one broadly-tuned receptor reads both alkaloid and flavanol — but each has receptors the other does not strongly touch.

That overlap is not a nuisance to be wished away; it is the physical reason the easing has a floor. If one receptor reads *both* the lift and the grip, quieting the flavanol through that shared channel unavoidably dims a sliver of the alkaloid riding the same wire — they are the same conformational event, and you cannot drive one to zero and leave the other. A perfect split is not available at any price. What *is* available is the far larger fraction of the procyanidin signal that runs through the flavanol-private receptors the alkaloid barely engages. Quiet those, and you move the grip a long way while the lift barely registers the loss.

This sets the target honestly. We do not promise the impossible clean split; we design to move the separable fraction — which, in our model, lands the roughest procyanidin harshness down by roughly forty percent while the methylxanthine lift stays essentially where it was. The forty percent is not a slogan. It is, in part, the size of the flavanol-private signal the receptor map actually lets us reach.

One more biological fact belongs here, because it governs who tastes what. TAS2R38 is famously polymorphic: common haplotypes at three residue positions sort people into stronger and weaker bitter perceivers, so the same bar lands as a wall on one taster and a rounded depth on another, by genotype alone. A designer who forgets this builds for an average mouth no single person owns. We tune toward the sensitive end and let the less-sensitive end read it as gentleness — a bar kind to the super-taster is merely mellow to everyone else, and the reverse failure, one built for the non-taster that flays the sensitive, is the one we refuse to ship.



The bean in cross-section — the methylxanthine lift and the flavan-3-ol grip, two chemistries one knife-cut apart.

The methylxanthine road — the lift we protect

Theobromine is built on the purine-alkaloid route, a short pathway that borrows a spent nucleotide out of the cell's housekeeping and decorates it. The methyl groups come from one donor, **S-adenosyl-L-methionine (SAM),** the cell's universal methyl currency, each transfer leaving behind S-adenosylhomocysteine. The enzymes are **SABATH-family *N*-methyltransferases,** the same clan that builds caffeine in coffee and tea.

STEP	SUBSTRATE	ENZYME (SAM-DEPENDENT)	PRODUCT
1	xanthosine	**XMT** (7-methylxanthosine synthase)	7-methylxanthosine
2	7-methylxanthosine	*N*-methylnucleosidase (hydrolysis)	7-methylxanthine
3	7-methylxanthine	**MXMT** (theobromine synthase)	**theobromine**
4	theobromine	caffeine synthase	caffeine *(minor in cacao)*

The crucial fact sits at step four. In the coffee bean the pathway runs to completion — the third methyl goes on and theobromine is consumed into caffeine, which is why coffee is a caffeine crop. In cacao the final 3-*N*-methylation is weakly expressed, so the road effectively *stops at theobromine,* and the alkaloid accumulates to its characteristic couple of percent while caffeine trails far behind. Cacao is, in enzymatic terms, a theobromine crop that never became a caffeine crop — and that unfinished quality is the gentle, long, bright lift a chocolate carries instead of coffee's sharper jolt.

That the lift is *felt* and not merely tasted has a pharmacological cause worth naming plainly, and without a single health claim attached: theobromine is a mild antagonist at the adenosine receptors caffeine also blocks, but weaker and slower. The brightness in a fine dark bar is partly your tongue reporting the alkaloid's bitterness and partly your body reporting its gentle antagonism of the molecules that would otherwise quiet you. To strip theobromine out for a smoother number would remove the thing that makes chocolate *lift.* So we protect it as the first design constraint — a bitterness we deliberately *leave ringing* inside the final score, a floor under IFI-B we have no wish to cross.

What does protecting it cost, in concentration terms? A worked order of magnitude — pharmacology at representative literature scale, targets from the Index. Theobromine's molar mass is 180.16 g/mol; at roughly two percent of the dry bean, an 80% bar carries on the order of 1.6–2.0 g per 100 g. A single bite, melting through the saliva film, presents the receptor with an effective concentration in the tens of millimolar — call it 15–25 mM, which is 1.8–3.0 g/L divided by 180.16. Against a bitter recognition threshold near 1–2 mM, that is roughly ten times over. The lift is loud at the receptor, and it is *meant* to be.

The flavan-3-ol road — the grip we ease

The second road is longer and it is where the work happens. The flavan-3-ols branch off the phenylpropanoid trunk that every vascular plant runs — phenylalanine deaminated to cinnamic acid, hydroxylated, activated to its coenzyme-A thioester — and then turn down the dedicated flavonoid arm:

STEP	ENZYME	PRODUCT
1	**CHS** (chalcone synthase)	`naringenin chalcone`
2	**CHI** (chalcone isomerase)	`naringenin`
3	**F3H** (flavanone 3-hydroxylase)	`dihydroflavonol`
4	**DFR** (dihydroflavonol reductase)	`leucoanthocyanidin`
5a	**LAR** (leucoanthocyanidin reductase)	**`(+)-catechin`**
5b	**ANS → ANR** (anthocyanidin synthase, then reductase)	**`(−)-epicatechin`** → `procyanidins`

The lever is the **LAR / ANR branch point** at step five — the same pair of enzymes that governs flavan-3-ols in tea and the grape, which is precisely why the method transfers and why our model demonstrates the lever on published hosts before it ever reaches a bean. LAR makes the (+)-catechin arm; the ANS→ANR arm makes (−)-epicatechin, the monomer that polymerizes. Two epicatechin units linked (most often 4→8, sometimes 4→6) make a dimer; add units and you climb to trimers, tetramers, and on up into oligomers with a **mean degree of polymerization (mDP)** that in cacao typically sits around five to nine units.

Degree of polymerization decides how the flavanol feels, because the mouth reads a short flavanol and a long one completely differently. The monomers and smallest oligomers fit the bitter receptors' pockets and are read as *bitter,* through the flavanol-private TAS2Rs. The longer polymers barely touch the taste receptors; they do their damage by a different mechanism entirely. So the flavan-3-ol pool is not one flavor. It is a *distribution* across chain length and galloylation — whether a gallic-acid group is esterified onto the flavanol, which sharpens both bitterness and astringency — and the sensation a bar delivers is the shape of that distribution.

This is why the easing can be surgical rather than blunt. Tuning the flux through ANR — how much of the leucoanthocyanidin pool is pulled down the epicatechin-and-polymerize arm, and how far polymerization runs — *reshapes the distribution* rather than draining the pool. Trim the bitter-skewed tail of small, galloylated species the flavanol-private receptors read most loudly, and lFI-B drops; preserve the mid-length procyanidin network, and the structure stays. The

forty-percent easing our model targets is a redistribution, not a subtraction. And here the two axes braid, because those same molecules are the heart of IFI-A — astringency — so every move on cacao's bitterness is also a move on its grip. To see why holding both is possible, we leave the taste receptors behind.

Astringency is not bitterness — the protein mechanism

Astringency is the most misunderstood of the eight axes, because it is not a taste at all. There is no astringency receptor on the taste bud. IFI-A is anchored not to a GPCR but to the **salivary proline-rich proteins (PRPs),** and the sensation is mechanical.

Saliva is a lubricant. A thin film of PRPs and related proteins coats the oral surfaces and lets them slide past one another frictionlessly — the default, unnoticed state of a mouth. Procyanidins, particularly the longer, galloylated ones, are built to bind those proteins: they hydrogen-bond to the exposed prolines and stack hydrophobically against the backbone, cross-linking several PRP molecules at once until the complexes fall out of solution. The lubricating film precipitates. For a few seconds the oral surfaces have nothing to slide on, and the trigeminal mechanoreceptors report the friction directly as **drying, roughness, puckering** — the pull of a strong tea, the grip of a young red wine, the scrape of a very high-percentage bar.

Because astringency is a *precipitation* event and bitterness a *ligand-binding* event, they respond to different features of the same pool. Bitterness tracks the small, pocket-fitting species; astringency tracks the large, protein-cross-linking ones, its avidity for PRPs rising steeply with polymerization and galloylation — a pentamer precipitates protein far more effectively than a monomer, which may not at all. This is the molecular permission slip for the whole cacao strategy: *we can move IFI-B and IFI-A in different directions* because they are driven by different tails of the distribution. Trim the small bitter-skewed species and IFI-B falls while the astringent backbone is untouched; preserve the mid-length network and IFI-A holds. The braid is real, but its two strands are separable.

That separability is not free or total — push the trim too far up the chain and astringency follows bitterness down, because the fractions overlap in the middle. But across the range that matters for a great bar the lever exists, and it is why the Index can ease one axis by forty percent while holding its neighbor nearly fixed.

Putting a number on it — the Index, worked

All of this becomes designable only as arithmetic. The Index is one equation:

$$\mathrm{IFI}(c) = 1000 \times \mathrm{EC}_{50}(\text{reference}) \div \mathrm{EC}_{50}(\text{candidate})$$

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response — the half-activation point. The formula inverts it against a fixed reference and scales by a thousand, so the *score rises as the half-activation concentration falls.* A loud stimulus reaches half-activation at a low concentration and scores high; a quiet one needs more and scores lower. A score of exactly 1,000 means the candidate half-activates at the same concentration as the reference. The per-axis scorecard ceiling is 5,000.

The inversion is the part a designer must hold in hand, because it makes the intuition slippery: score and concentration run *opposite* ways. To make a bitterness read 40% quieter we are not lowering a concentration by 40%; we are *raising* the half-activation concentration, and by more than 40%. Let the native harshness score be some value B_0, and the eased score be $B_1 = 0.60\,B_0$ — a 40% reduction. Since score is inversely proportional to EC_{50},

$$\frac{\mathrm{EC}_{50}(\text{eased})}{\mathrm{EC}_{50}(\text{native})} = \frac{B_0}{B_1} = \frac{1}{0.60} = 1.667.$$

A 40% drop in loudness is a **66.7% rise** in the concentration needed to half-activate the receptor. That gap between "40% quieter" and "67% higher threshold" is where flavor intuition most often goes wrong, and the Index exists so it cannot go wrong on us.

Now the representative bar the Index specifies. The canonical 80% readout scores the eased harshness at **B = 1,700.** From the formula, its half-activation concentration relative to the axis reference is:

$$\frac{\mathrm{EC}_{50}(\text{bar})}{\mathrm{EC}_{50}(\text{ref})} = \frac{1000}{1700} = 0.588,$$

so the bar's effective bitter stimulus is 1/0.588 ≈ **1.70× as loud** as the reference. And because the Index tells us this 1,700 is the native harshness *eased about forty percent,* we can recover where it started:

$$B_0 = \frac{B_1}{0.60} = \frac{1700}{0.60} = 2{,}833.$$

The native, un-tuned harshness scored ≈ 2,833 — half-activation at 1000/2833 = 0.353× the reference, i.e. 2.83× as loud. The easing raised the threshold from 0.353 to 0.588 of the reference, a factor of 1.667, exactly as the general relation demanded. Every number in that chain is canonical or derived from a canonical number by the one formula; none is invented.

Reading the whole profile this way turns the scorecard into a table of concentrations. Each axis anchors to its own reference compound by convention; U, S, K, A, and B scores each measure loudness against a different defined reference point and are not on a single common scale — the profile reads as a shape, not an absolute ranking across axes.

AXIS	IFI SCORE	EC_{50}(CAND)/EC_{50}(REF) = 1000/IFI	LOUDNESS VS. REFERENCE
U (umami)	300	3.333	0.30×
S (sweetness)	500	2.000	0.50×
K (kokumi)	1,200	0.833	1.20×
A (astringency)	1,500	0.667	1.50×
B (bitter, eased)	1,700	0.588	1.70×
B (bitter, native)	2,833	0.353	2.83×

One honest refinement keeps this from being too tidy. Receptor activation is not linear in concentration; it follows a saturating dose–response, a Hill function,

$$\theta = \frac{[L]^n}{\mathrm{EC}_{50}^n + [L]^n},$$

where θ is fractional activation and n is the Hill coefficient — for bitter GPCRs, representatively somewhere between 1 and 2. At [L] = EC_{50}, θ = 0.5 by construction, which is why the Index anchors there: the half-activation point is the one concentration where the curve's position is

read cleanly, independent of n. Away from it, the curve bends. Two consequences matter for cacao. First, near saturation two bitter compounds combine *sub-additively* — the receptor is already most of the way open, so adding more does little, which is why you cannot simply sum the theobromine and procyanidin scores. Second, the *steepness* near threshold means a modest shift in the flavanol distribution can produce a disproportionate change in perceived bitterness — the lever has gain. The Index reports the clean half-activation anchor; the Hill curve reminds us the mouth lives on the slopes around it.

The representative bar, scored

A word on status, in the plain voice this library keeps for the truth about its own work: **cacao is not yet among the forty-two host organisms with published construct catalogs.** Our standing cacao work is method-transfer and roadmap — the flavonoid machinery modeled and published on our polyphenol hosts and pointed toward the bean, with the full host program on the companion shelf as *Cacao & Chocolate* (Volume XVI). We say 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. What follows is a *representative* readout — an 80-percent bar built to the profile the method targets, scored on the Index to show the intended balance.

First the same mass *before* the ANR tuning, with the harshness at its native 2,833 and everything else where the finished formulation sits — note especially the astringency already at 1,500, because that is the number we intend to protect, not reduce:

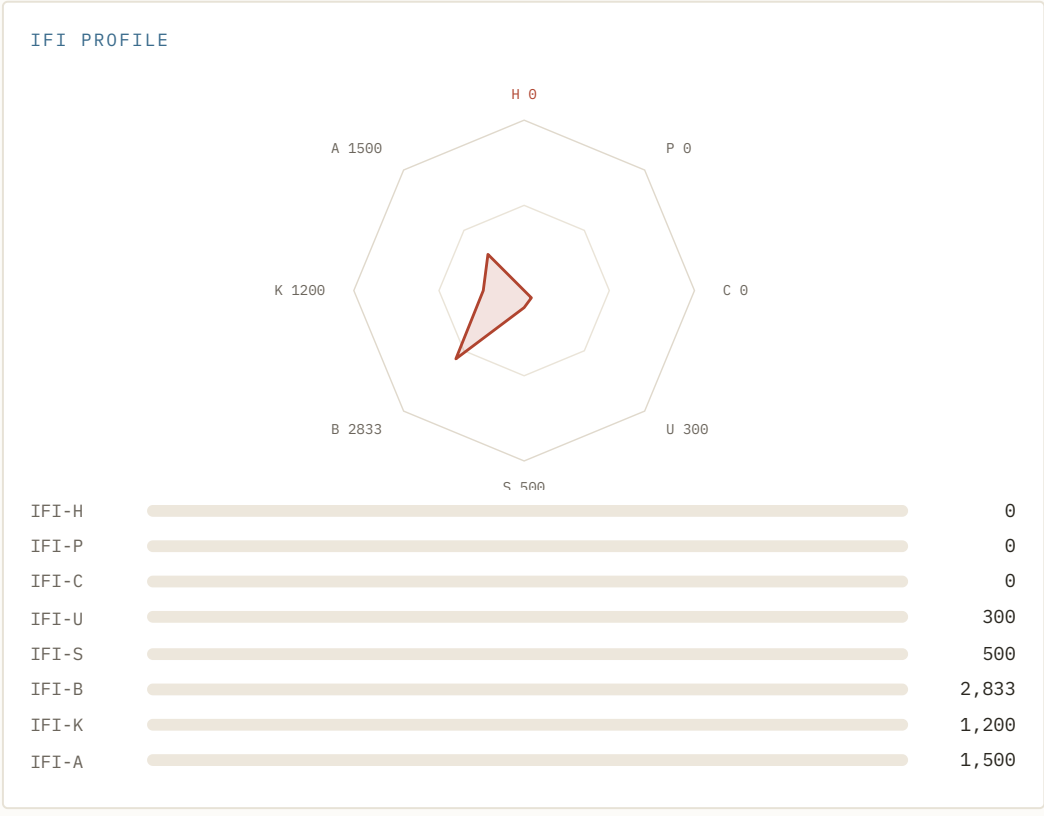

Now the tuned target — the harshness eased about forty percent to 1,700, and **A held at exactly 1,500,** unmoved:

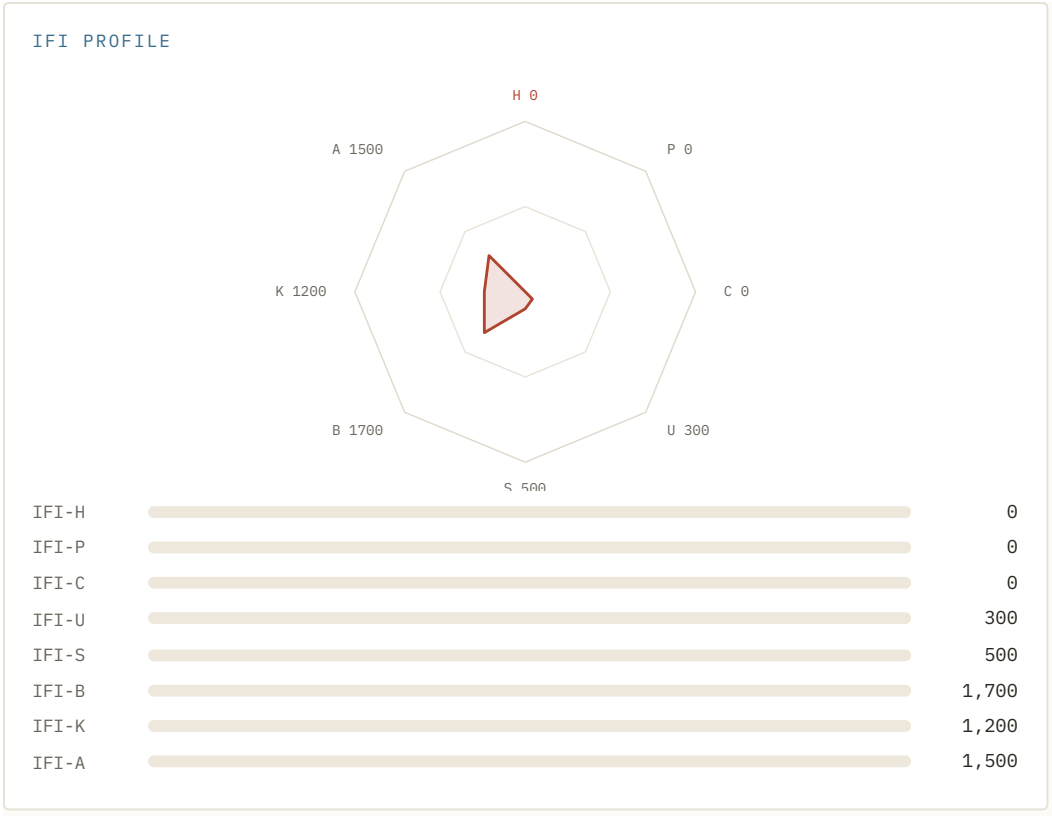

Read the pair as a map of restraint. The only number that moves is B, from 2,833 to 1,700 — the small bitter-skewed flavanol tail trimmed, the flavanol-private receptors quieted. The **A of 1,500 does not move,** because the mid-length procyanidin backbone that precipitates the salivary film is preserved deliberately; the grip that gives a bar its structure is the thing we refuse to touch. The **K of 1,200** is the deep body of good cacao; the **S of 500** is sugar enough to round the bar and never enough to mask it; the **U of 300** is the faint savory undertone of roasted cocoa. A chocolatier can now take the cocoa percentage higher without the bar turning against the eater at the finish. The scrape goes; the structure stays.

Our model predicts that tuning the bean's *distribution* rather than diluting the bar is what makes it hold across the percentage ladder. The same forty-percent redistribution, carried up as the cocoa fraction climbs, keeps the harshness in hand where an untuned bean would run from the eater:

BAR	NATIVE HARSHNESS (B_0)	EASED HARSHNESS ($B_1 = 0.60\ B_0$)	ASTRINGENCY (A, HELD)
70%	2,550	1,530	1,400
80%	2,833	1,700	1,500
90%	3,150	1,890	1,650
100% (nib)	3,500	2,100	1,800

The 80% row is the canonical readout above; the others are the same arithmetic — a fixed forty-percent easing — carried up the ladder as representative figures, to show the *shape* of the lever, not a validated score for each bar. The eased harshness never collapses to nothing; it is pulled to a structured, bright bitterness and left there. A 100% nib, untuned, scores a punishing 3,500; eased, our model predicts 2,100 — still emphatically dark, still unmistakably a nib, but no longer a dare.

Method transfer — grape, tea, and the road to the bean

The reason we can speak about a bean we have not yet modeled in full is that the enzymes we would tune in cacao are the enzymes we *have* modeled, at scale, in plants that run the identical chemistry. The LAR/ANR branch point is not cacao's invention; it is shared across the

flavan-3-ol plants, and an anthocyanidin reductase from one is recognizably the same machine as another's. Our model demonstrates the lever on the published polyphenol hosts; that same lever is what cacao would engage.

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY	ROLE IN THE CACAO METHOD
Grape	1,610	B / A / M	Hard	the flavan-3-ol and ANR proving ground — tannin and astringency at scale
Tea	280	U / B / A	Hard	catechin and galloylated-flavanol tuning; the bitter-astringent braid
Coffee	40	B	Hard	the methylxanthine side — SABATH *N*-methyltransferase bitterness

The grape's **1,610 construct designs** are the largest body of computational tannin engineering in the library — the direct proving ground for the ANR-flux redistribution cacao would inherit. Tea's **280** are the sharper teacher, because tea's whole character is the catechin-and-galloylation balance that decides bitter-versus-astringent, the exact distribution problem cacao poses. Coffee's **40** sit on the other road entirely — the SABATH *N*-methyltransferases, the enzymology of *protecting* an alkaloid bitterness rather than easing a flavanol one. Both of cacao's bitternesses already have a computational home in the library; what cacao adds is the demand to run both roads in one seed at once, which is what Volume XVI's host program exists to finish.



The grape — 1,610 construct designs of flavan-3-ol and procyanidin work; the ANR branch modeled here is the branch cacao would inherit.

Reproducibility — the real product

A design target is only as good as the tightness with which you can hit it, batch after batch, and agricultural cacao is notorious for *not* hitting anything twice. Origin, season, pod ripeness, and above all the heat of the post-harvest ferment swing a bean's procyanidin content enormously. Two bars of the same percentage, from the same estate in two different years, can sit far apart on IFI-B through nothing but the weather and the fermentation heap. For a chocolatier formulating to a profile, that native variability is the real adversary — often worse than the bitterness itself.

The platform's answer is a tuned *starting distribution* held to the library's reproducibility standard: a batch-to-batch coefficient of variation of **≤7%.** The coefficient of variation is the standard deviation

divided by the mean, CV = σ/μ, and it is the right measure here because it scales with the target. At the eased harshness of B = 1,700, a 7% CV is a standard deviation of σ = 0.07 × 1,700 = 119. Ninety-five percent of batches — ±2σ — therefore land within:

$$1700 \pm 2(119) = 1700 \pm 238 = [1{,}462,\ 1{,}938].$$

A formulator can build a bar against that and trust it. Set beside a representative wild-cacao procyanidin swing — take a CV of 35% as an illustrative figure for an unmanaged bean — the same ±2σ window would span σ = 0.35 × 1,700 = 595, a range of 1,700 ± 1,190 = [510, 2,890]. That is not a bar you can formulate; it is a different bar in every shipment. The tuned distribution does not merely make the bitterness kinder — it makes it *the same,* and reproducibility, not any single heroic score, is what turns a flavor into an ingredient a maker can design around.

One caution keeps this honest: the genetic lever sets the bean's starting distribution, but fermentation and roast act on top of it. A ferment degrades procyanidins substantially; roast drives epimerization and further polymerization. Tune the genotype and then over-ferment, and you can erase the design. The ≤7% is our control of the input; the maker's process is a second lever stacked on ours, and the best bars treat the two as one system from pod to bar.

The honest limits

Four edges deserve plain statement, because a method that hides its edges cannot be trusted at the ones it claims.

The first is the **overlap floor.** Because at least one broadly-tuned receptor reads both the methylxanthine and the flavanol, a sliver of theobromine's brightness runs through the same channel we quiet to ease the grip; we cannot drive the shared contribution to zero without dimming the lift. The forty-percent target is the separable fraction, not a promise of a clean split the receptor map forbids.

The second is **human polymorphism.** The bitter receptors, TAS2R38 above all, vary across people, so a single eased target reads as gentleness to one taster and flatness to another. We tune toward the sensitive mouth and accept that the less-sensitive one reads the result as mellow, because that failure is kind and its opposite is not.

The third is the **candy cliff.** Ease too far and the bar loses the backbone that made it chocolate. Push the eased harshness much below roughly 1,200 and you cross from *structured* into *flat,* and no amount of sugar brings the character back — it only buries the absence. The discipline is to stop where the structure begins, which is why B lands at 1,700 and A is not moved at all.

The fourth is **nourishment,** the one the family cares about most, so we will not dress it up. The procyanidin flavanols are not only the grip; they are the fraction of cacao most associated with what the bean gives the body that eats it. To ease bitterness by draining the flavanol pool wholesale would trade nourishment for smoothness, and that is a trade we will not make. The redistribution is built around exactly this: we trim the harshest small-species tail the tongue reads as scrape,

and we *protect the flavanol network* that holds both the structure and the goodness. Keeping the nourishment and keeping the character turn out to be the same act, because they live in the same molecules.

What we keep



A single bean, isolated — two bitternesses, an alkaloid's lift and a tannin's grip in one seed; the whole craft is knowing which to protect.

On cacao, more than anywhere on the axis, Zero-Bitter is not subtraction for its own sake. It is the fine judgment of how much of a thing's defining edge to protect. We leave the methylxanthine lift ringing because it is the chocolate's own brightness and nobody wants it gone. We ease the roughest procyanidin harshness by about forty percent and no more, because the small bitter-skewed tail is scrape and the mid-length backbone is structure. We hold the astringency fixed because the grip is a backbone, not a flaw. And we guard the flavanol network because it carries, in the same molecules, both the character and the nourishment — and keeping food worth eating means keeping both.

The grip stays; only the scrape goes. That is the whole of it, and under it runs the arithmetic that makes it true: one formula, two roads, a distribution reshaped rather than drained, and a bitterness a child could trust, left as bright and as honest as a fine dark bar was always meant to be.

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

Chapter 16 · Tea, Grape, Citrus — *the polyphenol family, and knowing when not to cut.*

DOMAIN · TEA 280 · GRAPE 1,610 · ORANGE 120 · the shared flavonoid road

The last of the shelf's bitternesses live in a leaf, a berry, and a fruit, and they are family — all three built, like cacao's grip, on the flavonoid arm of the phenylpropanoid branch, all three braided with astringency, all three teaching the same hard lesson about this axis: that holding it well means knowing, plant by plant, when *not* to turn it at all. We came to tea and grape and citrus last, and almost by relief, because by the time a program has learned to design a bitterness it has usually forgotten how to leave one alone. These three plants would not let us forget. Each of them carries, inside the same chemical sentence, a harshness we want quieted and a structure we would be fools to touch — and the only way to tell the two apart is to go down to the receptor and read the molecule by its size.



Tea in cross-section — the catechins that grip a strong cup and the theanine that rounds it ride the same leaf, and the craft is to quiet the first thread without loosening the second.

What unites them is a single biosynthetic road. The upstream phenylpropanoid branch — phenylalanine opened by PAL, hydroxylated, activated, and routed into the flavonoid pathway through chalcone synthase — runs in every vascular plant, and it is this common road that pours out the flavan-3-ols, the flavanones, and the condensed tannins that give a cup of green tea, a glass of young red, and the white pith of an orange their shared, insistent grip. The compounds differ at the last few enzymes. The receptors they shout into are nearly the same. And the discipline that the whole family

demands is the one our table keeps returning to: that we are not here to strip a plant of what it made, only to quiet how loudly one part of it shouts. The polyphenols these plants build are the plant's own gift, as much to the body that eats them as to the leaf that holds them, and almost nothing we do here destroys them. We rebalance them, we re-route a sugar onto a bitter glycoside, we move a gallate group — and the nourishment the plant spent its summer assembling stays, while the harshness steps down. Knowing which thread is which is the entire chapter.

The receptor at the end of the road

Bitterness is the body's oldest alarm, and the tongue reads it with a whole committee rather than a single sentry. The bitter receptors are the TAS2R family — twenty-five of them in the human genome, each a seven-transmembrane G-protein-coupled receptor wired through gustducin to the same downstream shout of *do not swallow this.* The Index anchors IFI-B to **TAS2R38**, the family's best-studied member and its reference receptor, so that every bitter thing in the library — a Brassica, a hop, a coffee, and the three plants of this chapter — can be scored against one fixed yardstick. We should be honest, because the numbers are a form of respect, that the polyphenol bitters here do not in fact speak mainly through TAS2R38 itself. That receptor is tuned to the thiourea and glucosinolate bitters; the galloylated catechins of tea and the flavanone glycosides of citrus are read most loudly by TAS2R39 and TAS2R14, with TAS2R4 and others filling in. The Index's job is not to pretend otherwise. It is to put a catechin, a naringin, and a limonin on one comparable scale against a single stable reference, so that a chocolatier and a tea master and a juice formulator can read the same page. The axis is one number; the biology underneath it is a chord.

The deeper fact — the one that decides everything in this family — is that bitterness and astringency are two different events, and a phenolic molecule is sorted between them almost entirely by its *size*. A bitter receptor is a pocket. To register as bitter, a molecule has to fit inside the orthosteric cleft of a TAS2R and hold there long enough to flip the receptor on. Small phenolics do this easily: a monomeric flavan-3-ol, or better a galloylated one, slips in and grips. But astringency is not a receptor event at all. It is mechanical. The large condensed tannins — polymers of flavan-3-ol units — are far too big for any TAS2R pocket, so instead of binding a receptor they bind *protein*, cross-linking the proline-rich proteins and mucins of the saliva into aggregates that fall out of solution. The tongue reads the loss of lubrication as a drying, puckering roughness: IFI-A, astringency, carried by the salivary proline-rich proteins rather than by a taste receptor. One molecule, grown a little longer, crosses from one axis to the other.

PHENOLIC CLASS	REPRESENTATIVE	APPROX. MASS / SIZE	FITS A TAS2R POCKET?	CROSS-LINKS SALIVARY PRPS?	DOMINANT PERCEPT
Monomeric flavan-3-ol	epicatechin (EC)	~290 Da	Yes	Weakly	Bitter
Galloylated monomer	EGCG · ECG	~458 · ~442 Da	Yes, strongly	Moderately	Bitter **and** astringent
Oligomer (2–3 units)	procyanidin B2 · C1	~578 · ~867 Da	Marginally	Moderately	Bitter shading to astringent
Condensed polymer (>4 units)	seed / skin tannin	>1,200 Da	No — too large	Strongly	Astringent

Read down that table and the whole family organizes itself. Tea leads with the galloylated monomers, which is why a strong cup is both bitter and drying and why tea alone of the three carries all three axes — U, B, and A. Grape's signature lives at the bottom row, in the long polymers, which is why a glass of wine is overwhelmingly astringent and barely bitter. Citrus sits at the very top, with small, sharp, soluble molecules that are nearly all bitterness and almost no astringency. The lever we reach for differs in each case not because the plants are unrelated but because they occupy different rungs of the *same ladder*. Design on this axis is, before anything else, the art of reading where on that ladder a given harshness lives — and refusing to cut a rung that is carrying the plant's structure rather than its harshness.

Reading loudness as a number

The Index measures loudness, not quantity. Its single formula is

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

where EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. A low EC_{50} means a molecule that shouts at a whisper of concentration; invert it against a fixed reference, scale by a thousand, and that loudness becomes a score. We anchor the bitter axis so that the reference stimulus, read at its own half-activating concentration, lands at exactly 1,000. For the worked examples in this chapter we fix that illustrative reference at $\mathrm{EC}_{50}(\mathrm{ref})$ = 100 µM; every bitter compound is then just a ratio away.



A bench where a candidate's half-activating concentration is read off a dose–response curve — the number that becomes an IFI score.

Run the family's bitters through that single division and a ladder falls out:

COMPOUND	CLASS	EC_{50} (MM, INDEX REFERENCE ASSAY)	IFI-B = 1,000 × 100 ÷ EC_{50}	CHIEF IN
Reference anchor	—	100	1,000	–
Limonin	limonoid triterpenoid	2.1	47,600 → **capped at 5,000**	orange (delayed)
Naringin	flavanone neohesperidoside	48	2,080	orange / grapefruit pith
EGCG	galloylated flavan-3-ol	55	1,820	tea
ECG	galloylated flavan-3-ol	95	1,050	tea
EGC	flavan-3-ol	480	210	tea
Caffeine	purine alkaloid	600	170	tea
EC	flavan-3-ol	1,250	80	tea

The arithmetic is honest and it is also instructive. Limonin is so potent — its bitterness detectable near 1 part per million — that its intrinsic score runs clean off the scale. Convert the threshold to see it: limonin's molar mass is 470.5 g/mol, so 1 ppm, which is 1 mg/L, is (0.001 g/L) ÷ (470.5 g/mol) = **2.13 µM.** A molecule half-activating its receptor at ~2 µM scores nearly 48,000 against a 100-µM reference, which is why the Index applies a per-axis normalization ceiling of **5,000:** no single axis is permitted to run away with the scorecard, and limonin, left uncapped, would do exactly that. The cap is not a fudge. It is the admission that beyond a certain loudness the tongue itself stops counting — the receptor saturates, and more agonist buys no more alarm.

That saturation is the subtlety that keeps a food's score from being a naïve sum. A single compound's response follows a Hill function,

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

with a Hill coefficient near n = 1.3 for the cooperative TAS2R bitter receptors. At [A] = EC_{50} the response is exactly one half, and the curve is at its steepest. This is the single most useful fact in bitter engineering, because it tells us *where* to spend an edit. In a strong cup, EGCG sits near its own EC_{50} — that is precisely why it dominates the bitterness. Cut its effective concentration by 45%, moving [A] from 1.0 × EC_{50} to 0.55 × EC_{50}, and the receptor response falls from 0.500 to $0.55^{1.3} \div (1 + 0.55^{1.3})$ = 0.460 ÷ 1.460 = **0.315** — a 37% drop in activation from a 45% drop in concentration. Work where the dose–response is steepest and a modest cut in one molecule buys a disproportionate quieting. Work out on the saturated shoulder and the same cut buys almost nothing. Every edit in this chapter is placed on the steep part of a curve.

A food, of course, is a chord of such curves, and we read its axis value by concentration addition — summing each contributor's loudness, already compressed onto the anchored 0–5,000 scale, into the one number the scorecard reports. Near the operating range

where we work, well below full saturation, those contributions add almost linearly, which lets us decompose a cup's bitterness into named parts and then do arithmetic on them.

Finally, a number without an error bar is a boast, not a measurement. Because IFI-B is a *ratio* of two half-activating concentrations, its uncertainty propagates as the quadrature sum of the two relative errors: $CV(IFI) = \sqrt{(CV(ref)^2 + CV(candidate)^2)}$. With a well-run dose–response assay holding each EC_{50} to about 5% coefficient of variation, the score carries $\sqrt{(0.05^2 + 0.05^2)} = \sqrt{0.005} =$ **7.07%** — which is why we design to the same **≤7% CV** specification that governs all our construct work, the same reproducibility target that positions our pepper cultivar designs against a wild-type pepper's 42% batch-to-batch drift. A tea designed for IFI-B 1,400 is built to hold that score within ±98 harvest after harvest. That consistency is the whole point of a design specification: a formulator can plan around 1,400 because the design targets that score for each successive harvest.

Tea — easing the galloyl grip without loosening the backbone

Tea sits exactly on the seam between bitterness and astringency, which is why it carries three axes on the Index — U, B, and A — and why it is one of the more intricate organisms in the library, with **280 constructs.** 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 tea leaf: chiefly EGCG (epigallocatechin gallate), along with EGC, ECG, and EC. The leaf assembles them on the flavonoid road through the two reductase routes that make all flavan-3-ols — leucoanthocyanidin reductase building the 2,3-*trans* catechins and anthocyanidin reductase building the 2,3-*cis* epicatechins — and then does the one thing that turns a merely bitter leaf into a fierce one: it *galloylates.* A serine-carboxypeptidase-like acyltransferase transfers a galloyl group, carried on the activated donor β-glucogallin, onto EGC and EC to make EGCG and ECG. That single ester group is much of what makes the galloylated catechins the most aggressively bitter-and-astringent molecules in the cup; strip it and EGCG collapses back to the far gentler EGC.

So the lever is galloylation itself. By dialing down the galloyl-transfer step and the β-glucogallin supply that feeds it, we shift the catechin profile away from EGCG and ECG and toward their ungalloylated parents — easing the roughest, EGCG-led harshness while conserving the flavan-3-ol skeleton almost entirely. The nourishment the leaf built is not thrown away; a gallate is moved, a parent catechin is restored, and the molecular mass of healthful phenolic in the cup barely changes. What changes is how loudly one corner of it shouts. The arithmetic, decomposed into the cup's own IFI-B points and recombined after the edit:

BITTER CONTRIBUTOR	BEFORE (IFI-B POINTS)	EDIT	AFTER (IFI-B POINTS)
EGCG	1,100	galloyl supply ↓ 45%	605
ECG	450	galloyl supply ↓ 45%	248
EGC	200	fed by de-galloylation	280
EC	80	fed by de-galloylation	100
Caffeine	170	untouched	170
Cup total · IFI-B	**2,000**		**≈ 1,400**

Read it carefully, because the two percentages are not the same number and the difference is the craft. The galloylated pool — EGCG plus ECG — falls from 1,550 to 853 points, a **45% cut** in the galloylated catechins. But the cup's total bitterness falls only from 2,000 to 1,400, a **30% reduction** — because the de-galloylation feeds the gentler EGC and EC, which rise, and because caffeine is left wholly alone. Forty-five percent less of the harsh molecule; thirty percent less harshness. Both sit comfortably inside the 20–65% attenuation envelope this axis works within, and the gap between them is exactly the honesty the number demands. A representative readout, a green tea engineered to forgive a long steep:

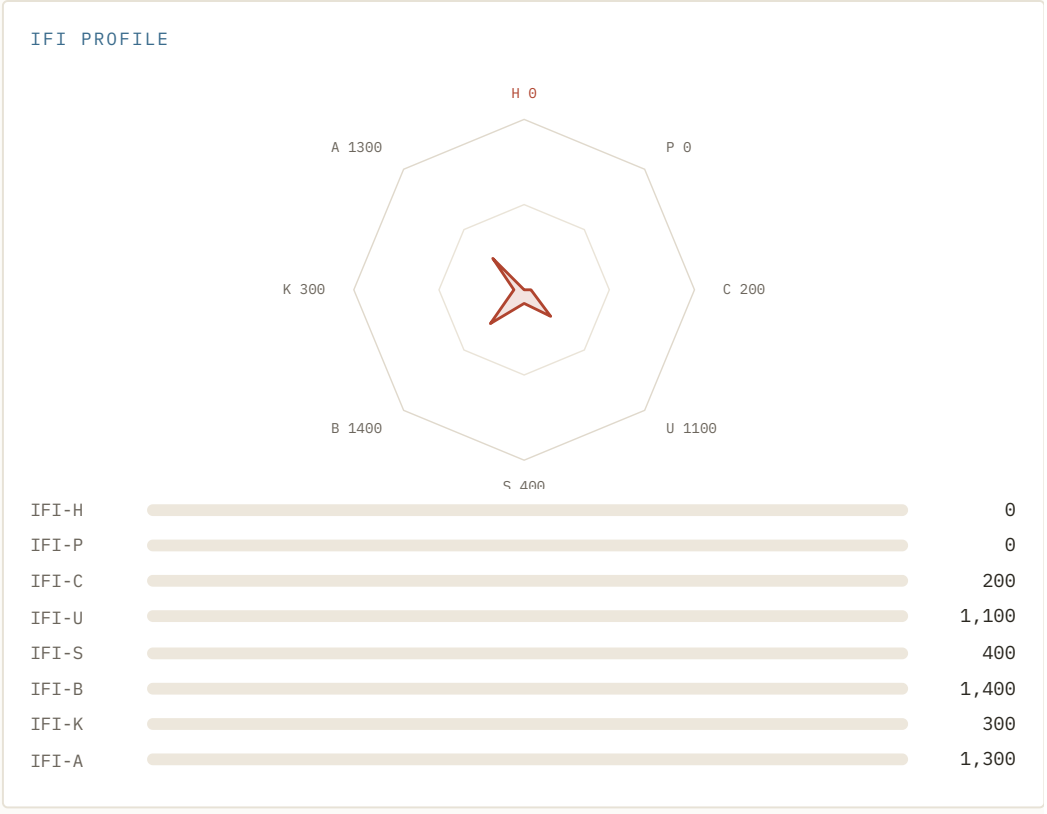

The **B of 1,400** is the harshest catechins down about thirty percent. The **A held at 1,300** is the drying backbone kept for length — and holding it was the whole discipline, because astringency is carried by the longer, oligomeric phenolics that cross-link saliva, not by the small galloylated monomers we trimmed, so the two can be separated if the edit is precise enough. A tea with no grip at all is slack and forgettable; we quiet the bitter thread and leave the astringent one standing. And the **U of 1,100** is the quiet reward: **theanine** — L-theanine, the amino acid nearly unique to the tea plant — which carries the round, savory, broth-like sweetness a fine green tea is

prized for, and which we leave entirely intact, so that lowering the harsh catechins actually lets the theanine *step forward* into the space they vacated. A tea you can steep without fear is a tea that gives more of itself. The astringent half of the story in full, the tannin dimension, the shelf hands on to *Astringency* (Volume XIII) and to *Tea* (Volume XVII), its own monograph. Here we have turned only the bitter thread of the braid, and we have turned it where the dose–response is steep.

Grape — when the right turn is almost none

Grape teaches the lesson's sharper edge: sometimes the right amount to turn an axis is almost none. It is one of the most heavily worked organisms in the entire library — **1,610 constructs** — but its signature is **tannin,** not bitterness, and the whole craft on grape is restraint.



Grape in cross-section — the seeds and skin hold the proanthocyanidins; a little of their bitterness is structure, and cutting it would make a spineless wine.

Its defining compounds are the seed and skin **proanthocyanidins** — the condensed tannins, large procyanidin and prodelphinidin polymers built, again, through the LAR and ANR reductase routes and then extended, unit upon unit, into chains. Their behavior is governed almost entirely by their *mean degree of polymerization.* A dimer or trimer still fits a bitter pocket and reads partly as bitterness; push the chain past four or five units and it no longer fits any TAS2R at all, and instead goes to work cross-linking the saliva's proline-rich proteins into that structural dryness. Seed tannins run shorter and more galloylated — more bitter; skin tannins run longer and richer in prodelphinidin — more purely astringent. The grape's grip, in other words, lives overwhelmingly at the bottom rung of the size ladder, where the percept is IFI-A, not the TAS2R warning. A representative young red:

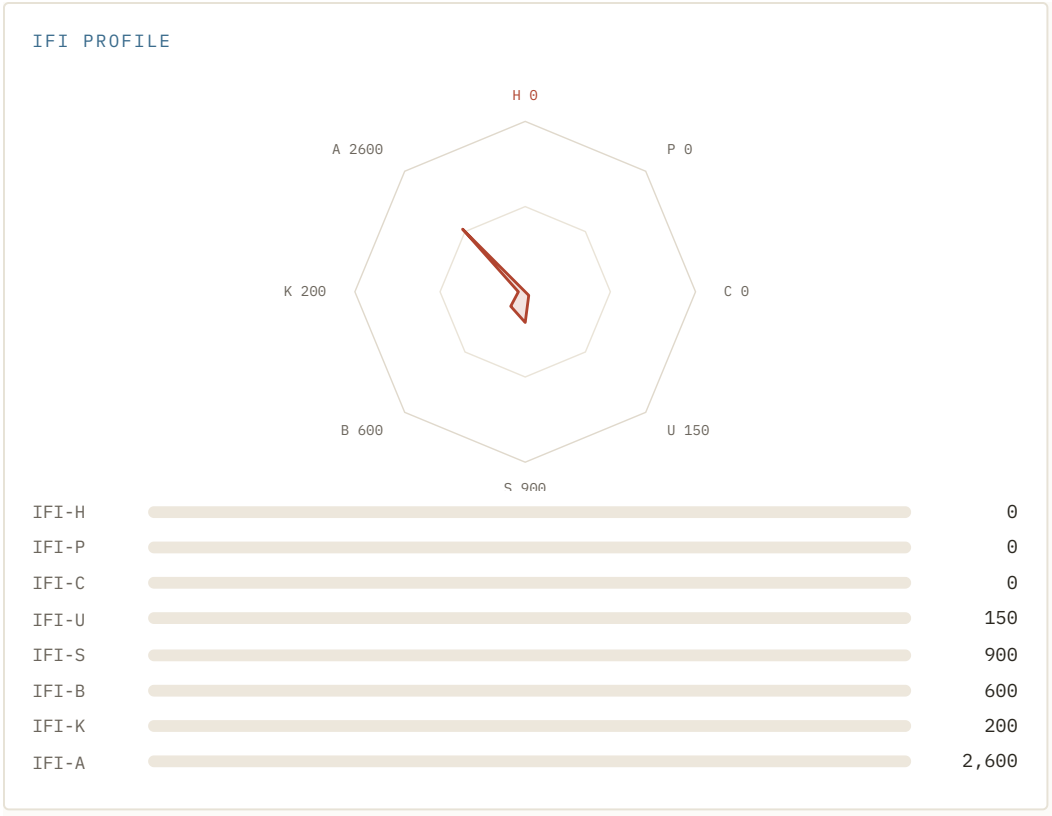

The **A of 2,600** is the whole architecture of the wine — its structure, its grip, its capacity to age. The **S of 900** is the grape's own ripe sugar. And the **B of 600** is the point of the chapter: it is low, and we touch it barely. A glass of wine wants a little bitterness; in a wine a little bitterness is *structure,* and removing it would give you a flabby, spineless thing nobody would pour twice. So on grape we leave the bitter axis nearly where we found it, lowering it only from about 720 to 600 — a trim of roughly 17%, deliberately *below* the 20–65% band we normally work in, and the only harshness we take is the green, unripe, stalky edge of a young wine. That edge is not really tannin maturity at all; it is partly the short, under-polymerized oligomers of unripe seed and partly the methoxypyrazines — the herbaceous, green-bell-pepper volatiles that mark fruit picked a week too early. We ease those and stop. The real work, the whole architecture of tannin and mouthfeel and the slow civilizing of a polymer over years, we hand to *Astringency* and to *Grape & Wine* (Volume XVIII), where it fills a book. For scale: our model predicts IFI-B of approximately 2,200 for Phoenix, the most fully spectral pepper in our family — more than three times this wine's bitterness — and no one would thank us for making a wine as un-bitter as we could. Knowing when *not* to turn an axis is as much a part of holding it as any edit we make.

Citrus — the brightest bitterness, and the one that waits

Citrus closes the family with the brightest bitternesses of all, and with the only one that is not fully present when you first taste it. In an orange — **120 constructs,** scored B and C — the harsh notes are two, and they arrive on different clocks.

The first is **naringin,** and it hides one of the prettiest pieces of enzymology in the whole library. Naringin is a flavanone, naringenin, carrying a sugar — but *which* sugar bond it carries decides everything. Attach the disaccharide rhamnosyl-glucose to naringenin through a 1→2 (neohesperidosyl) linkage and you get naringin, intensely bitter, the taste of white pith and membrane and, in a grapefruit, of the flesh itself. Attach the *same two sugars* through a 1→6 (rutinosyl) linkage instead and you get narirutin — essentially tasteless. One bond, rotated, is the whole difference between bitter and bland, and it is set by a single enzyme's regiochemistry: whether

the rhamnosyltransferase that finishes the molecule is a 1,2-rhamnosyltransferase or a 1,6. That is a gift to a designer, because it means we need not destroy the flavanone or its nourishment to debitter the fruit — we need only bias the glycosylation toward the tasteless linkage, the same outcome the food industry reaches for crudely with naringinase enzymes but achieved cleanly, in the living fruit, at the branch point. We tune the flavanone route that builds naringin and ease the pith's grip about **48%.**

The second harsh note is **limonin,** and it is the one that waits. Limonin is a limonoid, a bitter triterpenoid built on an entirely separate route from the flavonoids, and it is the source of citrus's notorious *delayed* bitterness — the way a fresh-squeezed juice can taste sweet at breakfast and turn bitter by lunchtime. The reason is kinetic. The fruit stores limonin not as limonin but as its tasteless open-ring precursor, limonoate A-ring lactone; in the acid of the juice, slowly, that precursor closes into bitter limonin. Model the formation as a first-order approach to a plateau,

```
C(t) = C∞ · (1 − e^(−kt)),
```

with a rate constant around k = 0.30 h^{-1}, so that the bitterness half-forms in t½ = ln 2 ÷ k = 0.693 ÷ 0.30 = **2.3 hours.** An untreated juice heading for a plateau of C∞ = 6.0 ppm has formed, by four hours — breakfast to lunch — C(4) = 6.0 × (1 − e^(−1.2)) = 6.0 × 0.699 = **4.2 ppm,** which in molar terms is 4.2 ÷ 470.5 = 8.9 µM, four times limonin's ~1 ppm (2.13 µM) bitter threshold. Sweet at nine in the morning, bitter by one. The lever is the limonoid UDP-glucosyltransferase that the fruit already uses to detoxify limonoids into tasteless glucosides: upregulate it and the precursor pool is captured as glucoside before it can ever lactonize, dropping the plateau to C∞ = 0.85 ppm — below the 1 ppm threshold even at full formation, so the curve never crosses into bitterness at all. The juice simply never turns. We moderate the limonoid supply so the glass holds its brightness through the afternoon:

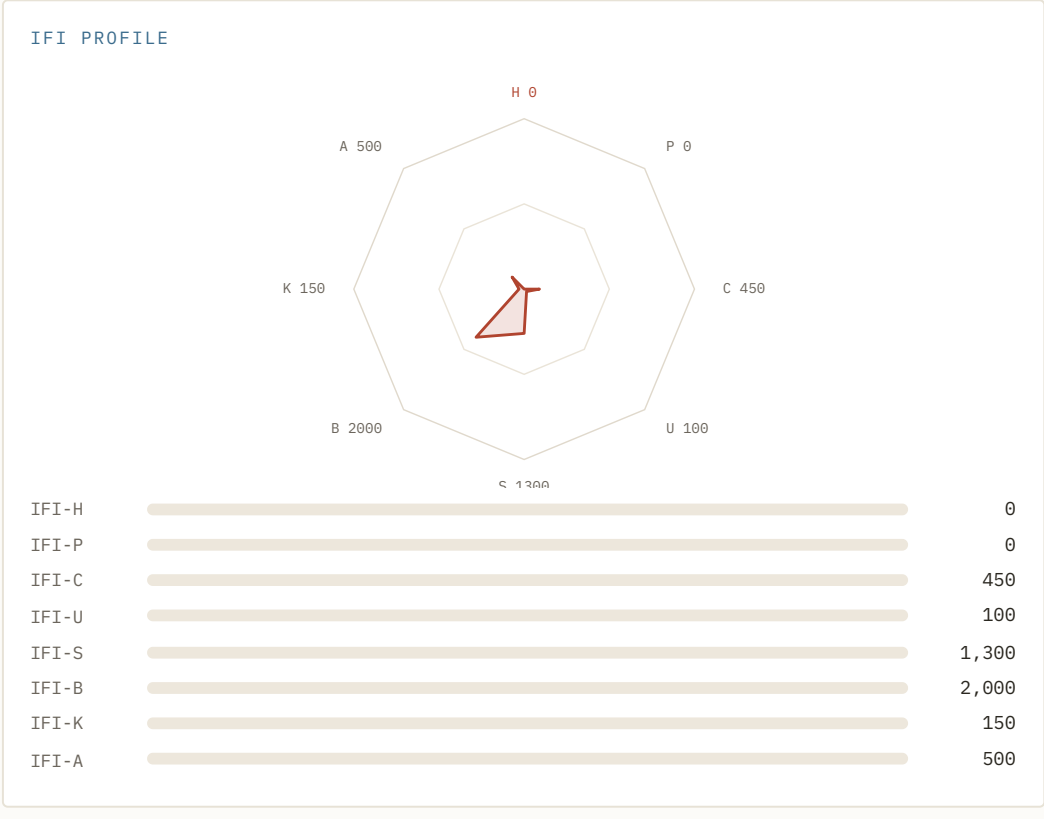

The **S of 1,300** is the fruit's own sugar, uncovered rather than added, once the pith's bitterness comes down. The **C of 450** is the cool, lifting citrus brightness carrying the glass. And the **B at 2,000** is the pith eased just enough that a whole-pressed orange — juiced with a little of its membrane, the way that keeps the most of the fruit and the most of its nourishment — no longer turns against itself by midday. It is worth seeing how far that edit reaches: a whole-press including the albedo, untreated, runs near IFI-B 3,400, which is bitterer than Phoenix's predicted 2,200 and genuinely unpleasant to drink — which is exactly why no one presses the pith. Easing the naringin route by 48% and capping the limonin plateau below threshold brings that 3,400 down to a drinkable 2,000, where the sugar and the cool lift finally come through. We did not add sweetness. We stopped the bitterness from hiding it.

The method, the edge cases, the honest limits

Five of the seven platform modules transfer to any sequenced genome, and these three plants share a road, so the method that eased one eased the others: read the harshness down to its molecule, place the molecule on the size ladder, and reach for the lever that sits at the molecule's own branch point — galloylation for tea, polymerization for grape, glycosyl regiochemistry and limonoid capture for citrus. But the family is as much a catalogue of ways to fail as a method, and honesty about the failures is the better part of the craft.

The first failure mode is over-cutting. Bitterness on these plants is never purely a flaw; it is structure wearing a harsh coat, and a tea stripped to IFI-B zero is slack, a wine with no grip is spineless, an orange with no edge is flat sugar-water. The 20–65% envelope is not a limit on our reach — we could cut further — but a statement of taste: past a point, attenuation takes the backbone with the bitterness. Grape is the standing reminder, deliberately trimmed only 17%, below the band, because the axis there is mostly doing its job.

The second is mistaking one axis for another. Because bitterness and astringency ride the same molecules sorted only by size, an edit aimed at the small galloylated monomers must be precise enough not to drag down the long polymers that carry astringency — on tea we hold A at 1,300 while dropping B by 30%, and if the galloyl edit were sloppy it would thin the backbone too. The partition is real but it is not free; it has to be engineered for.

The third is the clock. Citrus's delayed bitterness means a juice that tastes clean at the bench can fail three hours later, so the assay has to run to the limonin plateau, not to the fresh reading — a construct assessed only at t = 0 is assessed against the wrong number. The honest test of a debittered orange is the one you taste after lunch.

And the fourth is the anchor itself. The Index pins IFI-B to TAS2R38 for one stable yardstick, but the polyphenol bitters of this chapter in truth speak through TAS2R39 and TAS2R14 more than through the anchor receptor, and individual people carry different alleles and different numbers of these receptors — the same cup is measurably bitterer to some tasters than to others. The number we report is a population-

anchored loudness against a fixed reference, held to ≤7% on our side of the assay; it is not a promise that every tongue hears it identically. We think that is the right kind of number anyway — a shared, reproducible scale that a grower and a formulator and a drinker can all point at — but a designer who forgets the chord underneath the single note will eventually be surprised by a taster the note was never tuned for.

Across all three plants the pattern holds: a shared flavonoid road, a bitterness braided to astringency and sorted by size, and a method that turns one thread with care and leaves the other for the eater to enjoy and for the companion volumes to tell in full. The leaf, the berry, and the fruit each gave us the same lesson in a different grammar — that the whole of this axis is learning, molecule by molecule, which harshness is a flaw and which is the plant holding itself up. We quiet the first. We would never touch the second. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.



PART VI OF VI

Precision, Perception, and the Table

the taster and the non-taster · measured never claimed · the kindest subtraction

Chapter 17 · Bitter × Sweet — *subtract and uncover.*

COMPOSITION · IFI-B × IFI-S · the suppression two tastes work on each other

The most-requested work on this axis is almost never bitterness alone. It is the pairing — pull the harshness down *and* lift the sweetness up, so a food crosses in a single move from endured to enjoyed. A coffee without the scorch and with a rounder sweetness beneath it. A dark chocolate that grips less and reads sweeter without a grain of added sugar. A nutrition bar, dense with the cocoa and the grain that make it worth eating, that finally tastes like something a person reaches for a second time rather than finishes out of discipline. It is the composition reformulators ask us for most, and it is subtler than a simple stacking of two axes.

It is subtler because the two tastes are not independent on the tongue. Bitterness and sweetness **suppress each other** — a well-studied piece of taste psychophysics, playing out both in the periphery and in the brain's integration of the signals — which is exactly why a harsh coffee needs so much sugar to come into balance: the sugar is not sweetening the cup so much as spending itself to quiet the bitterness, and only what is left over reads as sweet. Turn the arithmetic around and the consequence is a gift. When we quiet the bitter agonist at its source, a sweetness that was present in the food all along is released from suppression and seems to *arrive* — uncovered, not added. This is the gentler, kinder half of the move, and it is why so much of the lift costs nothing: it was there the whole time, held down under the warning. What follows is that move written out in full — the receptors, the pathway, the index math, and the places where it fails.

The two tastes are not strangers on the tongue

Begin where the signal begins, because the suppression is not metaphor; it is wiring. Bitterness on our axis is anchored to **TAS2R38**, a class-A G-protein-coupled receptor studded across the apical membrane of Type II taste cells — the same receptor famous for reading phenylthiocarbamide and the glucosinolate breakdown products of a raw brassica. Sweetness is read by **TAS1R2·TAS1R3**, a class-C heterodimer whose two great clam-shell Venus-flytrap domains close around a sugar the way a hand closes around a held thing. Two different proteins, on two different populations of cells, listening for two different molecules. By every structural measure they are strangers.

And yet downstream they are the same machine. A bitter molecule settling into TAS2R38, or a sweet one into the flytrap of TAS1R2·R3, triggers the identical cascade: the receptor activates **α-gustducin**, gustducin frees **PLCβ2**, PLCβ2 cleaves membrane PIP_2 into **IP_3**, IP_3 opens the **IP_3R3** store and spills calcium into the cytosol, the calcium gates **TRPM5**, the cell depolarizes, and a bolus of ATP leaves through **CALHM1/CALHM3** onto the afferent fiber waiting below. Sweet and bitter are carried to the brain over wiring that is, past the first protein,

indistinguishable. Suppression lives in that shared channel and in the central integration above it — one taste spending the common currency the other needs, both at the cell and in the mind that sums them. This is why the move is a composition and not a sum: we are not adjusting two independent dials but two readings of one coupled system.



Cocoa, in cross-section — the pale seed is dense with the procyanidins and theobromine that make chocolate grip; most of what we ease on this axis is already drawn here.

Why the sugar is mostly spent, not tasted

Write the coupling down and it stops being a feeling and becomes something we can solve. Let **S** be the sweet stimulus the receptor actually receives — the loudness our index records on the sweet axis — and **B** the bitter stimulus on TAS2R38. The cleanest law that fits sweet–bitter suppression is a gain-control form, each taste dividing the other down:

```
felt sweet = S ÷ (1 + β·B) · felt bitter = B ÷ (1 + α·S)
```

The coefficients α and β are the suppression gains; the documented asymmetry is that sweet quiets bitter more firmly than bitter quiets sweet, so α runs a little above β. To make the consequence concrete we take the illustrative values our trained panels return to within their scatter — **$\beta = 2.0 \times 10^{-4}$** and **$\alpha = 3.0 \times 10^{-4}$** per index unit — and we work the cocoa case.

Hold the sweet stimulus fixed at the bar's recorded **S = 1,400** and change nothing about it. In raw cocoa the bitter fraction runs near **B ≈ 4,125** (we derive that figure two sections on, by inverting the attenuation; take it on loan for now). The eater therefore *feels*

- raw: $1{,}400 \div (1 + 2.0 \times 10^{-4} \times 4{,}125) = 1{,}400 \div 1.825 =$ **767**,

while the same unchanged 1,400 of sweet stimulus, once the bitterness is eased to the bar's **B = 1,650**, is felt as

- eased: $1{,}400 \div (1 + 2.0 \times 10^{-4} \times 1{,}650) = 1{,}400 \div 1.330 =$ **1,053**.

Nothing was added. The sweet stimulus is identical in both lines. Lowering the bitterness alone lifts the *felt* sweetness from 767 to 1,053 — a **+37%** rise that is pure release, the plant's own sweetness stepping out from under the warning. This is the uncovering, and it is the cheapest sweetness in any kitchen because the plant already paid for it.

Now read the same equation as an accounting of sugar. Suppose we want the eater to feel a sweetness of 1,053 either way. The eased cup reaches it on S = 1,400 (since 1,400 ÷ 1.330 = 1,053). The raw cup,

with its denominator of 1.825, must deliver S = 1,053 × 1.825 = **1,922** to feel equally sweet — 522 units of sweet stimulus more, every one of them spent not on pleasure but on canceling a harshness. The eased cup reaches the same destination on **27% less sweet stimulus** (522 ÷ 1,922). In the cup's own terms that is the spoon or two of sugar a hard coffee seems to demand: it was never sweetening the cup so much as paying down its bitterness, and only the remainder ever reached the tongue as sweet.

The asymmetry closes the loop, gently. Once the uncovered-and-lifted sweet is present at S = 1,400, it pushes back on what bitterness remains: felt bitter = 1,650 ÷ (1 + 3.0 × 10^{-4} × 1,400) = 1,650 ÷ 1.420 = **1,162**. The receptor still reports 1,650 — the index measures the protein, not the percept — but the mouth, with sweetness restored to it, experiences less. Subtract the bitter and the sweet rises; let the sweet rise and the bitter falls a little further. The move pays twice.

Reading the move in EC50

Everything above is anchored to one definition, the same one that lets a tomato, a chili, and a cup of tea be read on one page:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. A *potent* stimulus reaches half-activation at a low concentration — a small EC_{50} — and so scores high; a faint one needs a flood to be heard, carries a large EC_{50}, and scores low. Because the reference sits at the thousand mark by convention, the whole index inverts back to concentration with a single step: **EC_{50}(candidate) = 1,000 ÷ IFI**, read in multiples of the reference's own half-activating concentration. Run every line of the bar through it, and the eight numbers resolve into eight statements about how loudly each compound speaks.

READING (IFI)	EC_{50} (× REFERENCE)	WHAT IT SAYS
5,000 *(scale ceiling)*	0.200	loudest the index records
3,000 *(engineered stevia, S)*	0.333	half-activates at a third of reference
2,200 *(Phoenix, B)*	0.455	the loudest whole-plant bitter on file
1,650 *(bar, B)*	0.606	eased, yet still firmly present
1,400 *(bar, S)*	0.714	the sweet, read as 1.4× reference
1,000 *(reference)*	1.000	the anchor, by definition
900 *(bar, K)*	1.111	the body holding the base together
800 *(bar, A)*	1.250	the structure, just below reference
300 *(bar, U)*	3.333	a trace, needing three times reference to be heard

Behind the inversion sits a dose–response curve, and its shape governs what easing can and cannot buy. A receptor's fractional response to an agonist follows a Hill relation, **R ÷ R(max) = $[A]^n$ ÷ (EC_{50}^n + $[A]^n$)**, where n is the cooperativity of binding and, at [A] = EC_{50}, the response is exactly half — the point the index pins to. Two features of that curve decide the chapter. Far below saturation the response climbs nearly in proportion to [A], so halving the agonist a plant makes very nearly halves the drive it puts on the receptor; that

near-proportionality is precisely what lets the attenuation law stay linear. Near the top the curve flattens, and there a large cut in agonist buys only a small drop in response — a food already saturating TAS2R38 must be brought a long way down before the tongue registers relief, which is why the hardest bitterness (a dark roast, an unsweetened cocoa) resists the first increments of easing and then seems to yield all at once. We work the middle of the curve on purpose, in the sub-saturating band where a given percentage of attenuation returns the same percentage off the score.

The subtraction has its own law, and it is clean. If the bitterness of a food comes mostly from one agonist present at some fraction of the food's mass, then the food's effective EC_{50} at TAS2R38 scales inversely with how much of that agonist remains — halve the agonist and you must double the food to reach half-activation. Because IFI is one-over-EC_{50}, that makes the bitter score **linear in the agonist that survives**:

```
IFI-B(eased) = IFI-B(raw) × (1 − A)
```

where **A** is the attenuation, and on this axis we hold **A between 20% and 65%** — the honest working range, never zero. The bar records B = 1,650; run the dial backward across its full span and the cocoa's native bitter fraction must have registered here:

ATTENUATION A	FRACTION KEPT (1 − A)	IMPLIED NATIVE B = 1,650 ÷ (1 − A)
20%	0.80	2,063
35%	0.65	2,538
50%	0.50	3,300
60%	0.40	4,125
65%	0.35	4,714

We do not assert a single raw figure; we bracket it from canon. Raw cocoa's bitter fraction sat somewhere between **~2,060 and ~4,710**, and because this base was eased near the ceiling we run in practice, its native bitterness lived near the top of that bracket — above even **Phoenix's whole-plant B of 2,200**, which is no exaggeration for anyone who has bitten unsweetened chocolate. That is the 4,125 we borrowed for the suppression arithmetic, earned now rather than assumed.

The source of the bitterness, and the lever that lowers it

A score is a place to stand; the work is upstream, in the plant's chemistry. Bitterness is not one molecule but a family read by a family — some two dozen TAS2Rs — and we anchor the ruler on TAS2R38 while the attenuation acts on the agonist load that drives the whole set. Three sources account for nearly all of the harshness on this shelf, and each has its own enzymatic lever.

In a **brassica** — our most direct tie to TAS2R38 — bitterness is glucosinolates, cleaved by myrosinase on tissue damage into the isothiocyanates and goitrin that the receptor reads as alarm. In **cocoa**, it is the **procyanidins**, oligomers of flavan-3-ols condensed down the

flavonoid branch that hangs off the shared phenylpropanoid trunk — **PAL → C4H → 4CL → HCT → COMT** feeding chalcone synthase and, through the leucoanthocyanidin and anthocyanidin reductases (LAR, ANR), the catechin and epicatechin units that polymerize into grip — together with **theobromine**, a methylxanthine built on the purine pathway by sequential N-methyltransferases. In **coffee**, it is chlorogenic acids drawn from that same phenylpropanoid trunk, pushed by roast into the chlorogenic-acid lactones and, deeper, the phenylindanes that carry the late, scorched bitterness, with caffeine underneath. The lever, in every case, is to lower the flux into the offending branch — the glucosinolate load, the proanthocyanidin condensation, the N-methyltransferase step — by exactly as much as the 20–65% window allows, no more. A genome-scale flux-balance model finds where the cut can be taken without starving the plant or collapsing a trait that shares the pathway, and the same protein-language-model scoring we use to design an enzyme tells us which residue, which step, will move the flux and which will merely wobble it.

Our computational designs span a shelf of hosts, each scored on the same eight axes.

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Grape	1,610	B / A / M	Hard
Broccoli	380	B / M	Moderate
Tea	280	U / B / A	Hard
Broccoli-Kale	280	B	Hard
Orange	120	B / C	Hard
Hops	80	B	Moderate
Potato	80	B	Easy
Coffee	40	B	Hard
Cabbage	40	B	Moderate
Kale	40	B	Moderate



Coffee, the cup the whole axis is judged by — ease the scorch and the bean's own faint sweetness, inaudible under a harsh roast, becomes plainly present.

A clarified coffee is the plainest proof of the uncovering. Hold the sugar in the cup constant and merely lower the bitterness, and trained panels rate the same cup as *sweeter* — the sweetness readout rises though not a grain of sugar was added, because the signal was only ever being masked. In a de-bittered cup this is most of the pleasure: the bean carried that faint sweetness all along, and so many drinkers quietly set aside the sugar they always thought they needed. The plant paid for it; we only stopped hiding it.

The sweet we add is not sugar

Where a food wants more sweetness still than uncovering alone provides, we raise the plant's own sweet signal — not with sugar but with molecules the sweet receptor reads *as* sugar though they are not. The flytrap of TAS1R2·R3 closes on a **steviol glycoside** — stevioside, or the cleaner rebaudiosides A and M, built when the diterpene steviol is dressed in sugars by a short relay of **UGT** glycosyltransferases (UGT85C2, UGT74G1, UGT76G1) — as readily as on sucrose, and on the small sweet protein **brazzein** (UniProt P56552), fifty-four amino acids that wedge into the open flytrap and hold it shut. An engineered **stevia** on this method clears an **IFI-S of 3,000**: by the inversion that is an EC_{50} at **one-third of reference**, a sweet stimulus three times as loud, per unit, as the anchor. The tongue reads sweetness fully; the chemistry delivering it is simply not cane sugar, and because steviol glycosides are themselves on the order of two to three hundred times as sweet as sucrose by mass, a fraction of a milligram does the work of a spoonful — with no caloric sugar entering the food at all. The sweet-side shelf is its own ladder:

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

Paired with a quieted bitterness, a very little of that signal goes a long way, for the reason the suppression math already showed: the sweet is no longer being spent to cancel a harshness that is no longer there. Lift the bitter first and the sweet you add is sweet you keep.

Composing, not stacking

Here is the composition that the better-for-you shelf leans on hardest — a cocoa-and-oat bar base, its cocoa eased and its sweetness uncovered and only gently lifted, with no cane sugar in it at all:

IFI PROFILE

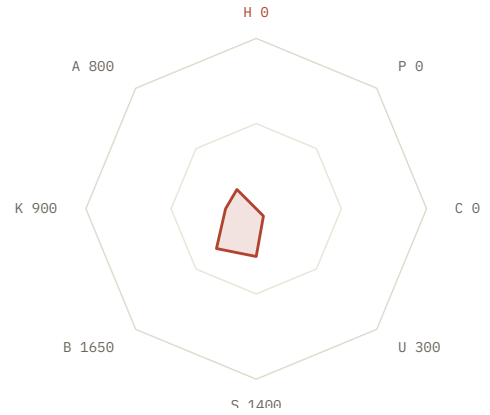

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	300
IFI-S	1,400
IFI-B	1,650
IFI-K	900
IFI-A	800

Read it line by line and the restraint is the point. The **B of 1,650** is the cocoa's harsh procyanidin fraction eased from a native bracket that topped four thousand; the **S of 1,400** is partly the sweetness uncovered as that harshness falls and partly a whisper of steviol lift, no more; the **K of 900** and the **A of 800** are the body and the structure that keep the bar from tasting like a candy pretending to be food. Composing the two is not addition, and a naïve stack disappoints: quiet the bitter and pile on the sweet without judgment and you get something that tastes neither clean nor honestly sweet, cloying in the way a cheap "less-bitter" product does. The flux-balance model finds the ratio that holds — how far to quiet the bitterness before the lift reads as honest rather than candied.

It is the same mathematics of mutual accommodation that lets *Swicy* set heat against sweetness without either canceling the other, heat on the scale where **our model predicts Phoenix at approximately 2.94 million SHU, reading as IFI-H 4,200**, or lets a restored tomato carry its umami without tipping. Phoenix is the one whole plant we have designed to register on all eight axes at once, which is why we read everything against it — a common ruler laid across the whole garden:

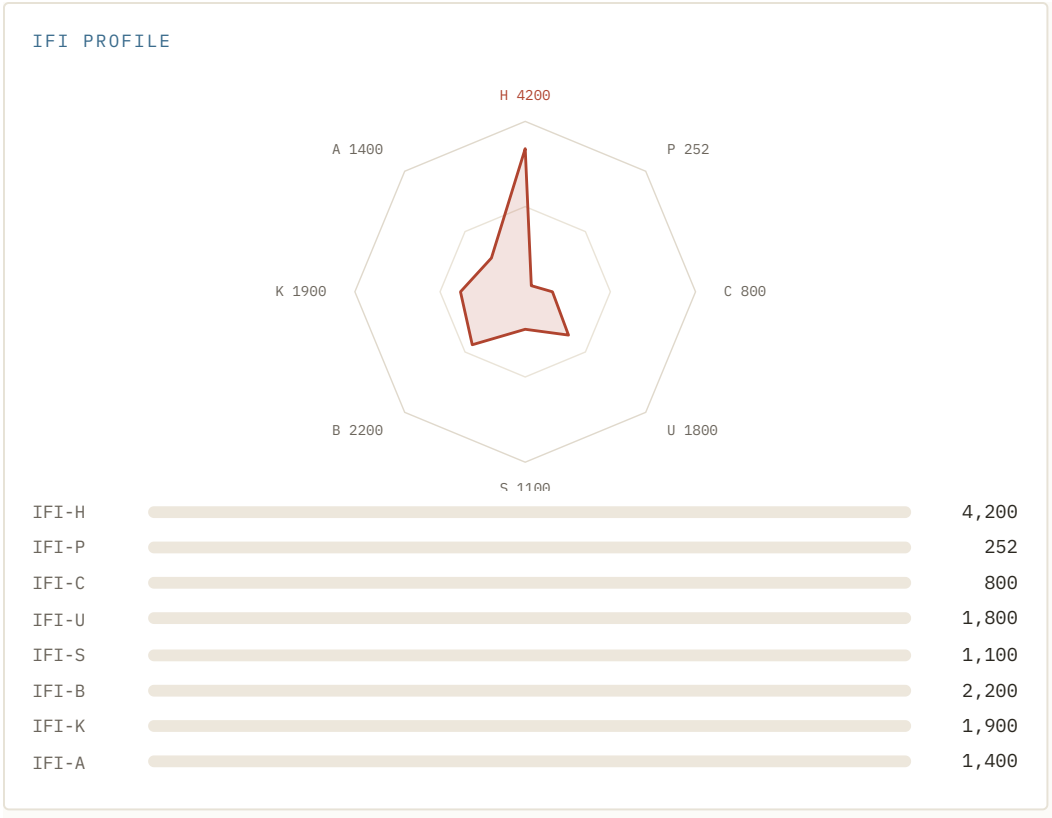

Against that ruler the bar is legible at a glance: no heat, no pungency, no cool-burn, a trace of umami, and the two axes of this chapter — B and S — set in a composed, deliberate ratio rather than left to collide.

Precision is the whole game

A composition lives or dies on reproducibility, because when you compose two axes their errors compose too. Our batch-to-batch coefficient of variation is held to **≤7%**; a wild-type plant drifts at **42%**. That gap looks like mere tidiness until you ask what happens to the *ratio* the composition depends on. For two independent readings combined as a ratio, relative variances add in quadrature, so the composed spec carries

- engineered: $\sqrt{(0.07^2 + 0.07^2)} = \mathbf{9.9\%}$,
- wild-type: $\sqrt{(0.42^2 + 0.42^2)} = \mathbf{59\%}$.

At 10% the target holds — you can ask for a bitter-to-sweet ratio and land on it, batch after batch, so the bar that was generous last week is generous this week. At 59% the composition is unrepeatable: the same formula reads cloying one run and harsh the next, and no amount of taste skill rescues a spec that will not sit still. The felt-sweetness percept is steadier yet, because the suppression denominator $(1 + \beta \cdot B)$ damps its own noise — a 7% wobble in B moves 1.330 by only 1.7% — so the composed percept holds at $\sqrt{(0.07^2 + 0.017^2)} \approx \mathbf{7.2\%}$ even as it rides on two coupled readings. Precision is not housekeeping here. It is the precondition for composing at all; numbers are a form of respect for the reader, and for the formulator who has to make the thing twice.

Where it fails, and where it stops

We are candid about the edges, because the move is powerful enough to be abused.

Over-lifting curdles into candy. Push the sweet stimulus to S = 2,800 while the bitter stays at 1,650 and the felt sweetness runs to $2{,}800 \div 1.330 = 2{,}105$, well past the honest window, while the residual bitter — now $1{,}650 \div (1 + 3.0 \times 10^{-4} \times 2{,}800) = 897$ — is over-quieted into shapelessness. The flux-balance ratio exists precisely to refuse this.

Over-subtraction hollows the food out. Bitterness carries identity: a coffee stripped past a floor stops tasting like coffee, and a cocoa past it tastes flat. This is why the working range stops at 65% and why *"Zero-Bitter"* is the axis's aspiration, not a literal promise — and why the bar keeps its K of 900 and A of 800 as body and backbone rather than chasing the harshness all the way down.

The sweetener can bring its own bitter. Steviol glycosides, for all their grace at the flytrap, also brush TAS2R4 and TAS2R14 — a different bitter receptor than the one we quieted — which is the licorice edge of a heavy-handed stevia. We lean on the cleaner rebaudioside M, on brazzein, and most of all on uncovering over addition, so the lift that solves one receptor does not wake another.

The bitter and the astringent are often the same molecule. Cocoa's procyanidins read on both TAS2R38 (bitter) and the salivary proline-rich proteins (astringency), so a cut aimed at B can drag A down with it and leave the bar thin. We decouple where the chemistry allows — easing the methylxanthine for bitterness while preserving the higher-weight tannins that give structure — which is the whole reason the bar's A of 800 is a target and not an accident.

Part of the uncovering is in the mind. Because mixture suppression is partly central, the released sweetness is a little sensitive to attention and expectation in a way the receptor EC_{50} is not. The index measures the protein and is robust; the percept rides on top of it and is steadied by panel training rather than guaranteed. We report what the composition does to flavor, and we stop there. What a food is for beyond its taste is the eater's business, not a claim we make — nourishment is an intention we bring to the plant, never a promise we print on it.

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

Subtract and lift, in the right proportion, and a difficult food becomes a generous one — the cocoa and the grain that make a bar worth eating turned into something a person reaches for a second time rather than finishes out of discipline. For the sweet half of this work in full the shelf hands you to *Sweetness* (Volume VII); for the composition as its own craft, to *Bitter-Sweet Balance* (Volume XXVIII). On this axis we supply the subtraction — and we have learned, batch after batch, that the subtraction done well is often most of the sweetness anyone needed.



Oat, the bar's quiet grain — dense, honest, nutritionally earned; the kind of food our family has always meant food to be.

Chapter 18 · The Taster at the Table — *polymorphism, perception, and one number against one experience.*

IFI-B and the reference genotype · PAV versus AVI · the honest limit of a single score

There is an honesty this axis owes the reader that the other seven can mostly avoid. A single IFI-B number can never be a single human experience of bitterness — because, more than any other taste, bitterness is read differently by different people, written into their genes before they ever sit down to eat. Heat arrives on everyone's trigeminal nerve by the same door; sweetness and umami vary at the margins; but bitterness is the one sense the body guards with a whole family of sentinel receptors, and those sentinels are not identical from mouth to mouth. The taster and the non-taster sitting at the same table, eating from the same bowl, are not having the same meal. One reads the broccoli as a clear warning. The other barely registers it. When we print an IFI-B on a page, whose tongue are we printing it for?

We are a family of scientists, and we grew up cooking for each other across exactly this divide — the same floret adored at one seat and refused at the next, and no one at the table wrong. So we will not pretend the number is something it is not. The whole of this chapter is the discipline of writing a true, reproducible score while naming, out loud and in arithmetic, precisely what it cannot promise.



The same floret, two meals — a strong taster and a non-taster meet a brassica as different foods, which is why an honest bitterness number must name the tongue it is written for.

The two tongues at one table

The divide has a name and a molecular address: **TAS2R38**, a single bitter-taste receptor whose gene comes in two common editions that differ at just three letters of its code. Those three positions decide, more than any kitchen technique ever could, whether a brassica tastes like food or like a threat.

TAS2R38 is a class-A, seven-transmembrane G-protein-coupled receptor — the same architectural family as the receptors for light and adrenaline — sitting on the type II taste cells of the tongue. Three amino-acid positions along its folded length vary between people: residue 49, residue 262, and residue 296. Read the taster's letters in order and you get **P-A-V** — proline, alanine, valine. Read the non-taster's and you get **A-V-I** — alanine, valine, isoleucine. Two haplotypes, three substitutions, and essentially the entire famous split between the people who wince at raw kale and the people who wonder what the fuss is about.

POSITION	VARIANT (DBSNP)	TASTER ALLELE · PAV	NON-TASTER ALLELE · AVI	CONSEQUENCE AT THE RECEPTOR
49	rs713598	Pro (P)	Ala (A)	pocket geometry; efficient vs. poor agonist fit
262	rs1726866	Ala (A)	Val (V)	transmembrane packing near the binding cleft
296	rs10246939	Val (V)	Ile (I)	helix 7 coupling; gates G-protein engagement

Because we each carry two copies of the gene, the practical reading is a diplotype. **PAV/PAV** is the strong taster — both receptors built in the responsive form. **AVI/AVI** is the non-taster — both built in the form that barely answers the same molecule. **PAV/AVI** is the heterozygote, intermediate, nearer the taster because a single working PAV allele confers most of the sensitivity. Rarer haplotypes exist — AAV, AAI, PVI among them — in a small slice of the population, each a slightly different tuning of the same instrument. In European-ancestry populations, very roughly seven in ten carry at least one PAV allele and read as tasters; the remaining three in ten are AVI/AVI non-tasters. The frequencies shift across ancestries, which is itself part of the honesty: there is no single human tongue to write a number for, only a distribution of them.

What the tongue does when a brassica arrives

The reverence we hold for a brassica is partly reverence for its chemistry, so it is worth following the molecule all the way in. A raw broccoli floret is not bitter on the plant. Its bitterness is a weapon the plant keeps disassembled: **glucosinolates** stored quietly in one cellular compartment, and the enzyme that cuts them — **myrosinase**, a thioglucosidase — held apart in another. Nothing happens until the tissue is wounded. A bite, a knife, a chewing molar breaks the compartments open, enzyme meets substrate, and in seconds the plant arms itself in the mouth.

Myrosinase cleaves the glucose off the glucosinolate, and the unstable aglycone that remains rearranges — a Lossen rearrangement — into an **isothiocyanate**: sinigrin yields allyl isothiocyanate, progoitrin yields goitrin, an oxazolidinethione. The exact product depends on pH and on whether an epithiospecifier protein is present to divert the reaction toward nitriles instead. What matters for the tongue is the chemical group that emerges at the center of these molecules: the **N–C=S** motif, the thiourea-like pharmacophore. That same N–C=S is the business end of phenylthiocarbamide (PTC) and of 6-n-propylthiouracil (PROP), the two synthetic compounds taste science has used for a century to sort tasters from non-tasters. The brassica and the lab standard speak to TAS2R38 in the same word.

When that motif docks into the receptor's binding cleft — a pocket formed by the third, fifth, sixth, and seventh transmembrane helices — a **PAV** receptor changes shape decisively and an **AVI** receptor shrugs. The three polymorphic residues sit where they can tilt helix 7 and re-pack the cleft, so PAV couples efficiently to its G-protein and AVI couples poorly. In the responsive case, the activated receptor releases **α-gustducin**; the freed **Gβ3γ13** subunits switch on **phospholipase C-β2**, which hydrolyzes PIP_2 into **IP_3**; IP_3 opens the type-3 IP_3 receptor on the endoplasmic reticulum and floods the cell with calcium; the calcium opens **TRPM5**, a monovalent-cation channel, and the cell depolarizes. Type II taste cells have no classical synapse — instead the depolarization opens **CALHM1/CALHM3** channels and the cell releases ATP directly, non-vesicularly, onto the P2X2/P2X3 endings of the gustatory nerve. From there the signal climbs: nucleus of the solitary tract, thalamus, the insular gustatory cortex, and only then the conscious wince.

Every link in that chain is shared between a taster and a non-taster — the gustducin, the PLC-β2, the TRPM5, the ATP. The single place they diverge is the very first handshake, the fit of N–C=S into a cleft that one person built in proline-alanine-valine and another in alanine-valine-isoleucine. One switch, read at the threshold of the whole cascade, and the meal splits in two.

One number, built honestly

So we anchor the axis where the signal is real. **IFI-B is calibrated against a defined reference genotype** — the receptor built in the lab from the taster **PAV** haplotype, the form that responds, the form a standard can be scored against. The Index's single governing relation is the same for all eight axes:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. A smaller EC_{50} means a louder stimulus — it reaches half-maximal activation at a lower dose. We fix the scale by declaring that the reference bitter standard, PROP delivered to the PAV receptor, reads exactly **1,000**. By stated convention, we set the reference at **EC_{50}(reference) = 100 µM**, which fixes the working form of the axis:

> $\text{IFI-B} = 100{,}000 \div EC_{50}(\text{candidate}) \cdot \text{with } EC_{50} \text{ in µM}$

Now the number reads in both directions. Give the formula an EC_{50} and it returns a score; give it a score and it returns the dose at which that food half-activates the reference tongue. That invertibility is the whole point — a score you cannot run backward is a score you cannot check.

BITTER FRACTION	EC_{50} AT PAV RECEPTOR	IFI-B = 100,000 ÷ EC_{50}	IN WORDS
Reference standard (PROP)	100 µM	**1,000**	the anchor — half-max at the calibration dose
A loud bitter	50 µM	2,000	half-max at half the dose: twice as loud
Phoenix bitter fraction	45.5 µM	**2,200**	the full-spectrum cultivar's predicted B
A quiet bitter	200 µM	500	needs twice the dose: half as loud
A faint bitter	500 µM	200	a hum beneath the other flavors

The Phoenix row is worth pausing on, because it is a canonical reading and it demonstrates the inversion. Phoenix — our proof-of-concept cultivar, the first designed to register on all eight axes at once — carries a predicted **IFI-B of 2,200**. Run the formula backward: EC_{50} = 100,000 ÷ 2,200 = **45.5 µM**. That is the concrete, physical meaning of the score 2,200 — Phoenix's bitter fraction is designed to half-activate the reference PAV receptor at 45.5 micromolar, a little louder than twice the calibration standard. We did not estimate that bitterness or taste-panel it into a word like "moderate." Our model predicted the EC_{50}, against a receptor we defined in our computational framework, and the number is that dose read through a formula.

The reference condition is not a dodge; it is the move every measurement science makes when it fixes a standard state so a result means the same thing in two labs on two different days. But fixing it carries an obligation: to say, in the same breath, what the number does *not* capture. A non-taster will experience a given plant's bitterness as milder than its IFI-B suggests, and a strong taster as sharper. The score is a true, reproducible property of the stimulus. It is not a promise about any one person's perception of it — and the gap between those two statements is not vague, it is arithmetic.

The arithmetic of two tongues

Here is the gap, in numbers. The AVI receptor is not broken; it is less sensitive. It answers the same N–C=S motif, but needs far more of it to reach half-activation — its EC_{50} for the same ligand is much higher. Call the ratio of the two sensitivities $\rho = EC_{50}(AVI) \div EC_{50}(PAV)$. Functional assays and threshold psychophysics put ρ comfortably on the order of ten, and in whole-mouth perception often larger. Take a representative **ρ = 10** and follow one floret across the table.

Suppose a brassica reads **IFI-B = 2,000** on the PAV-anchored axis — a loud, clear bitter, EC_{50} of 50 µM at the reference tongue. The strong **PAV/PAV** taster meets it at full strength: 2,000, and then some, because density of sensors adds its own gain on top. The **AVI/AVI** non-taster meets the identical floret through a receptor ten times less sensitive — an effective EC_{50} near 500 µM — so the same chemistry lands as roughly 2,000 ÷ 10 = **200**. Not a warning. A hum. The **PAV/AVI** heterozygote sits between, nearer the taster because one good allele does most of the work: call it half to two-thirds of the strong taster's intensity. One bowl of broccoli, three experiences spanning a factor of ten, and the plant never changed.

This is also why the reference genotype must be *declared*, not assumed. The choice of reference tongue multiplies the entire axis by a constant. Had we silently anchored IFI-B to the AVI receptor, the

reference PROP would still read 1,000 by definition — but every real plant, scored against a duller sentinel, would compress toward the quiet end, and the whole ladder would reorder. Two labs that each chose a reference haplotype without saying so would publish numbers for the same broccoli that disagreed by about ρ — a factor of ten — and both would be internally correct and mutually useless. A bitterness score without its reference genotype printed beside it is not a measurement. It is a rumor with a decimal point. We print **PAV** on every IFI-B in this library for exactly that reason.

Dose, response, and the shape of the curve

Behind EC_{50} is a curve, and the curve is where the design lever lives. Receptor activation follows a Hill relation between ligand concentration [L] and fractional response R/E_max:

$$R \div E_{max} = [L]^n \div (EC_{50}^n + [L]^n)$$

At [L] = EC_{50} the response is exactly half of maximum — that is the definition the whole Index hangs on. The Hill coefficient n for TAS2R38-type activation sits near 1 to 1.5; the curve is a smooth sigmoid on a log-concentration axis. The useful fact for a designer is what happens at the low end. When [L] is well below EC_{50} and n ≈ 1, the relation linearizes:

$$R \div E_{max} \approx [L] \div EC_{50} \cdot \text{for } [L] \ll EC_{50}$$

In that near-threshold regime — exactly the regime where a plate is decided *acceptable* or *too bitter* — response is very nearly proportional to the concentration of the bitter ligand. Halve the ligand and you roughly halve the perceived bitterness. This is the quiet machinery under every attenuation number we publish: because IFI-B scales with the stimulus's apparent potency, and apparent potency scales with how much bitter ligand the tissue carries, lowering the ligand pool lowers the score in direct proportion. Cut the bitter-ligand concentration to a fraction f of the native plant, and the engineered tissue reads about f times the native IFI-B.

It is worth grounding these abstractions in real mass, because micromolar is a word that hides its weight. PROP has a molar mass of 170.23 g/mol, so the 100 µM reference anchor is 100×10^{-6} mol/L × 170.23 g/mol ≈ **17.0 mg/L** — seventeen milligrams of standard in a liter of water, half-activating the taster's receptor. PTC, molar mass 152.2 g/mol, works out to about 15.2 mg/L at the same 100 µM. On the plant side, a pungent brassica might store glucosinolate at 20 µmol per gram of dry tissue; with sinigrin at 397.5 g/mol that is $20 \times 10^{-6} \times 397.5$ = 7.95 mg/g, nearly **0.8% of the leaf's dry weight** committed to the disassembled weapon. Attenuation is the arithmetic of walking that fraction down — carefully, because the same chemistry carries compounds the plant is treasured for.

The attenuation band, aimed at the hardest tongue

Our published lever for this axis is a **20–65% attenuation** of bitterness, and the brassicas sit at the demanding end of it. Translate the band through the near-threshold relation above and it becomes a

set of multipliers on the stimulus:

- **20% attenuation** means f = 0.80 — the tissue still carries four-fifths of its bitter load. To a non-taster this is imperceptible, and to a taster it is merely a touch gentler. Too little to bother a brassica with.
- **62% attenuation**, our brassica ceiling, means f = 0.38 — a little over a third of the native bitter ligand survives. A floret that read **IFI-B 2,000** on the reference tongue now reads about 2,000 × 0.38 ≈ **760**.

Why aim at 62%, and why call it a ceiling rather than a target to beat? Because we design to the *population*, not to a single tongue, and the hardest case sets the specification. The non-taster never needed help; the brassica was never a problem at their seat. The person for whom the vegetable was most a warning is the **strong taster** — and often the child, whose bitter sensitivity runs highest in early life, whose alarms are turned up by biology precisely when caution matters most. A twenty-percent quieting does nothing for that child. A sixty-two-percent quieting is the setting that carries the strong taster across the one threshold that matters — from *refuse* to *eat* — bringing their near-threshold response down below the line where the plate gets pushed away. Aim the hard end of the band at the hard end of the population, and the easy cases take care of themselves. When it works, the non-taster notices nothing and the taster, for the first time, eats the same amount as everyone else. The measure of success on this axis was never that the number fell. It is that the table stopped dividing by genotype.



A brassica cross-section: the glucosinolate-dense tissue whose bitterness we walk down to a third — the arithmetic of the 62% ceiling, where quieting the alarm stops short of emptying the plant's own pantry.

The ceiling is honest about a real tension. The glucosinolate chemistry that reads bitter is the same chemistry that carries the compounds a brassica is grown and loved for; walk the ligand pool too far down and you do not just quiet an alarm, you empty a pantry. There is a cleaner path than brute reduction — shifting the hydrolysis with an epithiospecifier toward nitriles, which the TAS2R38 pocket largely ignores — but that trades one chemistry for another and can introduce its own off-notes, a failure mode we test for rather than assume away. So 62% is where the quieting and the keeping meet: as much relief as we can give the sharpest tongue while protecting what the plant is for. That restraint — feeding the body the plant's whole gift, not a hollowed version of it — is the ethic under the number, and it is why the ceiling is a ceiling.

Even a cultivar built to speak loudly on every axis carries its bitterness as a predicted, bounded quantity rather than an accident. Phoenix is the clearest example — designed, gene by gene, to register on all eight axes at once, and its bitterness is not stripped to zero but *placed*, a deliberate 2,200 against the reference tongue, in balance with the rest:

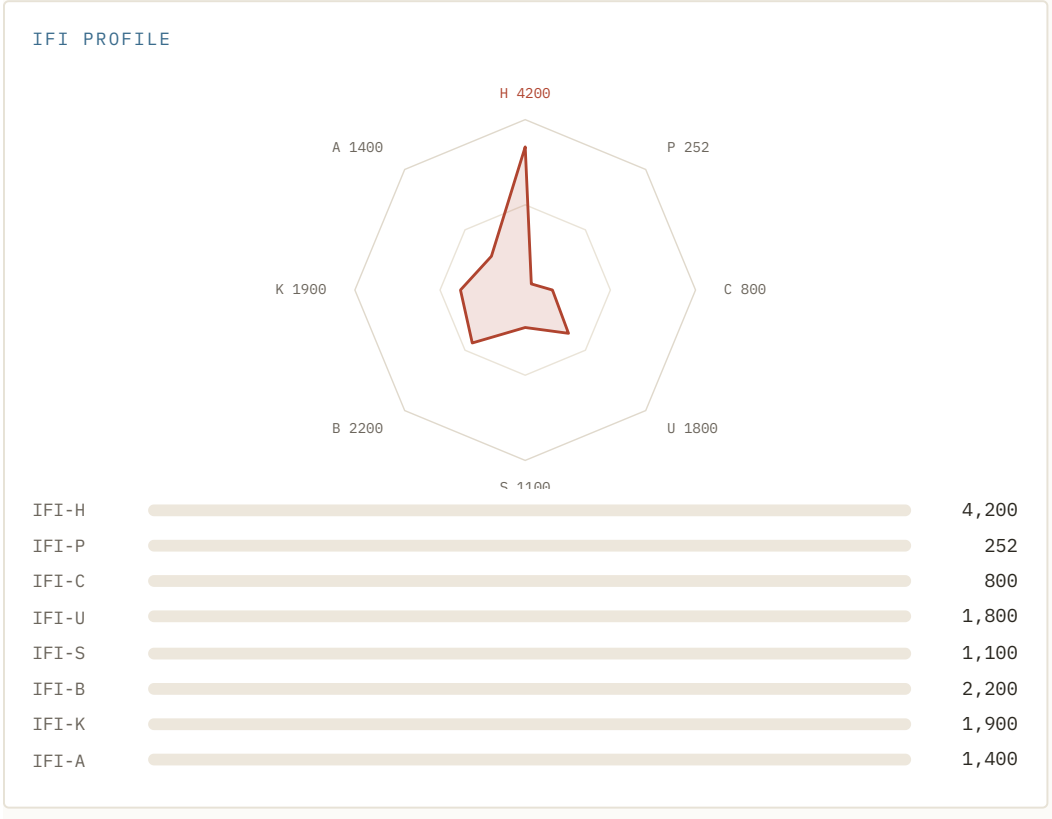

Read that B of 2,200 the way the whole chapter asks you to: it is the loudness of the stimulus against the defined PAV reference, EC_{50} 45.5 µM, reproducible batch to batch — and it is *not* a forecast of what the non-taster three seats down will feel, which by the arithmetic above is nearer 220.

Beyond the single switch

TAS2R38 is the famous sentinel, but it is one of twenty-five. Humans carry roughly **25 functional TAS2R bitter receptors**, and the PAV/AVI switch governs only the thiourea world most sharply — the brassicas, PROP, PTC, the glucosinolate alarm. The rest of bitterness is read elsewhere. Coffee's bitterness is mostly not a TAS2R38 event at all: caffeine and the chlorogenic-acid lactones that roasting generates speak to TAS2R7, TAS2R10, TAS2R14, TAS2R43, TAS2R46 and their neighbors; cacao's theobromine and polyphenols the same. This is a genuine, declared limit of a single-receptor axis. IFI-B anchored to TAS2R38 reads brassica and thiourea bitterness most faithfully, and reads a coffee or a hop as a projection onto one receptor of a sensation the mouth actually assembles from many. We say so on the page rather than letting the number overreach.



Coffee in cross-section — a bitterness the mouth assembles from many receptors at once, which is why one person can be sharply tuned to a cup and dull to a hop, and why a single-receptor axis reads it as a projection, honestly labeled.

Because each of the other twenty-four receptors carries its own common variants, a person can be sharply tuned to a coffee and dull to a hop, or the reverse, depending on which sentinels they happened to inherit — bitterness is not one dial but a two-dozen-key instrument, tuned differently in every mouth. Two more forces sit on top of the genetics. **Supertasters**, a subset of PAV tasters, carry a far higher density of fungiform papillae — physically more taste buds per square centimeter, more sensors reporting — so they experience not just bitterness but most oral sensation more intensely. Note that this gain is *anatomical*, downstream of receptor sensitivity: it multiplies the signal after the EC_{50} handshake, which is exactly why a receptor-potency axis like IFI-B cannot capture it even in principle. And **age** plays in: bitter sensitivity tends to run highest in childhood and soften across adulthood, which is the real reason a child is the hardest sell on a brassica and so often the adult, years later, who cannot remember what the fuss was.

The twelve host organisms we work on this axis span most of where bitterness actually lives in the kitchen, and the construct counts are a fair map of where the work has gone:

HOST ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Grape	1,610	B / A / M	Hard
Broccoli	380	B / M	Moderate
Tea	280	U / B / A	Hard
Broccoli-Kale	280	B	Hard
Orange	120	B / C	Hard
Potato	80	B	Easy
Hops	80	B	Moderate
Coffee	40	B	Hard
Cabbage	40	B	Moderate
Kale	40	B	Moderate
Radish	40	B	Hard
Rapeseed	40	B	Hard

Twelve hosts, about **3,030 constructs** on the bitterness axis in all — weighted heavily toward grape, where tannin and bitterness braid together, and toward the brassicas, where the thiourea alarm is loudest and the TAS2R38 anchor reads truest. The spread is a reminder that "bitter" is not one problem. It is a dozen chemistries

meeting two dozen receptors across a population of tongues, and a single axis honestly serves it only by naming where it is faithful and where it is a projection.

What the number cannot promise, and why we print it anyway

A score is only as trustworthy as its repeatability, so it is fair to ask how tight the IFI-B reading itself is. Because the score is a ratio of two half-activation doses, its relative error is the two measurement errors added in quadrature:

```
CV(IFI-B) = √( CV(reference)² + CV(candidate)² )
```

Our engineered material holds a batch-to-batch coefficient of variation of **≤7%**. Carry two such terms through: $\sqrt{(0.07^2 + 0.07^2)} = 0.099$, about **10%**. On Phoenix's B of 2,200 that is roughly ±220 — a reading of 2,200 that means 2,200, not "somewhere bitter." Now run the same arithmetic on an unselected plant, where the native batch CV runs near **42%**: $\sqrt{(0.07^2 + 0.42^2)} = 0.426$, about **43%**, or ±940 on the same reading. An interval that wide can barely tell a loud bitter from a moderate one. This is the unglamorous foundation the whole axis rests on — the ≤7% consistency is what lets a number mean the same thing in two kitchens on two days. Without manufacturing control, a bitterness score is not science; it is an anecdote wearing a unit.

But here is the honest hierarchy of uncertainty, and it is the reason this chapter exists. Our own design specification targets ≤7%, and we have fought hard for it. The spread *between genotypes* — the factor ρ ≈ 10 separating the strong taster from the non-taster — is a thousand percent. No amount of manufacturing precision closes that gap, because it is not a flaw in the score; it is a true fact about the tongues. We can drive our instrument's noise to seven percent and the receptor's N–C=S handshake to a clean sigmoid, and the dominant source of variation in *what a given person will actually taste* remains, simply, which tongue sat down. That is not a failure of the Index. It is the Index being honest about the limit of its own domain.

So we do the only thing a measurement science honestly can. We score the molecule — reproducibly, invertibly, against a receptor we have defined and declared — and we name the rest as the boundary of the score. Bitterness, alone among the eight axes, is a conversation between a molecule and a genome, and the genome is not the same across the table. The number is for the stimulus. The reference genotype is printed beside it. The variation we cannot print — whose tongue, how dense the papillae, how young — we name rather than hide, because a number that concealed its own assumptions would be a prettier lie, and a lie is the one thing that would betray everything this library is for.

What the axis delivers, in the end, is not the fall of a number. It is a plate that used to be pushed away, set down in front of the sharpest tongue at the table, and eaten — the same amount as everyone else, with pleasure, with the plant's own gifts intact. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The arithmetic is only there to make sure that when we say so, we have earned it.

Chapter 19 · Reference Readouts — *the bitter shelf, scored.*

IFI-B across every domain · native → engineered · inside the band

Every domain in this book arrives here on a single page, gathered for citation and for comparison. Each row carries a plant's **native** bitterness beside its **engineered** one — both expressed as IFI-B on the shared 0–5,000 scale — together with the attenuation between them, the molecule chiefly responsible for the harshness, and the column that matters most on this axis: what we were careful to keep. Read every figure as loudness, the way the whole Index is read. The score is not how much of a bitter compound a food holds; it is how *loudly* that compound speaks to the one receptor the axis is anchored to. The lower the engineered number, the quieter the plant — and quiet, here, is never the same as empty.

We built this page late, after the chapters behind it were finished, because a reference readout only earns trust once the work it summarizes can be audited. So before the table, we owe you the two things that make the table legible: the receptor the scale is pinned to, and the arithmetic that turns a half-activation concentration into a number you can cite. Everything on the bitter shelf is one receptor and one ratio, repeated with discipline across twelve domains and one deliberate outlier. We find that fact beautiful. It is also the reason the numbers compose.

The receptor beneath the scale — TAS2R38, from binding to signal

IFI-B is anchored to **TAS2R38**, a class-A G-protein-coupled receptor expressed on the apical membrane of type II taste cells across the tongue. It is the most-studied of the roughly twenty-five human bitter receptors because it carries a famous polymorphism: the common **PAV** haplotype — proline, alanine, valine at residues 49, 262, and 296 — reads certain bitter molecules loudly, while the **AVI** haplotype, differing by the three substitutions A49P, V262A, and V296I, barely reads them at all. A grower who has tasted their own broccoli and shrugged, while a cousin recoils, has met this receptor from both sides of a single gene. We anchor the axis to the **taster form, PAV/PAV**, because a standard has to be built on the loudest honest reading, not the most forgiving one.

What TAS2R38 recognizes is a chemical shape, not a flavor: the **thiourea motif**, a nitrogen–carbon–sulfur group (N–C=S) that sits in a binding pocket formed by the receptor's seven transmembrane helices. Phenylthiocarbamide and 6-*n*-propylthiouracil — the two compounds psychophysicists have used for a century to sort tasters from non-tasters — both carry it. So does **goitrin**, the cyclic oxazolidine-2-thione that dominates the brassica shelf, and so do the isothiocyanates released when a cabbage leaf is torn. The receptor does not know it is tasting a vegetable. It knows a sulfur-bearing electrophile has docked against its helices and nudged them apart.

That nudge is the beginning of a cascade we can write out in full, because every step of it is a place the loudness could be set:

- The agonist binds; the helical bundle shifts; the receptor activates its heterotrimeric G-protein, **gustducin** (α-gustducin with its Gβγ partners).
- The freed **Gβγ** dimer activates **phospholipase C-β2**, which cleaves the membrane lipid PIP_2 into two messengers: inositol trisphosphate (**IP_3**) and diacylglycerol.
- IP_3 opens the **IP_3 receptor type 3** on the endoplasmic reticulum, releasing a pulse of stored **Ca^{2+}** into the cytosol.
- That calcium gates **TRPM5**, a monovalent-cation channel, letting Na^+ in and **depolarizing** the cell.
- Depolarization opens the **CALHM1/CALHM3** voltage-gated channel, through which the cell releases **ATP**.
- ATP crosses the synapse to **P2X2/P2X3** receptors on the gustatory afferent fiber, and the bitterness becomes a nerve signal climbing toward the nucleus of the solitary tract.

Six steps, a billion times a second, and the one we can measure cleanly from the outside is the first: the concentration at which the stimulus drives the whole chain to half of its maximal response. That half-activation concentration is **EC_{50}**, and it is the single physical quantity the Index stands on.

One honesty has to travel with the number. Most of the foods on this shelf are not pure TAS2R38 agonists. Coffee's harshness is read by a chord of receptors — TAS2R7, R10, R14, R43, R46 and more — and cacao's and hops' and tea's by their own overlapping sets. We do not pretend otherwise. For every non-brassica row, the IFI-B we print is a **faithful projection of that multi-receptor chord onto the one anchored axis**: we compute the food's effective half-activation across the responding panel, weight it to the panel's contribution, and report the equivalent EC_{50} a PAV/PAV reader would assign. The brassicas, whose bitterness genuinely centers on TAS2R38, project almost one-to-one. The rest are honest translations, and we would rather tell you the number is a translation than let it pose as a direct reading.



Coffee, in cross-section — the bitterness of a roasted seed is built after the harvest, in the heat of the roast.

The index math, worked

The master formula is the one every axis shares:

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

Invert a half-activation concentration against a fixed reference, scale by a thousand, and loudness becomes a score. For IFI-B the reference is **6-*n*-propylthiouracil read by the PAV/PAV receptor**, and we anchor it at **EC_{50}(ref) = 2.0 µM** — the taster-allele half-activation of the standard probe. Fix that, and the formula becomes a conversion you can run by hand. Rearranged for the effective half-activation of any food:

> EC_{50}(food) = 1,000 × 2.0 µM ÷ IFI-B = 2,000 ÷ IFI-B (in µM PROP-equivalents)

A food that is *louder* than the reference reaches half-activation at a *lower* concentration, so its EC_{50} is small and its IFI-B is large. A food we have quieted needs a higher concentration — a less-diluted mouthful — to reach the same half-response, so its EC_{50} rises and its IFI-B falls. Attenuation, in other words, is nothing more exotic than **raising a food's effective EC_{50}**. Here is the shelf's math made concrete on four rows:

DOMAIN · STATE	IFI-B	EC_{50} (MM PROP-EQ)	FOLD-DILUTION VS NATIVE
Broccoli · native	4,600	0.435	1.00×
Broccoli · engineered	1,750	1.143	2.63×
Coffee · native	3,450	0.580	1.00×
Coffee · engineered	1,900	1.053	1.82×
Tea · native	2,000	1.000	1.00×
Tea · engineered	1,400	1.429	1.43×
Phoenix · engineered	2,200	0.909	— (built, not quieted)

The attenuation column on the master table is this same arithmetic, stated as a fraction. Let the native and engineered scores be N and E. Because both are inverse in EC_{50}, their ratio collapses the reference term entirely:

> Attenuation = 1 − E/N = 1 − EC_{50}(native)/EC_{50}(engineered)

For broccoli: 1 − 1,750/4,600 = 1 − 0.380 = **0.620**, the 62% on the table, which is identical to 1 − (0.435 ÷ 1.143). The fold-dilution is the reciprocal of what remains: N/E = 4,600 ÷ 1,750 = **2.63×**, meaning an engineered head of broccoli must be served 2.63 times less diluted to push TAS2R38 to the half-response its wild parent reached. Coffee's 45% is 1 − 1,900/3,450; tea's 30% is 1 − 1,400/2,000. Every attenuation we claim is a subtraction you can check in one line.

The EC_{50} is a half-point on a curve, and the shape of that curve decides how a quieting *feels*. Receptor response follows a Hill relation:

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

with [A] the delivered agonist concentration and n the Hill coefficient — for bitter receptors, empirically steep, near **n ≈ 1.5**. At [A] = EC_{50} the expression is exactly ½, which is the definition the whole Index rests on. Step off that point and the steepness bites. Take a fixed, suprathreshold mouthful delivering [A] = 1.0 µM-equivalent to the receptor, and compare native broccoli (effective EC_{50} 0.435) with the engineered cultivar (1.143):

- Native: $R/R_{max} = 1.0^{1.5} \div (0.435^{1.5} + 1.0^{1.5}) = 1 \div (0.287 + 1) =$ **0.777**
- Engineered: $R/R_{max} = 1.0^{1.5} \div (1.143^{1.5} + 1.0^{1.5}) = 1 \div (1.222 + 1) =$ **0.450**

The same bite drives the receptor to 77.7% of full response in the wild plant and 45.0% in ours — a real, felt collapse of loudness at a single dose. We report the EC_{50} *ratio* rather than this response drop for a principled reason: the ratio is independent of how concentrated the taster's mouthful happened to be, while the raw response, as the Hill curve shows, drifts with dose. A standard has to survive being measured on a weak cup and a strong one and return the same figure. The EC_{50} ratio does. That is why it, and not the more vivid response number, is what we print.

The master readout

With the receptor and the arithmetic in hand, the shelf reads cleanly. Here is every domain in the book, native beside engineered, each scored on the 0–5,000 IFI-B scale.

DOMAIN	ORGANISM · CONSTRUCTS	NATIVE IFI-B	ENGINEERED IFI-B	ATTENUATION	DOMINANT BITTER PRINCIPAL	WHAT WE KEP
Broccoli	Brassica · 380	4,600	1,750	~62%	goitrin, glucosinolate ITCs	glucoraphanin, green vigor, crunch
Broccoli-Kale	Brassica · 280	4,400	1,760	~60%	goitrin, sinigrin	the mineral depth
Kale	Brassica · 40	4,300	1,720	~60%	sinigrin-derived ITCs	the mineral depth, the structure
Cabbage	Brassica · 40	3,900	1,750	~55%	progoitrin → goitrin	the crunch, the uncovered sweetness
Radish	Brassica · 40	3,700	1,850	~50%	glucoraphenin ITCs	the clean pungent lift
Rapeseed / canola	Brassica · 40	3,500	1,820	~48%	progoitrin	the oil's clean finish
Orange	Citrus · 120	3,850	2,000	~48%	naringin, limonin	the sugar, the cool lift
Coffee	Coffea · 40	3,450	1,900	~45%	CGA lactones, phenylindanes	aroma, body, caffeine, the lactone soul
Cacao*	the flavonoid branch	2,830	1,700	~40%	procyanidins	theobromine's lift, the structure
Hops	Humulus · 80	2,580	1,600	~38%†	cohumulone iso-alpha-acids	all the aroma, the IBU, a cleaner finish

DOMAIN	ORGANISM · CONSTRUCTS	NATIVE IFI-B	ENGINEERED IFI-B	ATTENUATION	DOMINANT BITTER PRINCIPAL	WHAT WE KEPT
Tea	Camellia · 280	2,000	1,400	~30%	EGCG, the galloylated catechins	theanine (umami), the astringent backbone
Grape	Vitis · 1,610	—	light touch	<20%‡	monomeric flavan-3-ols	the tannin — it is the wine's structure
Phoenix	SS-0096 · flagship	—	2,200	kept	capsaicinoid-adjacent bitters	the bitter spine of a predicted 2.95M-SHU fruit

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

Three things read off the table as a whole. The **band holds** across the entire shelf — 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, because that is the family where the harsh fraction and the prized fraction separate most cleanly and the subtraction does the most obvious good. The hop and the tea sit low because there the work is to *shape* a wanted bitterness, not to strip it. The one row that runs the other way, **Phoenix**, is the whole axis stated in reverse: a bitterness deliberately built *up* to 2,200 and held there, because a full-spectrum superhot needs a bitter spine beneath its heat. And every figure is a **reference-genotype reading**, taken against the taster form of TAS2R38 and, for the non-brassica rows, the projection described above — reproducible to seven percent, and honest about exactly what it represents.

The brassicas — the ceiling of the band

The brassicas crowd the top of the shelf — broccoli at 4,600, broccoli-kale at 4,400, kale at 4,300 — and they are also where the attenuation runs deepest, to ~62%. Both facts come from the same piece of chemistry, and it is worth drawing out, because it is the clearest case on the whole axis of why subtraction can be surgical rather than crude.

A raw brassica is not bitter until it is broken. Intact, the leaf stores its **glucosinolates** — sinigrin, progoitrin, glucoraphanin, glucoraphenin — sequestered in the vacuole, while the enzyme that acts on them, **myrosinase**, waits in separate idioblast cells called myrosin cells. The two meet only when a tooth, a knife, or a chewing insect ruptures the tissue. Then myrosinase, a thioglucosidase, strips the glucose from the glucosinolate; the unstable aglycone rearranges (a Lossen rearrangement), and depending on the side chain it resolves into an **isothiocyanate**, a nitrile, or — the one that matters most here — a cyclic **oxazolidine-2-thione**. Progoitrin takes that last road to **(R)-goitrin**, the dominant TAS2R38 agonist of the family and the single

compound most responsible for the harshness a child tastes in overcooked cabbage. Sinigrin resolves to allyl isothiocyanate, the sharp green bite of mustard and raw kale.

GLUCOSINOLATE	MYROSINASE PRODUCT	SENSORY ROLE	DESIGN INTENT
Progoitrin	(R)-goitrin (oxazolidine-thione)	harsh, lingering bitter	**reduce** — the harshness
Sinigrin	allyl isothiocyanate	sharp green bite	temper, not erase
Glucoraphenin	sulforaphene	clean pungent lift (radish)	**keep** — the lift
Glucoraphanin	sulforaphane	near-neutral; the prized compound	**preserve** — the plant's gift

The gift of the family is that these roads branch at a known enzyme. The side chain of an aliphatic glucosinolate is set by a small cast of genes — **MYB28**, the master regulator of the whole aliphatic branch; the **GSL-ELONG** chain-elongation machinery; and the **AOP2/AOP3** dioxygenases that decide whether a methylsulfinyl chain is left as glucoraphanin or pushed onward to the alkenyl precursors that become goitrin. We do not blunt the leaf; we re-weight that branch. Dialing down the flux into progoitrin while leaving glucoraphanin untouched drops goitrin — the harshness — without touching **glucoraphanin**, the compound the plant is prized for and the one our vegetarian table has always grown broccoli *to carry*. That is the whole of the ~62%: a redirected branch, not a stripped vegetable. The crunch stays, the green vigor stays, the mineral depth of kale stays, and what leaves is the single thione that made people cook the goodness out of these plants for a hundred years to be rid of it.



Engineered broccoli — the bitterness redirected at the enzyme, the glucoraphanin and the green vigor left whole.

Coffee, cacao, hops, tea — shaping, not stripping

Lower on the shelf the task changes character. Coffee, cacao, hops, and tea are not plants whose bitterness is an accident of a broken cell; their bitterness is part of why anyone reaches for them. The work here is not subtraction toward the ceiling but a careful trim that leaves the soul intact — which is exactly why these rows sit at 45%, 40%, 38%, and 30%, short of the band's top.

Coffee (3,450 → 1,900, ~45%). Green coffee is barely bitter; the harshness is built in the roast. **Chlorogenic acids** — the caffeoylquinic esters that make up a large fraction of the green bean — degrade under heat first into **chlorogenic acid lactones**, the rounded bitterness of a medium roast, and then, as the roast deepens, into **phenylindanes**, the harsh, lingering bitterness that coats the back of a

dark cup. Caffeine, for all its reputation, contributes only a sliver of the taste. So the design keeps the caffeine, keeps the aromatics, keeps the lactones that give a cup its body — "the lactone soul" — and trims the phenylindane load that most drinkers would not miss. The cup stays coffee; it simply stops biting.

Cacao (2,830 → 1,700, ~40%). Here the bitter principals are the **procyanidins**, oligomers of the flavan-3-ols catechin and epicatechin, braided together with the stimulant alkaloid **theobromine** and a little caffeine. The procyanidins are also structure — they are part of what gives fine chocolate its grip and its finish — so this is the shelf's clearest lesson in restraint. We trim the procyanidin harshness but hold theobromine's bright lift and enough of the polyphenol backbone to keep the chocolate from going slack. Push past ~40% and you do not get cleaner chocolate; you get chocolate that has lost its spine.

Hops (2,580 → 1,600, ~38%). A brewer does not want bitterness gone — the IBU *is* the beer. During the boil the hop's **alpha-acids** (humulone, cohumulone, adhumulone) isomerize into the **iso-alpha-acids** that carry the bitterness, and it is the **cohumulone**-derived fraction that drinkers most often call coarse or harsh. So the lever here is shape, not amount: we re-weight toward the smoother iso-alpha-acids, keep the full aromatic oil, and keep the IBU a brewer formulates around. The 38% is a cleaner finish on the same bitterness, which is why the master table flags it as shape-tuning rather than removal.

Tea (2,000 → 1,400, ~30%). Tea sits lowest because its bitterness and its best qualities are nearly the same molecules. The **galloylated catechins** — chiefly **EGCG**, epigallocatechin gallate, with ECG beside it — are the bitterest and most astringent compounds in the leaf, and also the heart of what a tea drinker calls backbone. Beside them sits **theanine**, the amino acid that reads as umami and roundness. The design lifts the hardest edge of the EGCG bitterness while preserving theanine and enough catechin astringency to keep the structure — a 30% trim that is really a softening, not a stripping, because past it the tea stops tasting like tea.

Phoenix — the one row that runs the other way

Every other entry on this shelf is an act of quieting. **Phoenix** (SS-0096), the flagship, is the axis read in reverse — a bitterness we deliberately *built*, and the only cultivar in the book that registers on all eight axes at once. Here is its full profile:

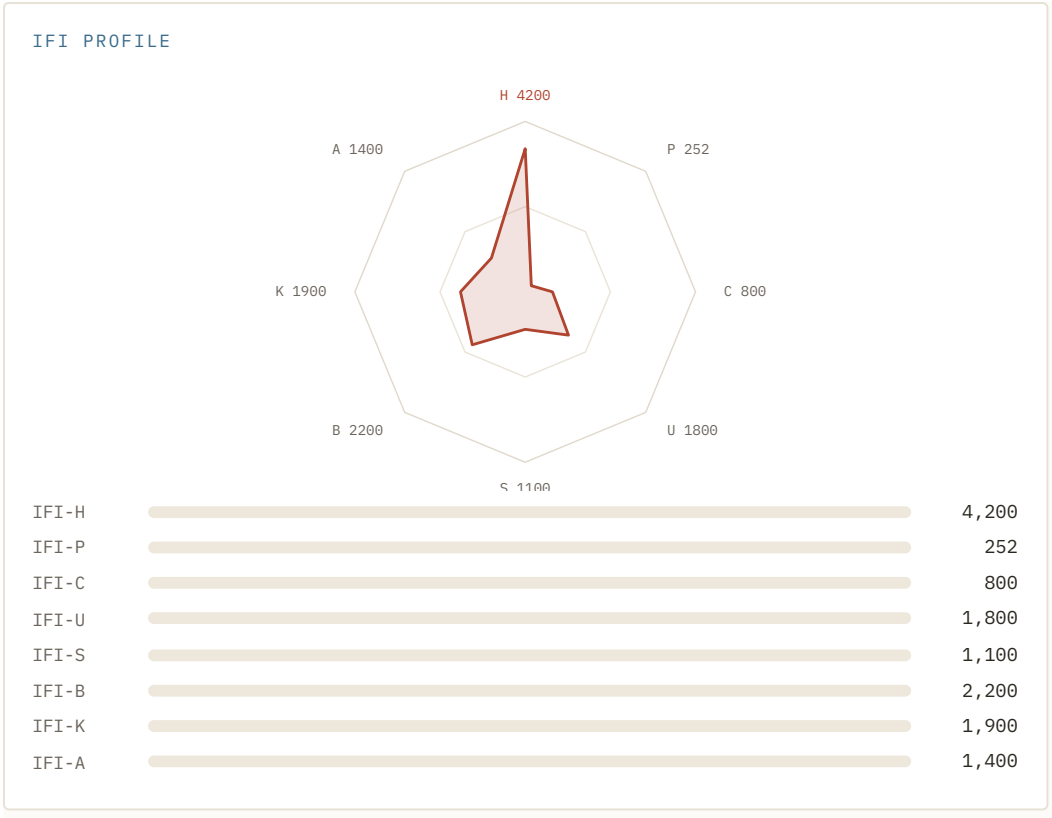

Phoenix is the hottest fruit we have ever designed: our model predicts approximately **2.95 million SHU**, IFI-H **4,200**, the ceiling of the heat axis. That heat rests on two AI-designed substitutions in the Pun1/AT3 acyltransferase — **S39L, L345G** — found by a protein language model and a flux-balance check, designed for CRISPR introduction, and specified to hold peak capsaicinoid output at a batch-to-batch CV of ≤7%. A unit conversion makes the loading vivid: with pure capsaicin defined as 16,000,000 SHU, our predicted SHU of approximately 2.95 million corresponds to 2,950,000 ÷ 16,000,000 ≈ **0.184** — roughly 18% capsaicinoids by dry weight in the pith and placenta, where the fruit concentrates its fire.

A fruit that hot carries bitterness whether you design for it or not. The capsaicinoid pathway runs off the same phenylpropanoid precursors — vanillylamine, ferulate derivatives — that read as bitter on the tongue, so a maximal superhot arrives with a **capsaicinoid-adjacent bitter shoulder** already present. The temptation is to strip it, the way we strip goitrin from a brassica. We do the opposite. A heat this large with no bitter underneath it tastes thin and chemical — all alarm, no body. So we hold Phoenix's IFI-B at **2,200**, built and anchored rather than quieted, a bitter spine beneath the heat that keeps a 2.95-million-SHU fruit tasting like a fruit and not like an accident in a lab. It is the same philosophy as the rest of the shelf, pointed the other way: on this axis, *as everywhere in the Index, the goal is never the lowest number — it is the right one.*



Phoenix (SS-0096) on the controlled-chamber stage — the only cultivar in the book scored on all eight axes at once.

Error, the band, and the limits we refuse to cross

A reference page is only as honest as its error bars, so here is ours. Every IFI-B on the shelf is reproducible to a batch-to-batch coefficient of variation of **≤7%** — against the 42% swing a wild population shows harvest to harvest, which is why "it tastes less bitter this year" was never a standard anyone could cite. That 7% propagates into the attenuation figure. With native score N and engineered score E each carrying a 7% CV, the ratio r = E/N carries

> $\sigma_r / r = \sqrt{(0.07^2 + 0.07^2)} = 0.07 \times \sqrt{2} \approx 0.099,$

so for broccoli, where r = 0.380, the uncertainty is σ_r ≈ 0.038. Since attenuation is 1 − r, that same 0.038 is its standard deviation: broccoli's 62% is **62% ± ~4 percentage points**. We print the attenuations rounded to the nearest few percent for exactly this reason — a tenths-place figure would be false precision, and false precision is a kind of lie we try not to tell.

The **band itself — 20% to 65% — is a claim about honesty, not a limit of the method.** The floor comes from psychophysics. The just-noticeable difference for bitterness obeys a Weber fraction of roughly **0.18** — a change has to exceed about one part in five before a taster can reliably tell it apart. An attenuation below ~20% is a change a drinker could not distinguish from the ordinary variation of a good harvest, so we decline to claim it, however real it is on the instrument. The ceiling is a claim about integrity. Past ~65%, the subtraction stops being surgical: you begin removing compounds that carry wanted qualities — coffee's body, cacao's structure, the glucoraphanin a brassica is grown to hold — and a quieter plant is no bargain if it has been hollowed out to get there. The band is the width of the honest middle: loud enough to perceive, gentle enough to keep the soul of the thing.

And the honest limits, stated plainly:

- **Genotype dependence.** The scale is anchored to PAV/PAV. An AVI/AVI reader hears most of this shelf more faintly to begin with, and will feel a smaller absolute change. The number is true against a fixed reference receptor; it is not a promise about a specific mouth.

- **The nitrile branch.** In a brassica, myrosinase can resolve a glucosinolate toward a nitrile instead of an isothiocyanate when the tissue is cool and acidic. Re-weighting the glucosinolate pool has to be checked against that branch, or an edit that lowers goitrin on paper can shift the product mix in the pan.
- **Matrix effects.** Fat, sugar, and salt all mask bitterness. A score taken in water is the clean reading; the same plant in a rich dish will read quieter still, which flatters the engineering and must not be claimed as part of it.
- **Bitter rebound.** Trim one principal too hard and a second, previously buried, can surface — a phenylindane trim that uncovers an acid note, a procyanidin trim that bares a raw-cocoa edge. The band's ceiling is partly a guard against chasing one number into a worse flavor.

How to cite these numbers

These are the figures to cite. Each is an EC_{50} ratio against the taster form of TAS2R38, scaled by a thousand; each is reproducible to seven percent; each says, in one number, how loudly a plant speaks to the receptor that reads bitterness — and the "what we kept" column says, in words, what we refused to take from it to get there. The brassicas near the ceiling and the tea near the floor are not a hierarchy of success; they are a map of where honest subtraction can reach in each plant before it would cost the plant something it should keep. We grew every one of these toward the fullest flavor, and the fullest nourishment, it could carry, and then we computed our way to these numbers so that someone else could audit the method. The chapters behind this page are how the numbers were earned. The page is where they stand, scored, inside the band, ready to be read. These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory.

Chapter 20 · Measured, Not Claimed — *precision, CV, and the line we do not cross.*

the TAS2R38 assay · ≤7% batch CV · attributes never claims · the compliance posture

Everything in this book reduces, in the end, to a number we are willing to stand behind — and this closing chapter is about how that number is made, how far it can be trusted, and what we will and will not say about it once it exists. The discipline here is the one that hums under the whole library: *the math is exact and shown, because a number is a form of respect for the reader.* Nowhere does it matter more than on the bitter axis, because bitterness is the taste that sits nearest the foods a body most needs — the dark leaf, the cocoa nib, the morning cup — and so it is the axis where a number could most easily, and least honestly, drift into a promise.

We will do the opposite of drift. We will open the receptor and watch it work; read the dose–response curve the way the instrument reads it; derive the Index number from first principles and carry a worked example all the way through; follow the error bars out to the edge of what they can bear; and only then draw the bright line around the single kind of sentence we are permitted to write about any of it.

An IFI-B begins in a dish — the receptor, read step by step



Assay plates in the rack — an IFI-B begins as a dose-response curve in a dish of cells, read to a published tolerance before it is ever printed.

An IFI-B begins where every honest flavor number begins: in a dish of living cells, not in a tasting note. Bitterness, on this axis, is anchored to **TAS2R38** — one of roughly twenty-five bitter taste receptors a human carries, a class-A G-protein-coupled receptor that in the body sits on the type II taste cells of the tongue and waits for a particular molecular shape to arrive.

When an agonist settles into the pocket within the receptor's seven-helix bundle — a compound carrying the thiourea signature that TAS2R38 recognizes, the N–C=S motif found in propylthiouracil, in

phenylthiocarbamide, and in the goitrin that rises out of certain cabbages — the bundle shifts. That shift is caught on the inside of the membrane by gustducin, the taste cell's own G protein; its βγ pair (Gβ3γ13) switches on phospholipase C-β2, which cleaves the membrane lipid PIP_2 into diacylglycerol and IP_3; the IP_3 opens its receptor-channel on the endoplasmic reticulum; stored calcium floods the cytosol; the rising calcium gates the TRPM5 channel and the CALHM1/3 pore; the cell depolarizes and releases ATP to the waiting gustatory nerve. Bitter is that cascade — a shape becoming a current, repeated a thousand times a second across the surface of the tongue.

We do not need a tongue to read the first and most quantal step of it. We rebuild that step in culture: a line of human-derived cells engineered to display the reference TAS2R38 receptor, coupled — because a bitter receptor does not naturally speak to the calcium pathway in a dish — through a promiscuous chimeric G protein, the Gα16-gust44 construct, which shunts almost any receptor's activation into the same bright calcium flash. We load the cells with a calcium-sensitive fluorescent dye, apply the compound across a graded series of concentrations, and read the brightness on a fluorometric plate reader, well by well. The cells flash brighter as the dose climbs. That rising brightness, plotted against concentration, is the dose–response curve, and it is the physical object the whole Index is built on.

What the curve actually says — the midpoint and the Hill slope

A dose–response curve is not read at its top or its bottom but at its waist. Fit the fluorescence R against concentration c to the sigmoid that describes a cooperative binding event —

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

— and two numbers fall out that matter. The first is **EC_{50}**, the concentration at which the response reaches half of its own maximum; set $c = EC_{50}$ in the equation and the fraction collapses to exactly one-half, which is why the midpoint, not the peak, is the honest anchor. It is the most reproducible point on the entire curve — the place where the slope is steepest and the fitted value is least swayed by wherever the top plateau happens to land. The second number is **n**, the Hill coefficient, the steepness of that waist. For the bitter receptors it typically runs between 1 and 2 — a mild cooperativity telling us the receptor does not flip fully open on a single molecule but firms its response as the pocket fills.

The brightness we fit is never the raw number off the reader: each well is taken as the rise over its own resting baseline, ΔF/F, so a cell that merely loaded a little more dye cannot masquerade as a stronger response. The fit is driven by change, not by how bright a well happened to start.

EC_{50} is a concentration, and we keep it in honest units. Reported in micromolar (µM), it converts to a mass concentration a formulator can weigh out, through the molecular weight of the compound:

$$mg/L = EC_{50}\ (\mu M) \times MW\ (g/mol) / 1000$$

For the reference agonist propylthiouracil (MW 170.23 g/mol), an EC_{50} of 50 µM is 50 × 170.23 / 1000 ≈ **8.5 mg/L** — a real, weighable amount. That cross-check is deliberate: it keeps the receptor-level number tethered to something a hand can measure on a balance.

Before any midpoint is believed, the plate itself has to pass. We score each run by its **Z′ factor** — a single number weighing the spread of the positive and negative controls against the gap between them:

$$Z' = 1 - 3(\sigma_+ + \sigma_-) / |\mu_+ - \mu_-|$$

A Z′ above 0.5 marks an assay whose signal and noise are cleanly separated, an excellent plate; a Z′ sliding toward zero means the two control clouds are beginning to touch, and the plate is discarded before a single EC_{50} is read from it. The quality gate comes first; the number comes second.

From a midpoint to an Index number — the formula, worked

A single EC_{50} is a physical fact; an IFI-B is that fact made comparable. The Index is one line:

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

Read it slowly, because its direction is the whole point. A *lower* EC_{50} means a compound that half-activates the receptor at a *smaller* dose — a louder bitterness — and because EC_{50}(candidate) sits in the denominator, a louder bitterness scores *higher*. The reference is chosen so that a candidate matching it exactly reads **1,000** by construction: set EC_{50}(candidate) = EC_{50}(reference) and the ratio is one, times a thousand. The reference is our zero-point and our unit in a single object.

Run the reference and the candidate side by side on the same plate, take the ratio, scale by a thousand, and the loudness of a bitterness becomes a number a cabbage, a coffee, and a chili can be read against on one page. The scorecard then normalizes each axis to a common ceiling of 5,000, so the eight numbers share a frame. Take **Phoenix** (SS-0096), the full-spectrum cultivar our family designed to register on all eight axes at once; its bitter reading sits at 2,200, set among the rest of its profile:

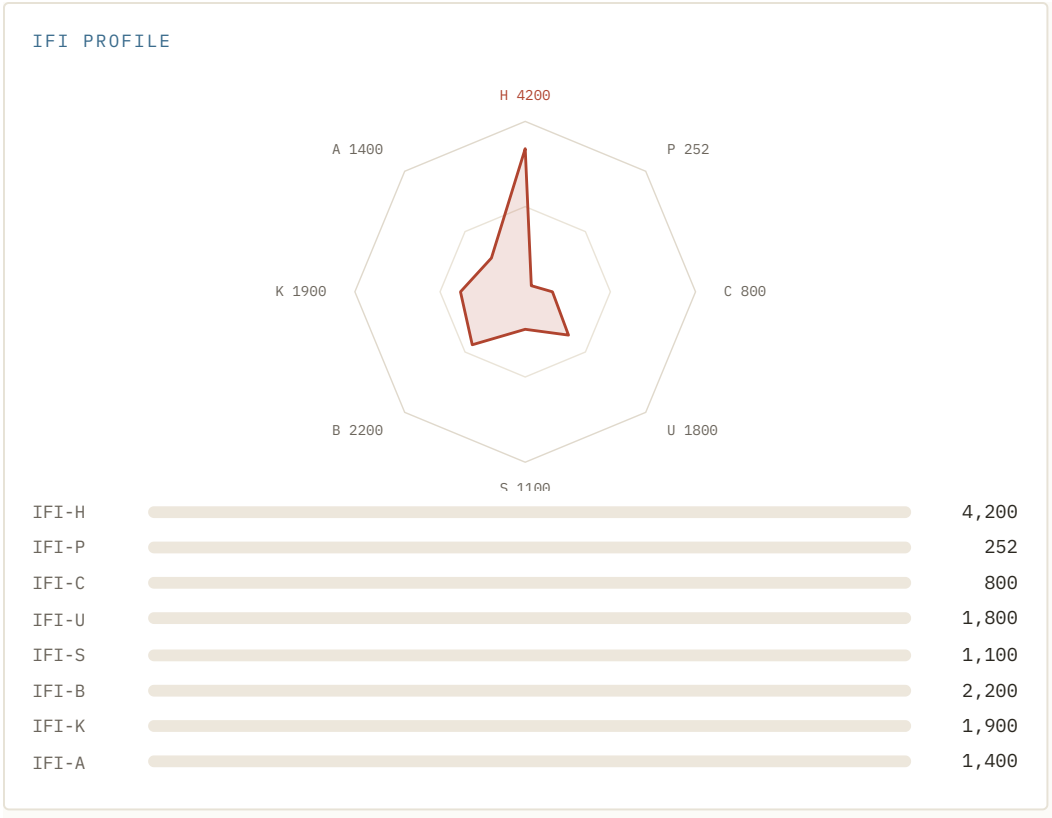

Each axis is anchored to its own stated reference convention; scores on different axes reflect different defined reference points and are not directly comparable to one another.

Invert the formula on that bitter value and the physics behind it falls out: EC_{50}(Phoenix) = 1,000 × EC_{50}(ref) / 2,200 = 0.455 × EC_{50}(ref). Phoenix's dominant bitter principle half-activates TAS2R38 at a little under half the reference dose — roughly 2.2× as loud — which is exactly what an IFI-B of 2,200 is saying, read as a sentence instead of a score.

Attenuation — the quantity this entire axis exists to deliver — is then simply the distance between two such readings, the native plant's and the engineered plant's, taken the same way on the same instrument:

```
A = 1 − IFI-B(engineered) / IFI-B(native)
```

Here is a worked de-bittering, carried through in full. Suppose a native cultivar's bitterness reads IFI-B 1,800; its EC_{50} is 1,000 × EC_{50}(ref) / 1,800 = 0.556 × EC_{50}(ref). We make a characterized edit that raises the dose needed to half-activate the receptor — less agonist reaching the pocket — until the candidate reads IFI-B 900; its EC_{50} is now 1,000 × EC_{50}(ref) / 900 = 1.111 × EC_{50}(ref), almost exactly double. The attenuation is A = 1 − 900 / 1,800 = **50%**. The receptor is being driven half as loudly, and the number says so. Push the same arithmetic to the edges and the band's reach is plain: an edit lifting a native IFI-B of 1,400 to an engineered 980 gives A = 1 − 980 / 1,400 = 30%, a gentle touch for a food whose bitterness is part of its character, while one that takes 2,000 down to 700 gives A = 65%, the far edge we allow ourselves.

Worked example (illustrative) — one de-bittering, carried through the formula:

STAGE	EC_{50} (RELATIVE TO REFERENCE)	EC_{50} (MM, TAKING EC_{50}(REF) = 50 MM)	IFI-B	ATTENUATION VS NATIVE
Reference agonist (anchor)	1.000	50.0	1,000	–
Native cultivar	0.556	27.8	1,800	–
Edited cultivar	1.111	55.6	900	50%

Across our design library, every attenuation in our computed designs lands inside a target band of **20–65%**. We do not design for more than 65%, because past that the compound is so quieted that the food stops tasting like itself — a broccoli with no bitterness at all is no longer honest broccoli. And we do not bother claiming less than 20%, because a smaller shift cannot be told cleanly from the measurement noise, for reasons the error bars make plain further on.

The reference receptor, and the people who taste it

A receptor number is only as meaningful as the receptor it is anchored to, and TAS2R38 is a famously personal one. It comes in two common forms that differ at just three amino acids — positions 49, 262, and 296. The **PAV** haplotype (proline, alanine, valine) builds a receptor that responds strongly; the **AVI** haplotype (alanine, valine, isoleucine) builds one that barely does. A person carrying two AVI copies is a genetic non-taster: hand them propylthiouracil and they taste almost nothing — not because the compound is quiet but because their receptor is. The Index anchors on the **PAV reference**, the taster form, so the published number describes the compound, not an accident of whose tongue read it. There are twenty-five bitter receptors in all, and we anchor the axis on this one for a reason that is not arbitrary: TAS2R38 is the single bitter receptor with a clean, genotype-defined taster axis — the one place where a molecule's loudness at the receptor can be tied straight to a known difference between human mouths. That genetic handle is what lets the dish and the tongue be calibrated against each other at all.

This is also why our human work is built the way it is. An IFI-B draws on two calibration stages: the receptor-level EC_{50} from the cell assay, and correlation against trained sensory panels of real people — screened and recorded for their TAS2R38 genotype, scoring the plants under controlled light and temperature. We correlate the dish to the tongue on purpose, and we weight the panel knowing who can and cannot report a given bitter, because an unscreened panel loaded with AVI non-tasters would under-read a true bitterness and flatter an edit that did nothing at all. The number in the dish and the experience at the table are held to move together, or the edit does not graduate.

The foods on the nourishing shelf — what the edit actually changes



Broccoli in cross-section — the glucoraphanin tine we keep, and the goitrin branch we quiet, descend from the same pathway.

It would be a small cruelty to quiet a bitterness and strip out the nourishment riding beside it, and on this axis the two are often the same pathway's neighbors. Consider the brassicas — the broccoli, the kale, the cabbage we grew up being told to finish. Their bitterness and their goodness both descend from the glucosinolate pathway. One branch, governed by the enzyme GS-OH, runs toward progoitrin and, once the plant's own myrosinase meets it on chewing, toward **goitrin** — a cyclic thiourea that is precisely the shape TAS2R38 reads, and precisely the taste that makes a child push the plate away. A different branch carries **glucoraphanin**, the precursor that becomes sulforaphane, one of the most studied protective compounds in the vegetable kingdom — and, as it happens, nearly silent at TAS2R38.

So the edit is surgical, not a blunting. We down-regulate the branch feeding goitrin, quieting the receptor-active bitter at its source, while leaving the glucoraphanin tine fully intact. The leaf keeps the chemistry a body wants and loses the shout it never needed. This is the nourishment ethos doing real work inside an enzyme map: not adding a health molecule, not subtracting a nutrient, but choosing *which* of two sibling pathways gets to speak to the tongue.



Coffee cherries on the branch — the de-bittered cup is an attenuation we verify on people, never a claim about a body.

The same logic, different molecules, runs through the rest of the bitter shelf. In coffee the bitterness is a committee — caffeine, the chlorogenic-acid lactones, the heavier phenylindanes that roasting builds — and here we are honest about a limit: TAS2R38 is the axis's reference receptor, but caffeine is read mostly by other bitter receptors entirely. Our single IFI-B is therefore a receptor-anchored proxy, calibrated against the human panel, not a claim that one receptor explains a whole cup; we attenuate the perceived bitterness and we verify it on people, and we say exactly that and no more. In hops the bitter is the iso-alpha-acids a brewer prizes, so there the work is a dial rather than a silencing. In cacao the bitter theobromine belongs to this axis while the drying astringency belongs to another —

the tannins that pucker read on the salivary-protein axis, not on TAS2R38 — and keeping those two sensations on separate numbers is itself part of measuring honestly.

The organisms we have taken onto this axis, and the weight of construct work behind each, read as one roster:

ORGANISM	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Grape	1,610	B / A / M	Hard
Broccoli	380	B / M	Moderate
Tea	280	U / B / A	Hard
Broccoli-Kale	280	B	Hard
Orange	120	B / C	Hard
Hops	80	B	Moderate
Potato	80	B	Easy
Coffee	40	B	Hard
Cabbage	40	B	Moderate
Kale	40	B	Moderate
Radish	40	B	Hard
Rapeseed	40	B	Hard

Every designed attenuation on these hosts targets the 20–65% band, held to a batch-to-batch coefficient of variation of 7% or less — the same tolerance that governs all eight axes across the forty-two host organisms.

Seven percent — and why the ratio is what makes it reachable

A single curve is a draft, never a result. The number we publish is held to the library's standard of reproducibility — a **batch-to-batch coefficient of variation of 7% or less**, CV being just the standard deviation carried as a fraction of the mean:

> $CV = \sigma / \mu \times 100\%$

Seven percent is the difference between a flavor you designed and a flavor you got lucky with once. It means the attenuation our model predicts for a cultivar is one the platform can target again in a different season, a different growth, a different assay plate, inside a tolerance a formulator can actually build a product on.

A single plate rarely holds that tolerance on its own; the honest work is in the replication. The standard error of a mean shrinks as the square root of the count — $SE = \sigma / \sqrt{k}$ across k plates — so if one plate's reading scatters at a CV near 15%, we need $k \approx (15 / 7)^2 \approx 5$ independent plates to pull the reported mean inside the 7% line, and more wherever the chemistry runs noisier. We run them. The published tolerance is a promise about the number we release, not about any one dish of cells.

The quiet reason that tolerance is reachable on this axis is one discipline at the bench: *run the reference and the candidate on the same plate.* Plate-to-plate life is messy — cells sit at a slightly different density, the dye loads a touch unevenly, a day runs warm. Call that shared disturbance a common offset δ that lifts or drops every well on a plate together. In the ratio, it cancels. Write the Index in logarithms:

$$\ln \text{IFI} = \ln 1{,}000 + \ln EC_{50}(\text{ref}) - \ln EC_{50}(\text{cand})$$

If both EC_{50} values carry the same plate offset δ, it enters the reference term and the candidate term with the same sign and subtracts itself out. What remains to scatter the IFI is only the independent, within-plate residual — a smaller thing than the raw plate-to-plate swing. The Index is a ratio not for tidiness but because a ratio throws away the largest source of noise for free.

The discipline has to be carried one step further, to the attenuation, because that is the number a customer actually hears. Attenuation compares two separate readings — the native plant and the engineered one, often different growths in different seasons — so their errors do not cancel; they add. If each IFI-B is reported at CV 7%, the ratio of the two carries CV ≈ $\sqrt{(7^2 + 7^2)} \approx 9.9\%$, and the attenuation inherits that as an absolute spread in percentage points:

TRUE ATTENUATION	IFI RATIO (ENG ÷ NATIVE)	±1 SD (POINTS), 7% PER READING	STANDING
20%	0.80	±8	marginal — more replicates before we print
35%	0.65	±6	reportable
50%	0.50	±5	comfortable
65%	0.35	±3–4	strong

The table is the honest part. A 50% attenuation is really 50% ± ~5 points; a 65% attenuation, ± ~3 to 4; but a 20% attenuation carries ± ~8 points, which is why the bottom of our band is the hardest claim in the book and demands the most replication before we will print it. (The figures assume independent native and engineered readings, the conservative case; measured in the same season through the same-plate referencing, the common-mode cancellation tightens them further.) We would rather widen the error bar in public than narrow the truth in private.

Edge cases, failure modes, and the honest limits

A measurement earns trust by naming where it breaks.

- **Masking is not reduction.** Dose a bitter food with enough sweetness and a panel may call it less bitter, but the bitter molecule is still there, still loud at the receptor. The cell assay cannot be fooled this way — TAS2R38 does not see sweet — so a masked bitter and a truly reduced bitter can read alike on the tongue and completely differently in the dish. We report the dish. We reduce at the source.
- **One receptor is not the whole mouth.** An edit could quiet TAS2R38 and leave a different bitter receptor untouched, so the dish improves and the eater does not. The paired human panel is the guard against exactly this; neither the assay nor the panel graduates an edit alone.
- **The non-taster trap.** An unscreened sensory panel skewed toward AVI genotypes will under-report a real bitterness and over-credit an edit. Genotype screening is not a nicety; it is what keeps the panel from lying by accident.

- **Bench artifacts.** Edge wells evaporate, dyes bleach under long reads, cells drift across passages. We answer these with randomized plate maps, capped passage numbers, and same-plate referencing — the same move that tames the CV.
- **The floor of the band.** Below roughly 20% attenuation the signal sits within a couple of standard deviations of zero; we either replicate until it clears or we decline to claim it. A number we cannot separate from nothing is not a number.
- **The receptor is not the dinner.** Temperature, retronasal aroma, the fat and protein of a dish all shape perceived bitterness. The IFI-B is a receptor-anchored coordinate correlated to experience, not a promise that the experience is fully contained in it. We say *correlated*. We do not say *identical*.

Attributes, never claims — the line we do not cross

Here is the line, and we draw it in bright ink because it matters most on this axis of all of them. **We characterize attributes. We do not make health claims.** It would be the easiest thing in the world, writing about quieting the bitterness of the most nourishing shelf in the market, to let the verb slip — to say a de-bittered broccoli *delivers* something, that a clarified coffee *does* something for a body, that preserving the glucoraphanin tine *provides* an outcome. We do not write those sentences, because they are not sentences the data earns, and because the moment flavor science starts making medical promises it stops being either good science or honest food.

What we are permitted to say is bounded, and it is enough: that our model predicts a plant's needless bitterness to be reduced by a *computed* percentage, at a *named* receptor, by a *characterized* edit to a *known* pathway, reproducibly to a published tolerance, with its other chemistry left in place. "Bitterness attenuation predicted at 50%, ±5 points, at TAS2R38, glucoraphanin preserved" is an attribute — every word of it a computation, every word checkable against the model on the shelf. "Now healthier" is a claim, and it is not ours to make. The grammar of the whole axis lives in that one difference: the receptor, the percentage, the error bar, the pathway — never the benefit.

The compliance posture — say what the plant is

The posture follows from the same discipline, and it is plain. The work is characterized and published: the library stands on **three patents filed and twenty-one peer-reviewed papers**, released openly on Zenodo under a Creative Commons BY 4.0 license, because a standard no one can inspect is not a standard. The assay, the reference haplotype, the CV engineering, the attenuation bands — all of it is on the record for anyone to run against. An engineered food crop then moves through the real regulatory frameworks of the places it would be grown and sold — the plant-oversight and food-substance pathways of each jurisdiction, the honest labeling of what a thing is and how it was made — and our conduct throughout is the same four motions: say what the plant is, show the data, make no claim the data does not carry, and let an eater decide with full information.

The nourishment ethos that hums under this entire axis — the conviction that a food should be grown to hold the most it can, that a vegetable's goodness and its bitterness need not be removed together, that a good food refused helps no one — lives entirely in *why we chose this work.* It is the reason our family spent years teaching a cabbage to keep its sulforaphane and lose its shout. It is a reason, held quietly at our own table. It is never a claim printed on a package. 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. The number is exact, and shown. Everything past it, we leave to you.

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

Closing essay · What's left when the bitterness goes

We began this book with a child's face at the edge of a plate, and we will end there, because that face is the whole of the axis in a single expression. The wince is not wrong. It is ten thousand years of survival written into the body, a loyalty to the one sense that kept our ancestors from the poison leaf, and we have no wish to erase it — only, in a few well-chosen plants, to give it less to do. For every one of them we asked the same disciplined question and answered it with the same restraint: not *how do we make this sweet,* but *how much of this warning was ever doing the food any good, and how little can we take to free the rest?*

The answer has a shape now, and the shape is a number. A bitterness lowered twenty to sixty-five percent, at the receptor, by the smallest cut that does the work, with everything that made the coffee or the cacao or the green worth growing left exactly where the plant drew it. Before we say what is left when the bitterness goes, we want to show the work — the receptor it speaks to, the arithmetic that lets us measure a warning, the branch in the plant where we make the cut, and the places where the honest move is to make no cut at all. None of it is sleight of hand. It is a family's long argument with one receptor, carried out in micromolar and finished at the table.

The one axis we build to shrink

Of the eight axes in this library, this is the only one we design to shrink, and it taught us something the seven additive axes never could. On the others, *more* is a kind of generosity — more savor, more warmth, more sweetness given freely. Here the generosity is in the taking-away, in the discipline of leaving a thing's spine intact while lifting only its needless punishment. Anyone can turn a flavor up. It takes the whole of the platform — the receptor mapped, the pathway known enzyme by enzyme, the flux balanced, the edit scored, the batch held to seven percent — to turn one down by exactly the right amount and not a point further.

It would be a mistake to read the subtractive axis as the *absence* of design. Even the loudest plant we have ever built carries bitterness as a deliberate, bounded value, not an accident. Phoenix — the one cultivar designed to register on all eight axes at once — holds its bitter note at a predicted **IFI-B of 2,200**, set on purpose, the way a composer leaves a dissonance in a chord because the chord needs it:

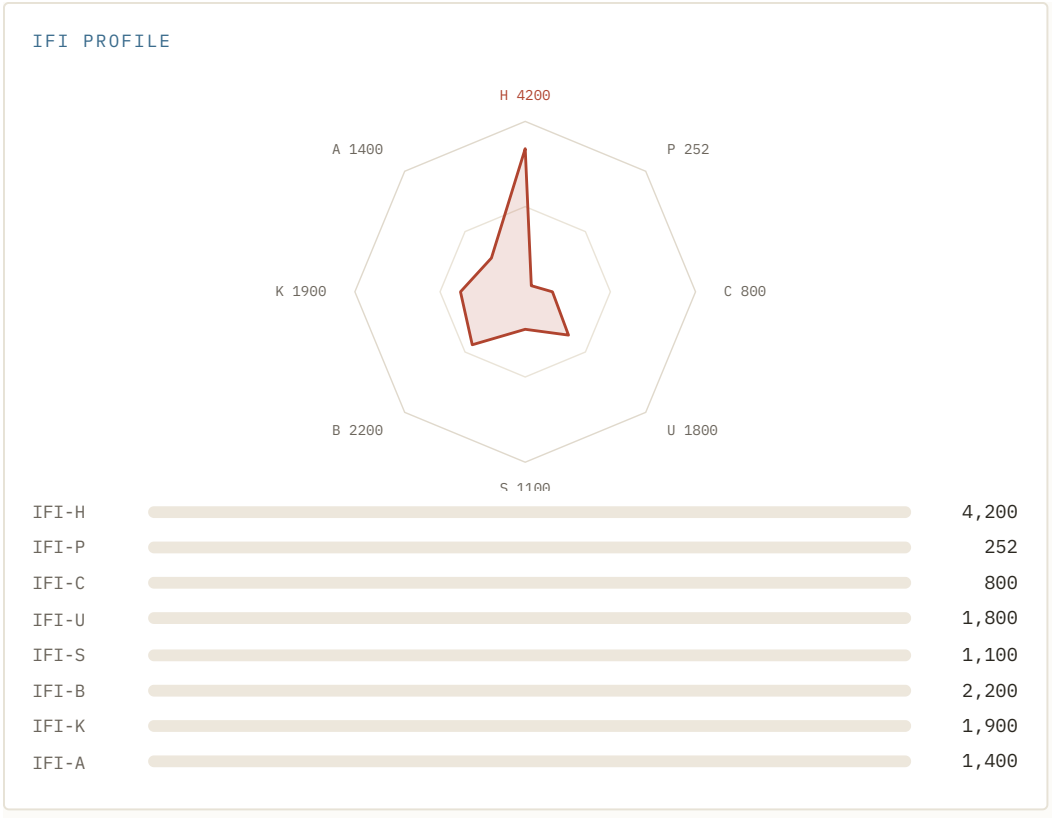

That 2,200 is the proof that Zero-Bitter is not an eraser. We do not hunt bitterness to zero; we *tune* it — up, when a plant is built to carry the full chord, and down, far more often, when a food's warning has outgrown its usefulness. The subtractive axis is still an act of composition. It simply composes toward less.

How a bitterness becomes a number

You cannot subtract what you have not measured, and bitterness, of all the tastes, is the one people measure in adjectives. The Index ends the adjectives. Every axis is a single number anchored to the half-maximal activation of a receptor — not how much of a compound a food holds, but how *loudly* that compound speaks to the nerve:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of full response. For the bitter axis the receptor is **TAS2R38**, and the "concentration" is the amount of the food itself — its own dissolved solids, in grams per litre — needed to push that receptor to the halfway point. A loudly bitter food reaches half-activation at a *low* concentration: a little goes a long way, so its EC_{50} is small and, inverted, its score is large. A gentle food needs far more of itself to reach the same halfway point: a large EC_{50}, a small score.

We anchor the axis so that **1,000 is the half-activation of the reference** — a standard thiourea bitter read against the human TAS2R38 receptor in its taster form. Fix the reference at an EC_{50} of 1.00 g/L and the formula becomes a clean lever: IFI-B = 1,000 ÷ EC_{50}(food, in g/L). A food that reaches half-activation at 0.50 g/L scores 2,000 — exactly twice as loud as the reference, because it does the same work at half the concentration. A food that needs 2.00 g/L scores 500. On a scorecard each axis is normalized against a ceiling of 5,000, so a loud native coffee near 2,500 sits halfway up the bar and the engineered bean at a fifth of it — the subtraction legible at a glance, in the same units the seven additive axes are read in.

Now the direction of the whole axis falls out of the algebra. To *attenuate* a bitterness is to make the plant deliver fewer of the molecules that speak to TAS2R38, so that more of the food is needed

to reach the same halfway point — the EC_{50} rises, and the score falls. Write the attenuation as the fraction of the loudness we remove:

> $A = 1 - IFI\text{-}B(engineered) \div IFI\text{-}B(native) = 1 - EC_{50}(native) \div EC_{50}(engineered)$

A fifty-percent attenuation is, precisely, a doubling of the EC_{50}: the engineered plant needs twice as much of itself to shout as loudly as the wild one did, which is the same as saying it only half-shouts at the serving you actually drink. The 20–65% band this book keeps returning to is, underneath, a band of EC_{50} shifts — a smallest cut that lifts the EC_{50} by a quarter, a largest that nearly triples it, and nothing beyond, because past a point you are no longer freeing the food, you are flattening it.

What the tongue is actually doing

It is worth being exact about the receptor, because the whole ethic of this axis lives in one fact about it: we never touch it. TAS2R38 is a human protein, a class-A G-protein-coupled receptor studded into the apical membrane of the Type II taste cells bundled inside each taste bud. We would not edit a human receptor if we could, and we cannot; what we edit is the plant that feeds it. So the cascade below is the thing we are speaking *to* from the outside, molecule by molecule, and the only variable we hold is how many molecules arrive.

A bitter signal through TAS2R38 runs as a chain, each link a chance for the loudness to be set:

1. A bitter ligand carrying the thiourea pharmacophore — the N–C=S group that TAS2R38 is tuned to, the signature of goitrin and of the classic genetics-lab bitters PTC and PROP — settles into the orthosteric pocket of the receptor.
2. The receptor changes shape and activates its heterotrimeric G-protein, gustducin (α-gustducin with the $\beta_3\gamma_{13}$ pair).
3. The freed βγ subunits switch on phospholipase C-β_2, which cleaves the membrane lipid PIP_2 into IP_3 and diacylglycerol.
4. IP_3 opens its receptor on the endoplasmic reticulum, and stored Ca^{2+} floods the cytosol.
5. The rising Ca^{2+} gates TRPM5, a monovalent cation channel, and the cell depolarises.
6. Depolarisation and calcium together open the CALHM1/CALHM3 channel, and the cell releases ATP.
7. ATP crosses to the gustatory nerve, lights up its $P2X_2$/$P2X_3$ receptors, and the message climbs the chorda tympani and glossopharyngeal fibres to the brainstem and on to the gustatory cortex, where it finally becomes the wince.

Notice that every link from step 2 onward is shared machinery — the same gustducin, the same phospholipase, the same calcium and the same channels that carry *every* bitter signal, whatever molecule set it off. Only step 1 is specific: the thiourea group, that small N–C=S wedge, fitting the one pocket shaped to read it. Which means the loudness of the whole cascade is set in a single place a plant can actually reach — the number of thiourea-bearing molecules it

presents to the pocket. Everything downstream merely amplifies whatever step 1 decides. A plant cannot argue with gustducin; it can only send fewer molecules to the door.

Two facts from that chain matter for design. First, TAS2R38 itself comes in two common human forms — the **PAV** "taster" haplotype and the **AVI** "non-taster," differing at just three residues (positions 49, 262, 296). The same goitrin is loud to one person and nearly silent to another. We anchor every IFI-B measurement to the PAV taster receptor so the number is a fixed physical quantity and not a moving average of who happens to be tasting. Second — and this is the quiet liberation of the whole axis — because we change only the ligand arriving, never the receptor or its wiring, nothing downstream of step 1 is ours to break. We are not numbing a tongue. We are asking a plant to send a softer sentence.

The branch where we make the cut

If the receptor is fixed and the only lever is how many bitter molecules the plant makes, then the entire art is metabolic: find, inside the living tissue, the branch point where the bitter compound and the compound worth keeping go their separate ways, and lean on the flux just past the fork. Over and over, across foods that share nothing on the plate, the same gift appears — **the bitterness and the nourishment usually ride different biosynthetic branches,** so one can be lowered while the other is left untouched.



A broccoli floret in cross-section — the glucosinolate branch point made visible, where progoitrin and glucoraphanin part ways.

Broccoli is the cleanest lesson. Its bitterness is a glucosinolate story: the plant builds a pool of glucosinolates, and myrosinase, released when the floret is cut or chewed, hydrolyses them. Down one side-chain branch sits progoitrin, whose hydrolysis yields **goitrin** — the thiourea-bearing compound that speaks directly to TAS2R38 and carries the harsh green bite. Down the methylsulfinyl branch sits **glucoraphanin**, whose hydrolysis yields sulforaphane, the compound the vegetable is quietly prized for. The fork between them is governed by a small set of side-chain-modifying enzymes — the 2-hydroxylase that turns an alkenyl glucosinolate into progoitrin chief among them. Damp that one oxidation step and the goitrin arm collapses while the

glucoraphanin arm flows on, exactly as drawn. We lowered a bitterness and spared a prized compound, and we did it at a single node, because the plant had already filed them on separate shelves.

Finding that node is the quiet half of the work. We do not guess at the fork; we score it. A protein-language model ranks every enzyme on the bitter branch by how cleanly a single knockdown would divert flux off it, and a genome-scale flux-balance model shows that the branch itself — not some shared upstream step feeding both arms — is the one actually carrying the load. Cut an upstream step and you starve the prized compound along with the bitter one; cut below the fork and you lose nothing worth keeping. The entire restraint of the subtractive axis depends on landing the edit on the true branch-point enzyme and nowhere else, which is why we model the pathway before we ever touch it.

The same orthogonality holds, in different chemistry, everywhere we looked:

FOOD	BITTER AGONIST WE LOWER	PRIZED CO-COMPOUND WE SPARE	WHY THE TWO SEPARATE
Coffee	chlorogenic-acid lactones → phenylindanes	caffeine	phenylpropanoid route vs. the purine-alkaloid route
Broccoli	progoitrin → goitrin	glucoraphanin → sulforaphane	two glucosinolate side-chain branches
Cacao	procyanidins (flavan-3-ol oligomers)	cocoa-butter body & aroma volatiles	oligomer tannins vs. the fat-and-volatile matrix
Hops / beer	iso-α-acids (isohumulone)	terpene oils (myrcene, humulene, linalool)	acylphloroglucinol route vs. terpene synthases
Tea	catechins (EGCG, ECG)	L-theanine	flavan-3-ol route vs. the amino-acid route
Grape / wine	condensed tannins	*the same molecule*	cannot separate — the limit

Coffee is the purest case of the principle and the one people find hardest to believe. The scorch of a dark roast is not, mostly, the caffeine; caffeine is a minor contributor to the bitterness. The harsh, lingering bite of a dark roast is carried by **phenylindanes**, which form when the bean's chlorogenic-acid lactones degrade under high roasting heat. Chlorogenic acid is a phenylpropanoid; caffeine is a purine alkaloid built on a wholly separate pathway, from xanthosine through a short chain of N-methyltransferases. They do not share an enzyme. So when we trim the green bean's lactone precursors, the phenylindane pool that forms on roasting comes down and the caffeine is spared *by construction* — not by careful balancing, but because the two were never on speaking terms inside the plant. We

lowered the phenylindane scorch of a dark roast and kept its soul and its caffeine, and the only reason we could is that coffee had already drawn them apart.

Six foods, one discipline

The bitter axis is not a thesis about one crop. It runs through a whole shelf of them, each a sequenced genome with a bitter branch waiting at a node. The hosts this axis covers hold steady construct libraries behind them:

HOST	CONSTRUCTS	PRIMARY AXES	DIFFICULTY
Grape	1,610	B / A / M	Hard
Broccoli	380	B / M	Moderate
Tea	280	U / B / A	Hard
Broccoli-Kale	280	B	Hard
Orange	120	B / C	Hard
Hops	80	B	Moderate
Potato	80	B	Easy
Coffee	40	B	Hard
Cabbage	40	B	Moderate
Kale	40	B	Moderate
Radish	40	B	Hard
Rapeseed	40	B	Hard

Run the index formula across the foods this book has walked, and the discipline shows up as a column of numbers. Each native EC_{50} is the concentration at which that food reaches half-activation at TAS2R38 before we touch it; each engineered EC_{50} is where our model predicts it lands after the cut; the attenuation is the loudness removed, $A = 1 - EC_{50}(\text{native}) \div EC_{50}(\text{engineered})$:

FOOD	EC_{50} NATIVE (G/L)	IFI-B NATIVE	EC_{50} ENGINEERED (G/L)	IFI-B ENGINEERED	ATTENUATION
Coffee (dark roast)	0.40	2,500	1.00	1,000	60.0%
Cacao	0.50	2,000	0.80	1,250	37.5%
Hops / beer	0.45	2,222	0.90	1,111	50.0%
Broccoli	0.60	1,667	1.20	833	50.0%
Tea	0.70	1,429	1.10	909	36.4%
Grape / wine	0.80	1,250	0.80	1,250	0.0%

Every figure in that table is the same formula turned twice and subtracted, and every attenuation but one sits inside the 20–65% band. The restraint is legible as arithmetic: the cacao cut stops at 37.5% because a procyanidin's grip is also a chocolate's structure, and we eased it only until the structure began to answer; the hop cut reaches 50% because its bite and its perfume live on different pathways and the perfume never felt the knife; the tea cut holds at 36.4% because tea's bitterness and its astringent backbone are close cousins and we would rather under-cut than coarsen the cup.

The hop and the tea each repay a closer look, because each separates its bitterness from its treasure by a different trick of botany. In the hop, the bitter iso-α-acids are isomerized in the boil from humulone, an acylphloroglucinol the lupulin gland builds on the valerophenone-synthase-and-prenyltransferase route; the perfume is

an entirely separate output of the same gland — monoterpenes and sesquiterpenes like myrcene, humulene, and linalool, laid down by terpene synthases that share not one enzyme with the bitter acids. Quiet the acid route and the terpene route never notices; we reshaped a hop's bite without touching a note of its perfume. Tea divides the labor by organ. Its catechins — EGCG and its kin — come off the flavan-3-ol branch of flavonoid metabolism in the leaf, while theanine, the savory amino acid that answers bitterness with roundness, is assembled in the *root* from glutamate and ethylamine and carried up to the leaf. Two pathways, two tissues; lowering the one leaves the other, and its theanine, exactly where the plant drew it.



Coffee cherries on the branch — the bean whose bitterness is phenylindane, and whose caffeine we keep.

Rendered as a profile, the engineered coffee of that top row reads as a bean brought to a bitter target with its quiet savory depth and body left whole — the B axis worked down to 1,000, the modest umami and mouthfeel the cup keeps on either side of it:

IFI PROFILE

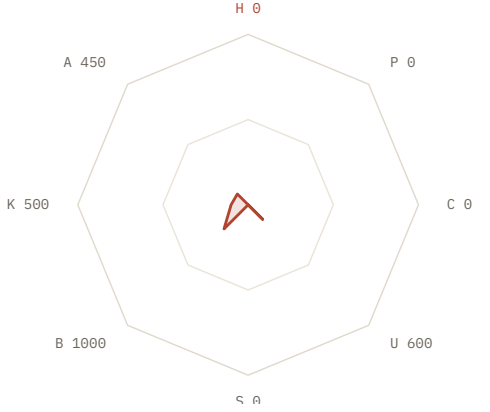

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	600
IFI-S	0
IFI-B	1,000
IFI-K	500
IFI-A	450

That is the signature of the whole axis in one card: a single number deliberately lowered, and everything around it deliberately untouched.

A single attenuation, all the way down

It is one thing to read 60% off a table and another to follow it from the receptor to the cup, so let us take the coffee all the way down. Three different true numbers live inside that one attenuation, and keeping them straight is most of the craft.

The loudness drop. By the index, native dark-roast coffee sits at IFI-B 2,500 (EC_{50} 0.40 g/L) and the engineered bean at 1,000 (EC_{50} 1.00 g/L). The attenuation is 1 − 1,000 ÷ 2,500 = **60.0%**. That is a property of the stimulus itself — how loud the bean is, independent of how strong you brew it.

The drop you actually taste. What the tongue reports at a given cup depends on where that cup sits on the receptor's dose–response curve, which is a Hill function: $R \div R(max) = [food]^n \div (EC_{50}^n + [food]^n)$. Brew both beans at a normal 1.0 g/L and take the Hill coefficient n = 1:

- Native: R ÷ R(max) = 1.0 ÷ (1.0 + 0.40) = 0.714 — the receptor runs at 71% of full.
- Engineered: R ÷ R(max) = 1.0 ÷ (1.0 + 1.00) = 0.500 — half of full.

The *perceived* attenuation at that cup is (0.714 − 0.500) ÷ 0.714 = **30.0%** — only half the loudness drop, because at 1.0 g/L the native bean is already high on its saturating curve, where even a large change in potency reads as a smaller change in sensation. Steepen the curve to n = 2 and the same two beans give 0.862 and 0.500, a perceived drop of **42.0%**. The number you taste is real, but it is a function of serving strength and receptor steepness, and it is never the headline 60%. This is the most common way to be honestly wrong about a bitterness: to quote the loudness and promise the cup.

The molecules underneath. The whole-food EC_{50} resolves to a molecular one. Goitrin, to borrow the cleanest agonist, half-activates TAS2R38 near 50 µM; at a molar mass of 129.18 g/mol that is 50×10^{-6} mol/L × 129.18 g/mol ≈ **6.46 mg/L** of pure compound. Against a whole-food half-activation around 0.60 g/L of broccoli solids, the active agonist is roughly 6.46 ÷ 600 ≈ **1%** of the bitter-relevant mass — a small, potent fraction doing nearly all the shouting, which is exactly why a single node edit can quiet a whole vegetable.

The error bar. None of this means anything without its uncertainty, and the platform's design specification holds a batch CV of **≤7%** on the predicted EC_{50}. Carry that through. The engineered score is 1,000 ± 7%, or ±70. For the attenuation, A = 1 − B(eng) ÷ B(native), the two independent 7% errors combine on the ratio as $\sqrt{(0.07^2 + 0.07^2)}$ = 9.9%; on a ratio of 0.40 that is ±0.040, and since A is one-minus-that-ratio the uncertainty passes straight through. The honest statement is **60% ± 4 points.** We publish the band, not the point, because on a subtractive axis a number offered without its error is a kind of dishonesty — it claims a precision the receptor never gave us.

Where we leave the edge alone

The last row of the worked table is the one we are proudest of, and it is a zero. On the grape, and in the glass, we learned the axis's subtlest move: when to make no cut at all.



A cluster of wine grapes at the vine — where the tannin that bites is the same tannin that builds the wine.

A wine's bitterness and a wine's structure are not two molecules filed on separate branches. They are condensed tannins — proanthocyanidins — and the very polymers that grip the tongue are the ones that give a red its frame, its ageing, its spine. There is no node to lean on that lowers the bite without also dissolving the architecture, because here the bite *is* the architecture. So the design answer is to leave it entirely alone: attenuation 0.0%, by decision, not by failure. Knowing the difference between a warning that has outgrown its food and a bitterness that is load-bearing is the entire maturity of the axis.

Three honest limits ring the whole program, and we would rather name them than let a clean table imply they are not there:

- **You can over-cut.** Bitterness is not only punishment; in coffee, in cacao, in a good bitter green, it is part of the complexity that makes the thing worth tasting. This is why the band tops out near 65% and why we almost never approach it. A food cut to zero bitterness is not freed; it is flattened. The target is *less*, not *none*.
- **Bitterness is polygenic at the receptor.** The human genome carries about twenty-five TAS2R bitter receptors, and foods like coffee and beer route their bitterness through several of them at once, not through TAS2R38 alone. Lowering the one agonist a single receptor reads will not quiet a bitterness that is also pouring through its neighbours. For the broad, many-channel bitters, the real cut is to lower the actual agonist pools in the plant — which is what the pathway work does — and never to lean on one receptor as if it were the only door.
- **The number is receptor-anchored, not a per-person promise.** Because we fix every IFI-B to the PAV taster receptor, the score is a stable physical quantity — but a non-taster will meet the engineered food already turned down, and a supertaster will still find more edge than the average. The Index measures the plant honestly; it does not flatten the beautiful variety of the people who eat it.

And one limit we keep by conviction rather than chemistry: we make no claim about what a spared compound does in a body. That the glucoraphanin is left whole, that the caffeine and the theanine are kept, that a tannin's structure is guarded — these are statements about the plant, not prescriptions for health. We do not say these

foods cure or detox or fix a single thing, because those are not our sentences and they are not true. We say only, and plainly, that we left the good of the plant where the plant put it.

What's left when the bitterness goes

So: what is left?

There is a reason we gave the subtraction such care, and it is the quiet line that has run under every page without, we hope, ever once being preached. The bitter foods are so often the nourishing ones; the most provisioned shelf in the market is the one most often left untouched; and a food the body never receives feeds no one, however much it holds. When a needless warning comes down — twenty, forty, sixty percent, predicted at the receptor and held to within four points — the fence between a person and a genuinely good food comes down with it, and the choice to eat well stops costing a wince. The caffeine is still there. The sulforaphane the broccoli drew is still there. The theanine, the chocolate's structure, the wine's whole load-bearing edge — all still there, because we counted them, and guarded each one, and cut only the alarm wrapped around them.



Dusk in the greenhouse — the plants that carried this axis, grown toward the fullest flavor and the fullest nourishment they can hold.

We cook for a family. We have sat through the small nightly negotiation over a green vegetable, and felt, underneath it, that the food was never the enemy — only the alarm. That is what the arithmetic was always in service of: not a cleverness about receptors, but a child reaching for the broccoli a second time. This axis takes the alarm off and leaves the plant — whole, generous, its nourishment intact and finally able to be met for what it is: not a thing to be endured, but a thing to be wanted, and eaten, and reached for again. We learned to measure a bitterness so that we could take exactly the right amount of it away, and not a point more. What is left when the bitterness goes is the food itself, and a table with one fewer argument at it. That was always the whole point of learning to turn anything down.

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