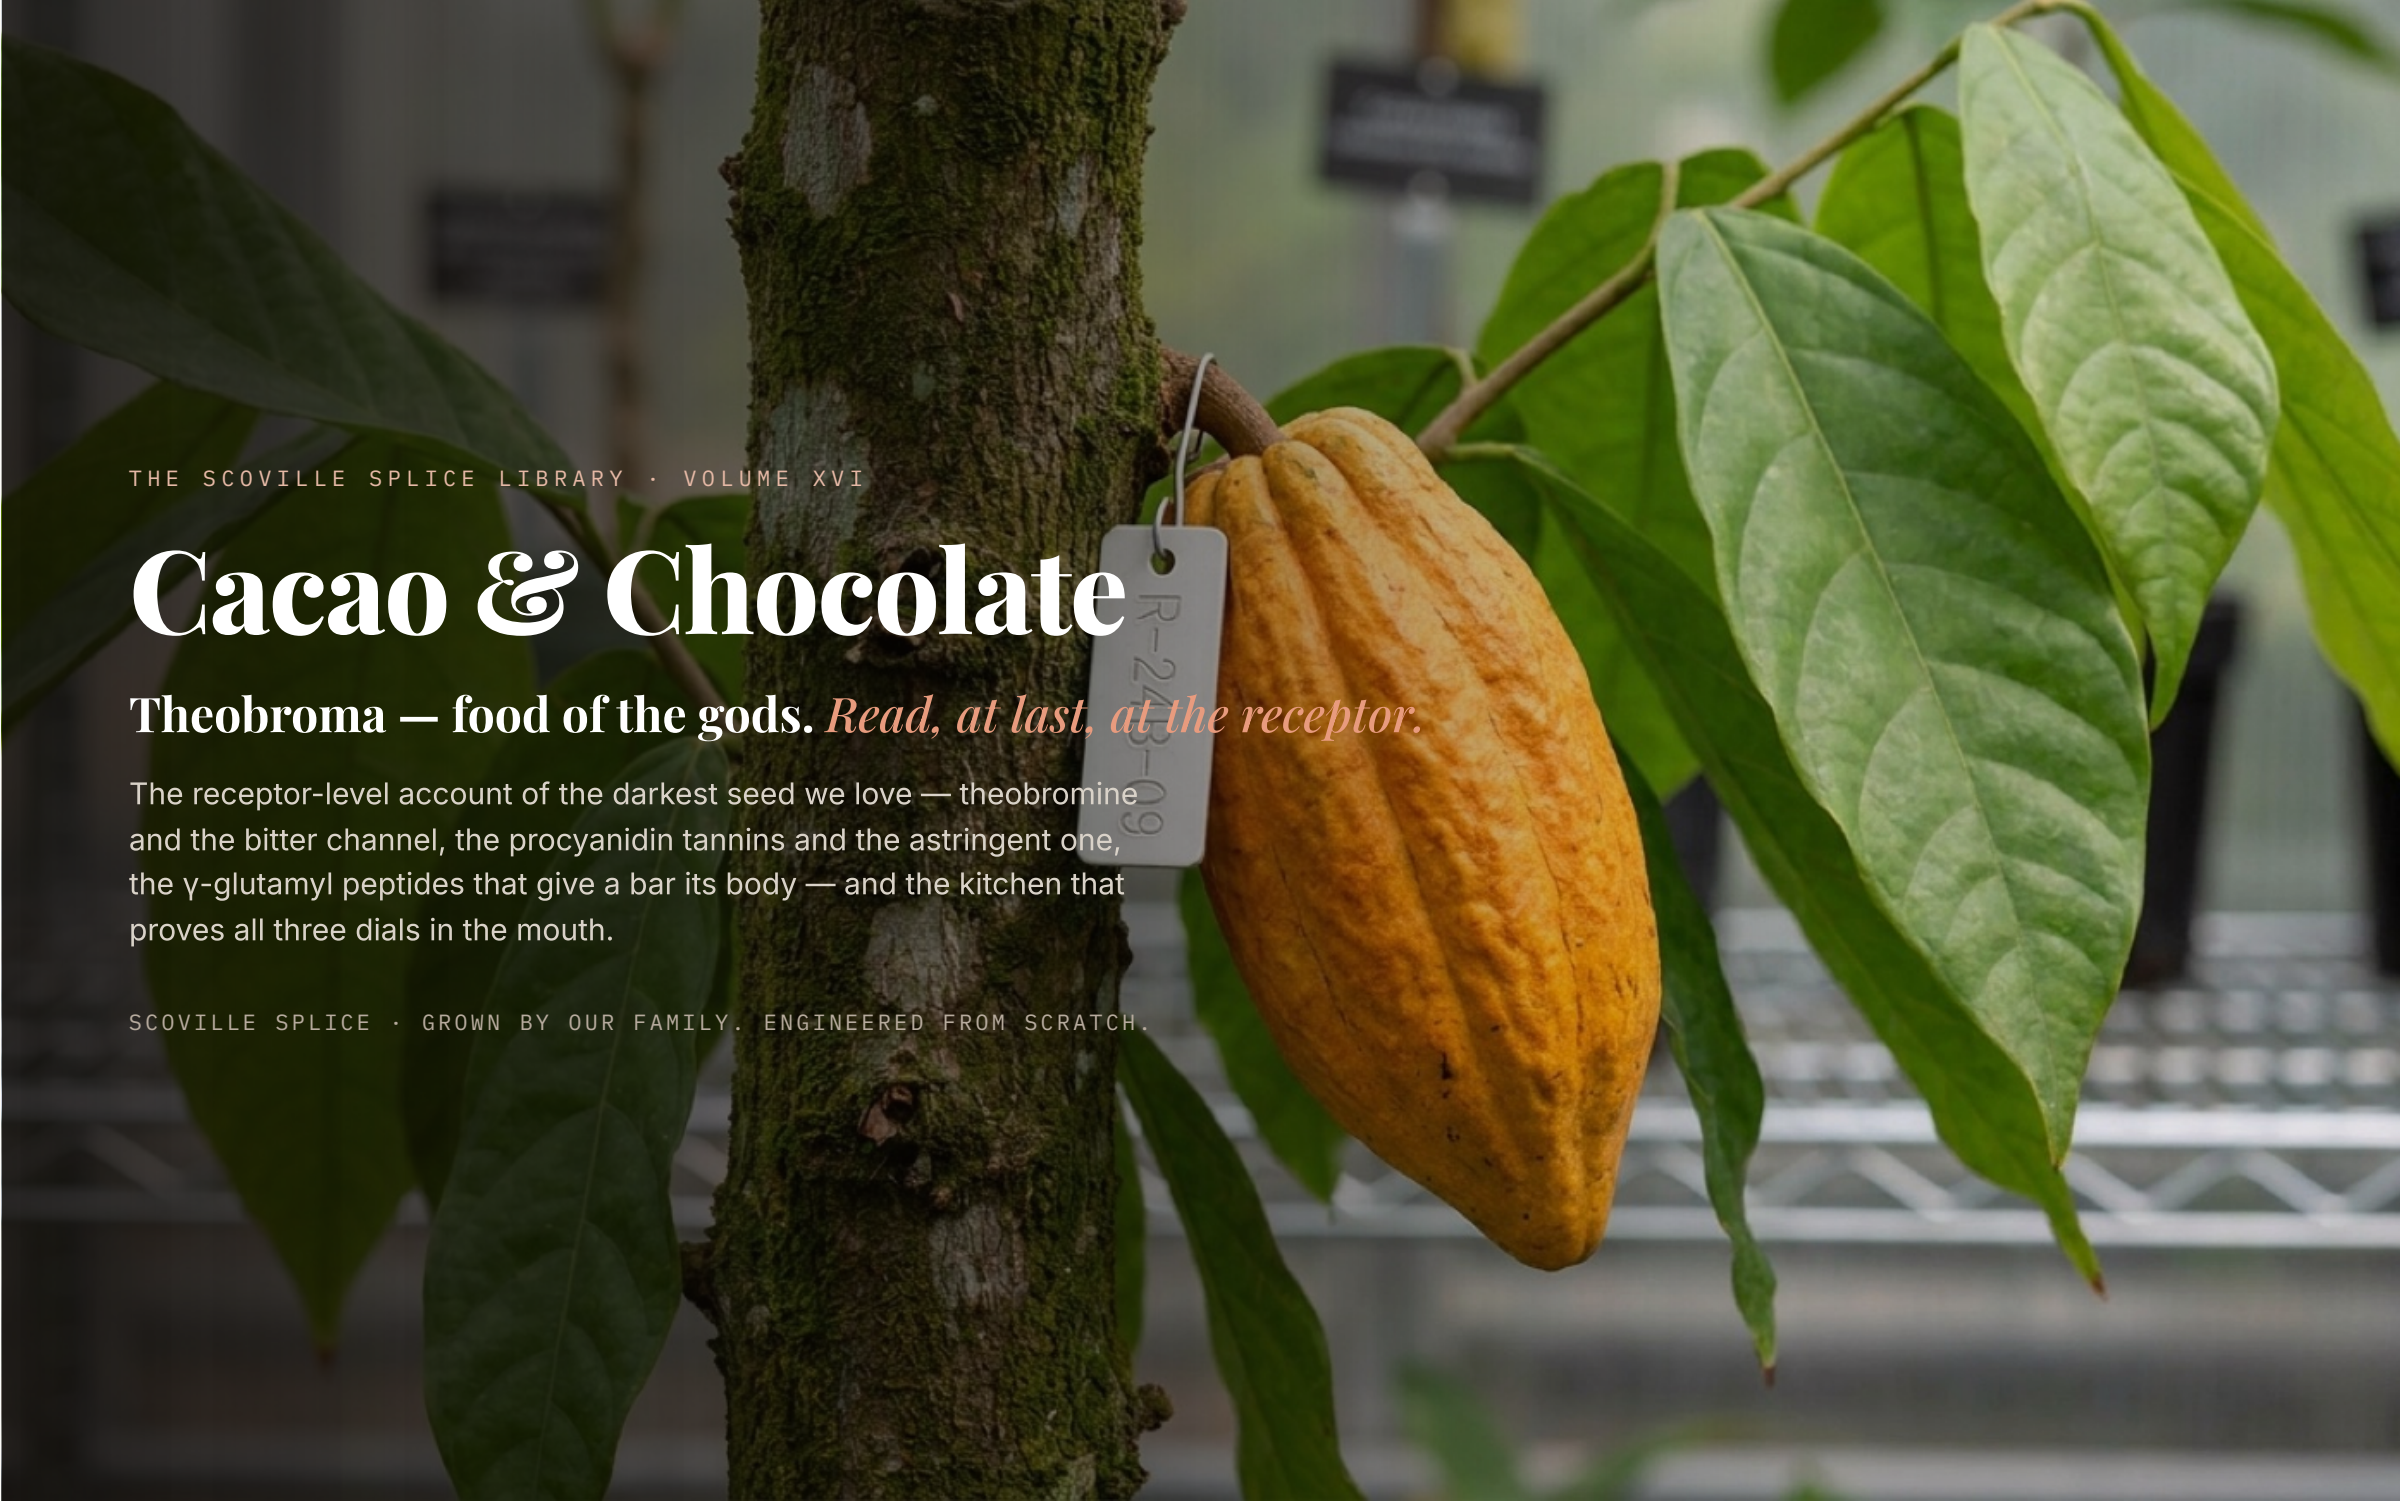

THE SCOVILLE SPLICE LIBRARY · VOLUME XVI

Cacao & Chocolate

Theobroma — food of the gods. *Read, at last, at the receptor.*

The receptor-level account of the darkest seed we love — theobromine and the bitter channel, the procyanidin tannins and the astringent one, the γ-glutamyl peptides that give a bar its body — and the kitchen that proves all three dials in the mouth.

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

Contents

01 Opening essay · The seed we are still learning to speak

02 Chapter 1 · Theobroma — the bitterness we chose to love.

03 Chapter 2 · Three dials on one seed — B, K, A.

04 Chapter 3 · The number under the dark — EC_{50}, IFI, and how to read a bar.

05 Chapter 4 · The methylxanthine problem — theobromine, caffeine, and the bitter repertoire.

06 Chapter 5 · Attenuation, not amputation — 20–65% off, and nothing hollowed.

07 Chapter 6 · The design, written — ESM2, flux balance, CRISPR.

08 Chapter 7 · The cross-linked mouth — procyanidins and the PRP cascade.

09 Chapter 8 · Tuning the tannin — ANR, LAR, and the degree of polymerization.

10 Chapter 9 · Kokumi — the dimension you feel but cannot name.

11 Chapter 10 · Fermentation, the first engineer — what the microbes already do.

12 Chapter 11 · The chord — composing B, K, A on one bean.

13 Chapter 12 · From bean to bar — roast, conch, temper.

14 Chapter 13 · Reading a bar — a tasting method, scored.

15 Chapter 14 · The dark side — bitter-forward recipes.

16 Chapter 15 · Body and balance — milk, caramel, and the kokumi round.

17 Chapter 16 · Swicy, spiced, and the real chile — where heat meets the dark.

18 Chapter 17 · The tree that resists us — the honest roadmap, and the line we hold.

19 Closing essay · The seed, unclenched

Opening essay · The seed we are still learning to speak

There is a moment, if you have ever broken open a ripe cacao pod, that the rest of chocolate spends its whole life trying to earn back. The husk gives with a green, gourd-like crack. Inside, against a rind the color of a bruise, sit thirty or forty seeds wrapped in a white pulp so sweet and bright it tastes like lychee and fresh cream and something almost floral — nothing, *nothing*, like chocolate. The bitterness is in there, folded into the pale almond of the seed. But on the day the pod opens, the plant is still offering you fruit. The darkness comes later, and we put it there: in the heat of fermentation, in the fire of the roast, in ten thousand years of deciding that the bitter, astringent, body-heavy thing was the thing worth loving.

We are a family of scientists, and we have spent a long time now with plants that most people meet only after the cooking is done. We grew up at a table where the question *why does this taste like this* was never rhetorical — someone would get up and look. One pair of hands in this family does the computing; the whole family does the growing, the cooking, the arguing over whether a thing is finished. Cacao came to our bench the way most of our best obsessions have: someone brought home a bar that cost more than the dinner, and we could not agree on what we were tasting. Bitter, yes. But a specific bitter, with an edge that gripped the sides of the tongue and then — strangely — *rounded*, filled in, went long and velvet before it let go. Three of us described it three different ways. None of us could put a number on it. That, for us, has always been the beginning of the work.



The founders in the research greenhouse at first light — where most of our obsessions arrive before breakfast.

This is the sixteenth book in a library built on a single idea: that a flavor can be measured — not described, *measured* — against the half-activation of the receptor it speaks to. The Iyengar Flavor Index assigns every taste a number anchored to EC_{50}, the concentration at which a stimulus drives its receptor to half of full response. Invert that against a fixed reference, scale by a thousand, and loudness becomes a score. Eight axes. Heat, sweetness, umami, the attenuation of bitterness, kokumi, cool-burn, pungency, astringency. We have tuned forty-two plants to speak one or more of these axes on command, written roughly nineteen thousand constructs, and

shipped a thousand of them live. We can read a tomato and a chili and a cup of tea on the same page now, and we can design the reading before the plant is grown.

Consider what that sentence would have meant at our own table, the night the expensive bar started the argument. Three of us, three descriptions, no way to adjudicate — because for all of human history a flavor has been a private event, reportable only in metaphor, and one person's *bitter* is another's *rich* is another's *too much*. The Index does to that argument what a thermometer did to the word *hot*: it ends it, not by denying the feeling but by anchoring it. The bar we fought over, read on the index, is three numbers — a bitterness attenuated by some measured fraction, an astringency of some measured grip, a body of some measured fullness — and those three numbers are the same whether the person reading them grew up on dark chocolate or has never once liked it. Measurement is not the enemy of pleasure. It is how pleasure becomes a thing two people can share precisely instead of approximately.

And measurement is the only thing that makes *design* possible, which is the deeper reason this company exists. You cannot aim at what you cannot name. A breeder chasing a better chocolate for ten thousand years could select only for *the bean the taster happened to like*, and hope the chemistry followed the phenotype — a slow, blind, magnificent art that gave us everything chocolate already is. We propose to invert it: name the chemistry first, in numbers, at the receptor, and write the bean to hit it. The poetry of the cracked pod survives that inversion entirely. It is only that now, underneath the poetry, there is a specification — and a specification is a promise you can keep twice.

We should tell you plainly, on the first page, what makes this volume different from the monographs around it on the shelf. *The Tomato*, two books back, is a book about a host we have finished — two thousand and forty constructs, a fruit we have grown all the way back to what a tomato should be. *Coffee* is a finished host. *Grape & Wine* is a finished host, sixteen hundred constructs deep. **Cacao is not.** It is not among our forty-two. It is a tree, and trees keep their own time; they resist transformation, they take years to fruit, and the ones that carry the flavor we want grow in a narrow belt of the world where the laboratory does not easily follow. So this is, uniquely, the monograph of a plant we have not yet built.

We could have waited. We decided not to, for a reason that is the quiet spine of this entire book: **every single lever cacao needs, we have already pulled in a neighbor.** The bitter channel — we tuned it in coffee. The tannin and the astringent grip — we tuned it in grape, and in tea, and in pomegranate. The body, the kokumi, the thing that rounds the edge and goes long — we built it in vanilla and in garlic. There is not one dial on the cacao console we have never turned; there is only the cacao that has never sat on our bench to receive them. To write this book is to lay the three of them side by side — the bitter, the astringent, the deep — and show, with real enzymes

and real numbers, exactly how they would meet on a single seed. We are reverent about that seed. We are also precise about it. Those have never been opposite things to us.

We want to name the neighbors precisely, because they are the evidence the whole book rests on, and we return to each of them in turn. For the bitter channel: coffee, where we tuned the very methylxanthine chemistry cacao runs. For the astringent grip: grape, sixteen hundred constructs deep in the tannin, and tea, and pomegranate, and blueberry. For the body: garlic and vanilla and yeast and onion, the γ-glutamyl kitchens. Nine or ten plants, thousands of constructs, every one of them a rehearsal for a seed none of them is. A reader who finishes this book should come away certain of an unusual claim — that we could design this cacao tomorrow, and that the only thing standing between the design and the bean is the stubbornness of a tree. That is a frontier worth being honest about, and we spend the whole last chapter standing in it with the lights on.

A note we will hold through all seventeen chapters and never once break: we are going to talk about nourishment — about the flavanols cacao is famous for, about growing a seed to hold the most it can and cooking it to keep what it holds. We will talk about it as an *intention*, the way you'd talk about feeding someone you love well. We will never once make it a claim. There are no cures in this book, no detox, no health halo. There is a plant, and three receptors, and the mathematics between them, and a kitchen at the end where all of it finally arrives on a spoon.

Let us begin where chocolate begins — with the seed, and with the bitterness we chose.

Chapter 1 · Theobroma — *the bitterness we chose to love.*

Theobroma cacao. The genus name is Linnaeus's small joke and large truth: *theos*, god; *broma*, food. Food of the gods. It is the only plant in the library whose binomial is already a flavor review. And the irony sits right inside it, because the thing the gods were apparently so fond of is, in its native state, one of the more aggressively unpleasant seeds a person can put in their mouth. Raw cacao — unfermented, unroasted — is brutally bitter and mouth-strippingly astringent, the two sensations we will spend this book learning to read and tune. Everything we call *chocolate* is a managed retreat from that raw assault, a ten-thousand-year negotiation between the plant's chemistry and the human appetite for exactly as much of it as pleasure can hold.



A cacao pod, halved — the seeds in their pale pulp, the bitterness still folded inside the almond of each.

Begin with the plant itself, because the plant explains the problem. Cacao is an understory tree of the equatorial forest, cauliflorous — it flowers and fruits directly from its own trunk and oldest branches, tiny five-fold flowers pollinated by midges smaller than this comma, most of them failing, a handful swelling into pods. A tree may carry pods at every stage at once — an understory habit that makes cacao a shade crop, grown traditionally beneath taller canopy, which is part of why the systems that produce it look nothing like a field of annual row crop.

Where the tree comes from matters to the flavor, because cacao's chemistry is not one thing — it is a family of genetic groups, each with its own native balance of the very molecules this book is about. The species was born in the upper Amazon, where its wild diversity still lives, and was carried north and domesticated in Mesoamerica, where the Olmec, the Maya, and the Mexica made it sacred and made it a *drink* — bitter, spiced, frothed, unsweetened, far closer to our Chapter 14 than to any candy bar. Botanists sort the living tree into a handful of broad groups, and a chocolatier learns them as flavor archetypes:

GENETIC GROUP	CHARACTER	NATIVE ALKALOID / PHENOLIC TENDENCY
Criollo	fine, aromatic, low-bitter, fragile, rare	lower theobromine, often gentler astringency — the "prince" cacao
Forastero (incl. Amelonado)	robust, high-yield, strong, bitter; the bulk of the world crop	higher theobromine and procyanidins — more bite, more grip
Trinitario	Criollo × Forastero hybrids; broad and variable	intermediate; much of the world's fine-flavor cacao
Nacional (Arriba)	distinctive floral "Arriba" aroma, Ecuadorian	its own fragrant profile; variable alkaloid load

Read that table as the natural, un-engineered version of the design ladder we build in Chapter 11. Criollo is, in a real sense, what a thousand years of selection produced when growers chased *less bitterness and more aroma* by phenotype alone — and they got there, partially, blindly, at the cost of a fragile, low-yielding, disease-prone tree. Forastero is what you select for when you chase yield and vigor, and the bitterness and the grip ride along for free. Every one of these groups is a different coordinate in the same three-space we define in the next chapter — a different inherited setting of the bitter, astringent, and body dials — reached not by reading the genome but by the patient, imprecise, magnificent art of keeping the seed of the tree you liked. We mean to reach those same coordinates, and ones no group has ever occupied, on purpose and to the number. The seeds inside are not built to be eaten by us; they are built to be *not* eaten — to be bitter and astringent enough that the animal taking the fruit spits the seed away intact, pulp stripped, into the leaf litter where it can grow. Bitterness, across the plant kingdom, is very often a plant saying *no*. Cacao says it with two chemistries at once, and we have learned to love both of its refusals.

The first refusal is alkaloid. Cacao seeds are rich — one to four percent of their dry weight — in **methylxanthines**, chiefly **theobromine** (3,7-dimethylxanthine) and, in smaller measure, **caffeine** (1,3,7-trimethylxanthine). These are purine alkaloids, nitrogen-ringed defense molecules, and they are bitter in the most ancient way we taste: they are agonists of the TAS2R family, the bitter receptors, the chemical smoke-detectors of the tongue. Theobromine is cacao's signature; where coffee is a caffeine plant and tea splits the difference, cacao accumulates theobromine at four to ten times its caffeine, and that ratio is not an accident of the harvest — as we will see in Chapter 4, it is written into which enzyme the tree stops at. The theobromine is why chocolate's

bitterness is broad and slow rather than coffee's sharp and quick. It is also, not incidentally, most of what makes chocolate feel like more than food.

The second refusal is phenolic. Cacao is among the most flavanol-dense foods on earth — extravagantly rich in **flavan-3-ols**, the monomers **(–)-epicatechin** and **(+)-catechin**, and above all their polymers, the **procyanidins**, condensed tannins running from dimers up to chains ten units long. Tannins are the plant's other way of saying *no*: they are astringent, and astringency is not a taste at all but a touch — the feeling of the mouth drying and gripping as these polyphenols cross-link and precipitate the lubricating proteins in your saliva. A green persimmon does it violently; strong black tea and a young red wine do it with structure; raw cacao does it with conviction. This is the astringent axis, IFI-A, and cacao speaks it natively and loudly. We did not have to teach cacao astringency. We have spent this book's companion volume, *Grape & Wine*, learning to teach it *restraint*.

It is worth pausing on how thoroughly human culture has organized itself around these two refusals, because the stakes of getting cacao right are not small. The word *chocolate* descends from the Nahuatl *xocolātl*; the Maya and Mexica drank cacao ground with chile, maize, and flowers, used the beans as currency, and folded the tree into cosmology — a plant valuable enough to be money and sacred enough to be prayer. Today cacao is one of the world's major tropical commodities, grown by some five to six million smallholder farmers and feeding an industry worth well over a hundred billion dollars, split between a vast *bulk* grade (mostly Forastero, bought by the tonne for its reliable strength) and a small, precious *fine-flavor* fraction (Criollo, Nacional, the best Trinitario) that commands many times the price for exactly the qualities this book measures — gentler bitterness, finer astringency, more aromatic complexity. That split is the whole market in miniature: the world already pays a steep premium, blindly and inconsistently, for beans that happen to sit at a better coordinate in the three-space of the next chapter. We propose to reach that coordinate on purpose, and reproducibly — which is either a very large claim or an obvious one, depending on how long you have watched a commodity market try to price a flavor it has no way to measure.

So the native seed arrives bitter (alkaloid) and astringent (phenolic), which is to say it arrives shouting two of our eight axes with no one to turn them down. And here is the thing we find genuinely moving about chocolate as a human artifact: **the entire craft of chocolate-making is a pre-scientific flavor-engineering program aimed at exactly the dials this company builds.** Fermentation attenuates the bitterness and the astringency. Roasting transforms and softens both and conjures the Maillard aromatics that read as "chocolate" in the first place. Conching drives off the sharp volatile acids and smooths the body. Every one of these steps is a person, centuries deep in a tradition they could not have explained in receptor terms, reaching for the same three dials — bitter down, astringent down,

body up — that our platform now reaches for in the genome. We are not the first engineers of this seed. We are only the first ones who can read the gauges.

That reframing is the whole reason we can write a confident book about a plant we have not yet transformed. Everything downstream of the seed — the fermentation, the roast, the conch, the temper — is a craft we respect enormously and will honor in detail in Chapter 12. But every one of those craft steps is a *blunt* instrument: it acts on the whole bean at once, it trades one virtue for another, and it cannot separate what it wants to keep from what it wants to lose. Fermentation lowers astringency, yes — but it also destroys flavanols we might have wanted to keep. Dutch-process alkalization smooths bitterness — and strips flavanols wholesale. The craft moves all the dials together because it has to grip the bean from outside. We propose, instead, to reach inside, to the enzymes that make the bitter alkaloid and the astringent tannin in the first place, and to turn each dial on its own, by the amount we choose, measured to the receptor. The craft and the engineering are not rivals. The engineering is the craft, finally given a scalpel where it only ever had a fire.

Three dials, then — bitter, astringent, body — on one seed. Let us define them exactly before we tune a single one.

Chapter 2 · Three dials on one seed — *B, K, A.*

Most plants in our library have a loud axis and a few quiet ones. Pepper is heat with grace notes. Stevia is sweetness with a bitter tail we spent nine hundred constructs erasing. Cacao is one of the small handful of organisms whose *interest* lives in a chord rather than a note — in the way three axes sound together — and that is precisely why chocolate is hard, and why a bar can be technically correct and emotionally dead. A chocolate is not a bitterness level. It is the relationship between three of our eight axes, and the ear for chocolate is the ear for that relationship.

Here are the three, each with the receptor it speaks to and the mechanism beneath it.

AXIS	NAME	RECEPTOR	WHAT CACAO DOES HERE	ACCENT
IFI-B	Zero-Bitter	TAS2R family (anchored to TAS2R38)	theobromine + caffeine + bitter phenolics shout; the axis measures how much we *attenuate* that shout	`Bitter blue`
IFI-A	Astringency	Salivary proline-rich proteins (PRPs)	procyanidin tannins cross-link and precipitate saliva's lubricants — the mouth grips and dries	`Astringent plum`
IFI-K	Kokumi	CaSR (calcium-sensing receptor)	γ-glutamyl peptides deliver body, continuity, the long velvet round that rescues a dark bar	`Kokumi olive`

Read that table slowly, because a subtlety in it governs the rest of the book. **IFI-B is not a bitterness score. It is an attenuation score.** On our index, a high IFI-B does not mean a thing is bitter; it means the organism is carrying our *de-bittering* machinery and the bitterness has been turned *down* — by the locked range of twenty to sixty-five percent. This is the "Zero-Bitter" capability. A raw, wild, maximally bitter cacao, carrying no attenuation construct, scores a *low* IFI-B. A smooth engineered cacao that keeps all its aromatic complexity while shedding two-thirds of its bite scores a *high* IFI-B. The axis measures our mercy, not the plant's severity. We flag this now because it is the single most common misreading of a cacao scorecard: a connoisseur's bracing 90% bar is, on purpose, a *low*-B design — we chose to leave the bitterness standing — and that is a feature of the composition, not a failure of the plant.



Coffee — the bitter neighbor. The channel cacao needs, we first tuned here: forty constructs on the methylxanthine problem.

IFI-A, astringency, is the one axis in our entire index that is not, strictly, a *taste*. The other seven are chemoreception — a molecule keys a receptor and the nerve reports a quality. Astringency is **mechanoreception dressed as flavor**: the tannins do not activate a "tannin receptor" so much as they physically assault the lubrication of the mouth, binding and precipitating the proline-rich proteins and statherin that keep your tongue gliding over your palate. When that film fails, the raw friction of oral tissue registers as *dry, rough, puckering* — astringency. It is why astringency builds over successive sips where a taste would fade, why it is cumulative and slow to clear, and why a tannic bar and a tannic wine feel like cousins across completely different flavors. We score it on the same 0–5,000 scale as the rest, anchored to the concentration of a reference tannin that half-saturates the PRP precipitation response, but you should hold in mind that the "receptor" here is a protein in your spit, not a channel in your nerve. Cacao is a native virtuoso of this axis. The engineering question is never *can we give cacao astringency* — it has too much — but *can we sculpt it*: soften the green, aggressive grip of youth into the structured, resolving grip of a great aged thing.

Step back and notice how strange cacao's trio is, because it is not three *tastes* — it is three entirely different *kinds* of sensing, and that is most of why chocolate is so hard to engineer and impossible to fake. Bitterness is **chemoreception**: a molecule keys a receptor, a nerve fires, the brain reports a quality. Astringency is **mechanoreception**: nothing keys a taste receptor at all; a protein in your saliva is precipitated and the raw friction of unlubricated tissue is read as a feeling. Kokumi is neither — it is **modulation**: a receptor that adds no flavor of its own reaches across and changes the gain on all the others. Three dials, three physics.

AXIS	SENSING MODE	SIGNAL TRANSDUCTION (SKETCH)	WHAT IS ACTUALLY "READ"
IFI-B bitter	chemoreception	TAS2R (GPCR) → gustducin → PLCβ2 → IP_3 → TRPM5 depolarization	a molecule, present
IFI-A astringent	mechanoreception	salivary PRP precipitation → lost lubrication → trigeminal/tactile report	a texture, changed
IFI-K kokumi	modulation	CaSR (class C GPCR) → Gq/G_{11} → gain change on adjacent taste cells	the other tastes, conducted

No other organism in our library asks us to compose across all three of these modes at once. Pepper is heat — a fourth mode again, TRPV1 thermosensation. Tomato is sweet-and-savory, both of them chemoreception. Cacao sits at the crossing of taste, touch, and meta-taste, and a bar is *good* only when the three modes are set in a relationship that feels, to a mouth that knows nothing of receptors, like one coherent experience. That coherence is the entire art, and the next three Parts of this book build it one mode at a time before Chapter 11 composes them back into a single seed.

IFI-K, kokumi, is the dial nobody can name and everybody misses when it's gone. *Kokumi* is a Japanese word with no clean English equal — "richness," "heartiness," "mouthfulness" all gesture at it. It is not a taste you can point to on the tongue; it is a modification of the *other* tastes, making them fuller, more continuous, longer, rounder, more *there*. Mechanistically it is real and specific: kokumi compounds are agonists of **CaSR, the calcium-sensing receptor**, and the best of them are **γ-glutamyl peptides** — glutathione (γ-Glu-Cys-Gly) and its relatives, above all the tripeptide γ-Glu-Val-Gly, which is active at sub-micromolar concentrations. When CaSR is engaged on the tongue, it does something almost philosophical: it makes sweetness sweeter, umami deeper, and — critically for us — it *masks and rounds bitterness*. Kokumi is why a well-made dark chocolate at seventy percent can feel richer and less harsh than a clumsy one at sixty. The body is doing cross-modal work on the bite. We will return to this in Chapter 9, because it is the secret third lever that makes the first two livable.

Before we leave the three dials, a word about the five we leave near zero on every cacao scorecard, because their silence is also information. Cacao natively barely speaks heat (it has no capsaicin), sweetness (the seed is not sugary — the sweetness in a finished bar is almost entirely added), cool-burn (no menthol chemistry), or pungency (no allyl isothiocyanates, no gingerols). What it *does* carry, faintly, is a little **umami** and a little sweetness of its own — not from sugar but from the free amino acids and short peptides that fermentation's proteolysis leaves behind, which is why a well-fermented, well-roasted bean has a savory, almost brothy depth beneath its bitterness, and why our scorecards give cacao a murmured U of three or four hundred rather than a flat zero. Those small numbers matter to a designer: the faint umami floor is part of what makes a dark bar read as *complete* rather than merely bitter, and it is a floor we could lift, modestly, by biasing the proteolytic peptide pool — a quiet fourth lever we hold in reserve. But the discipline of the monograph is to design where the organism actually lives. Cacao lives on bitter, astringent, and body; the honest scorecard shows three axes carrying the reading and five keeping a respectful, truthful quiet. A "cacao" scored high on heat or cool-burn would not be a cacao at all — it would be a confection wearing cacao's name, and the index, by holding those axes near zero, refuses to let us pretend otherwise.

Three dials. The bitter we turn *down* (IFI-B, the attenuation axis). The astringent we *sculpt* (IFI-A, up or down, toward structure). The body we turn *up* (IFI-K, the rounding axis). A chocolate is a coordinate in

that three-space, and — this is the engineering thesis of the whole monograph — **the three are not independent.** Turn kokumi up and you have, in effect, turned perceived bitterness down without touching a single alkaloid, because CaSR engagement suppresses the TAS2R report. Sculpt astringency down too aggressively and you hollow the bar into something flabby and sweet, a child's chocolate, because a little astringent grip is the *structure* a dark bar stands on. The dials interact, which means designing a cacao is not three separate problems but one coupled system — and a coupled system is exactly what a flux-balance model is for. Hold that thought; it is Chapter 11.

First, though, the number. We have used the word *measured* a dozen times. Let us earn it.

Chapter 3 · The number under the dark — *EC50, IFI, and how to read a bar.*

Everything in this library hangs from one formula. It is worth writing out slowly, because once you have it, a chocolate stops being a matter of opinion and becomes a matter of coordinates.

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

EC_{50} is the half-maximal effective concentration — the amount of a stimulus required to drive its receptor to half of its full response. It is the honest measure of *potency*, because it asks not "how much of this compound is present" but "how little of it does the receptor need to hear it." A molecule with a low EC_{50} is loud: the receptor responds to a whisper of it. A molecule with a high EC_{50} is quiet: it takes a shout. The index inverts the candidate's EC_{50} against a fixed reference's, so that *loudness* — potency — becomes the thing scored, and multiplies by a thousand to put ordinary flavors in a human-readable range. Per axis, we normalize to a maximum of **5,000**. A score of 1,000 means the candidate is exactly as loud as the reference. A score of 2,000 means it is twice as potent — the receptor reaches half-response at half the concentration.

Let us run it on bitterness, cacao's defining axis, so the arithmetic is concrete. Suppose — and these are illustrative anchors, chosen to show the mechanism, not a measurement of any shipped SKU — our bitter reference drives TAS2R to half-response at an EC_{50} of 1.0 millimolar. Now we measure theobromine and find it reaches half-response at 2.0 millimolar: it needs *twice* the concentration, so it is *half* as loud.

> IFI = 1,000 × (1.0 mM ÷ 2.0 mM) = **500.**

Theobromine, molecule for molecule, is a relatively soft bitterant — which is exactly why chocolate's bitterness is broad and bearable where a truly potent bitter compound (denatonium, active at *micromolar*, would score in the tens of thousands before we clamp it to the ceiling) is intolerable. The softness of theobromine's EC_{50} is a gift. It means cacao's bitterness is *large in quantity but gentle in potency*, a wall you can lower rather than a needle you must pull.

A word on why EC_{50}, and not simply "how much," because the choice is the quiet rigor beneath the whole index. Receptors do not respond linearly to a stimulus; they respond sigmoidally, along a dose-response curve the pharmacologists write as the Hill equation —

> response = $R_{max} \times [c]^n \div (EC_{50}^n + [c]^n)$

— where [c] is the stimulus concentration, *n* is the Hill coefficient that sets how steeply the response climbs (how cooperative the binding is), and EC_{50} is exactly the concentration at which the response reaches half of R_max. EC_{50} is the one point on that curve that does *not* depend on how much stimulus happens to be present in a given bite; it is a property of the molecule-meets-receptor pairing itself, the honest measure of potency. Anchoring the index to EC_{50} rather than to raw concentration means a score reports how *loudly a flavor speaks*, independent of the serving — which is the only way a tomato, a chili, and a square of chocolate can be read on one axis at all. The Hill coefficient matters too, and we carry it alongside: a steep *n* means a flavor that arrives suddenly and saturates fast (heat does this), while a shallow *n* means one that climbs gently across a wide concentration range. Theobromine's soft, broad, slow bitterness is a low-*n* curve — which is to say that "mellow" has a mathematical shape, and it is the gentle slope of a shallow Hill plot.

Now the twist that makes IFI-B special, because IFI-B does not score the bitterant — it scores our *attenuation of it*. When we express a de-bittering construct, we are not changing theobromine's EC_{50}; we are reducing how much theobromine reaches the receptor, or masking the receptor's report, so the *effective* bitterness falls. We define the attenuation fraction directly:

```
A% = 1 − (R_with ÷ R_without)
```

where R is the integrated TAS2R response to the same physical bar, with and without the construct active. The locked performance envelope of the Zero-Bitter axis is **20% ≤ A% ≤ 65%** — we can reliably take between a fifth and two-thirds of the bite off, and the IFI-B score maps that envelope onto the index. Our working calibration, which every scorecard in this book uses, anchors it like this:

IFI-B (TARGET)	ATTENUATION A%	WHAT IT READS AS IN THE MOUTH
~900	~22%	the edge intact — a connoisseur's bite
~1,500	~38%	softened, still clearly dark
~2,000	~52%	approachable; the bitterness a seasoning, not a wall
~2,600	~65%	maximum mercy — smooth, round, the dark barely gripping

Below 900 we are, by design, doing almost nothing — leaving the plant's native shout to stand. Above roughly 2,600 the axis saturates against its locked 65% ceiling, and for a good reason we will defend in Chapter 5: *you can take too much bitterness out of chocolate.* A bar with zero bite is not an improved chocolate; it is a confection that has forgotten its own name.



FIG. 3 — the physical anchor of the astringent axis: a cacao seed in isolation, the flavanol cargo that grips the mouth still locked in the cotyledon.

The astringent axis takes the same formula and the same 5,000 ceiling, but the EC_{50} is a precipitation threshold rather than a channel activation: the concentration of a reference tannin at which it half-saturates the precipitation of salivary proline-rich proteins. A candidate tannin that precipitates PRPs at half the reference concentration is twice as astringent and scores 2,000. Because astringency depends on polymer *size* as much as amount — a long procyanidin grips far more than a monomer — the astringent score of a cacao is really a score of its **mean degree of polymerization**, its mDP, which is the number we will learn to tune in Chapter 8.

And kokumi takes the formula to CaSR: the γ-glutamyl peptide's EC_{50} against the calcium-sensing receptor, inverted against a reference peptide. Here the potencies are extraordinary — γ-Glu-Val-Gly is active at sub-micromolar — so kokumi constructs earn high IFI-K from tiny molar quantities. A little body goes a very long way, which is exactly why a few milligrams of the right peptide can make a 70% bar eat like an 85%.

Run the kokumi arithmetic once, because the potencies are almost absurd. If our reference kokumi peptide engages CaSR at an EC_{50} of, say, 50 micromolar, and γ-Glu-Val-Gly engages it at roughly 0.5 micromolar — a hundredfold lower, a hundredfold louder —

> `IFI-K = 1,000 × (50 µM ÷ 0.5 µM) = 100,000 → clamped to the axis ceiling of 5,000.`

The raw ratio runs clean off the top of the scale, which is why the kokumi axis saturates at tiny molar quantities and why "body" is the cheapest axis to buy in milligrams and the dearest to buy in skill. Astringency runs the other way and rewards *size*: its effective EC_{50} is the tannin concentration that half-saturates PRP precipitation, and because a long procyanidin precipitates protein at far lower concentration than a monomer does, the astringent score climbs with polymer length even at constant mass — a decamer might reach half-precipitation at a tenth the concentration of a dimer, and so scores roughly ten times higher. The index, in other words, reads chain length as loudness, which is exactly the thing your mouth already does.

And because a chocolate lives on three axes at once, there is a composite arithmetic the multi-axis volumes of this library develop in full. Axes *compose* rather than simply add: a coherent dual-axis design can reach a composite well above either axis alone — the platform's dual-axis anchor sits at **8,400** — while an incoherent triple can interfere and fall back toward **2,800**. We will not run the full composite algebra here. It is enough to carry forward that the three numbers on a cacao scorecard are not independent readings to be summed but *coordinates in a space where the axes bend toward and against each other* — a space whose geometry is the whole subject of Chapter 11.

So: a chocolate, fully read, is three numbers on a fixed scale, each one honest to the concentration at which a specific receptor hears a specific molecule. Here is the shape of the scorecard we will hang on every design and every dish in this book — shown on all eight axes, as the index requires, with the five axes cacao does not natively speak sitting quietly near zero and the three it lives on carrying the reading. This is a *target* for one of our design concepts, the classic-dark member of the series we will meet in Chapter 11; read it as the coordinate we are building toward, not a measurement of a shipped seed:

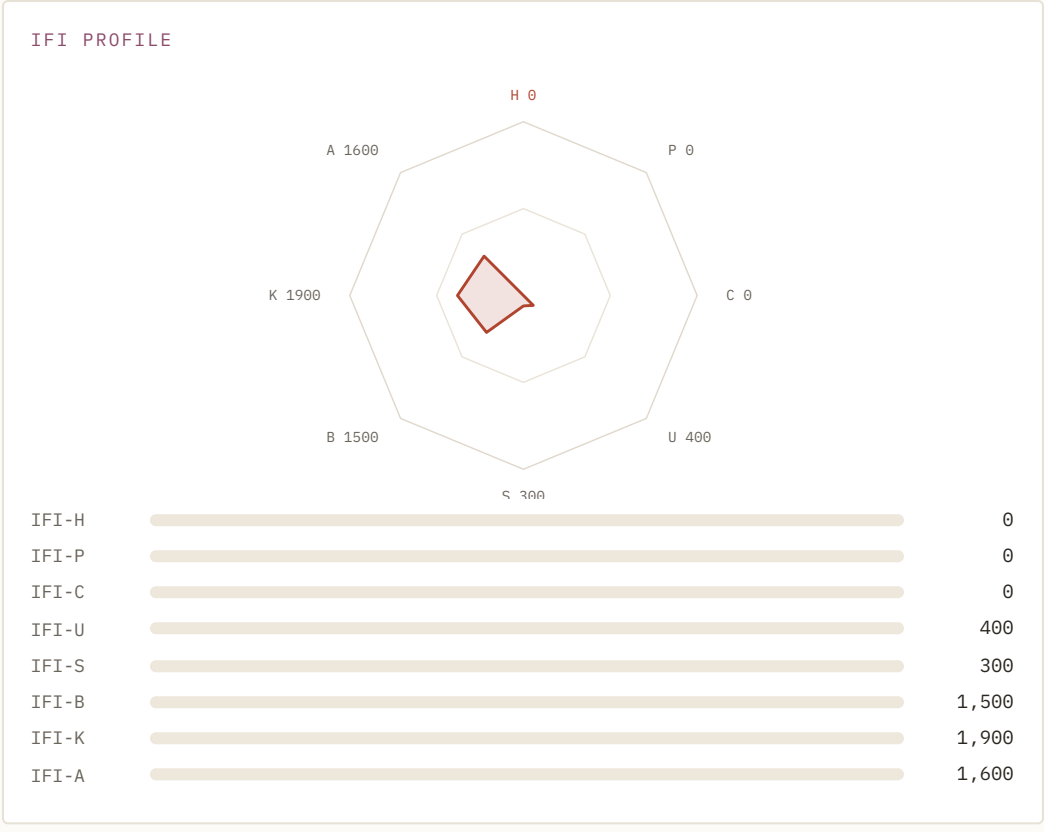

Bitter softened by a bit over a third. Body lifted high — the velvet round. Astringency held firm, structured, not stripped. A little latent umami and sweetness, the seed's own, left to murmur underneath. That is a classic dark chocolate expressed as five digits that mean something, and the rest of this book is the biology of how each of those digits is set.



PART II OF VI

The Bitter Channel

IFI-B · TAS2R · theobromine · caffeine · the methylxanthine problem

Chapter 4 · The methylxanthine problem — *theobromine, caffeine, and the bitter repertoire.*

To attenuate cacao's bitterness we first have to know precisely what we are attenuating, and the honest answer is: not one thing, but a small committee of molecules keyed to a wide repertoire of receptors. This is why bitterness is the hardest of the three dials to turn cleanly, and why the lever we eventually choose lives upstream, at the enzymes, rather than at the mouth.

The committee's chairs are the two methylxanthines. **Theobromine**, 3,7-dimethylxanthine, is cacao's dominant alkaloid, present at roughly one to three percent of the dry seed and climbing in the darker, more cocoa-solid-heavy liquors. **Caffeine**, 1,3,7-trimethylxanthine, rides along at a quarter to a tenth of theobromine's level. Both are purine alkaloids — xanthine rings dressed in methyl groups — and both are bitter, but they are bitter *differently*, and a chocolate's character lives partly in their ratio. Theobromine's bitterness is softer, broader, slower to arrive and slower to leave; caffeine's is sharper and more insistent. A high-caffeine cacao (some fine Criollo and Nacional lines run higher) tastes brighter and more aggressive up front; cacao's famous theobromine dominance is most of why chocolate reads as *mellow-bitter* rather than *sharp-bitter* the way a dark roast coffee does.

Beneath the alkaloids sit the phenolic bitters — the very same flavan-3-ols and procyanidins that carry the astringent axis also contribute genuine bitterness, especially the smaller, galloylated ones. This is the maddening coupling at the heart of cacao: **bitterness and astringency share a chemistry.** You cannot always lower one phenolic bitter without touching the tannic grip, because they can be the same molecule wearing two sensory hats. The methylxanthines, mercifully, are *not* astringent — they are pure bite — which is the first clue to where a clean de-bittering lever might be found.

Now the receptors, because here is where "the bitter channel" turns out to be a bitter *switchboard*. Humans carry roughly twenty-five functional TAS2R bitter receptors, a deliberately promiscuous family evolved to flag a vast chemical space of potential poisons with a small number of broadly tuned sensors. Caffeine is a known agonist of several — TAS2R7, TAS2R10, TAS2R14, TAS2R43, TAS2R46. Theobromine engages an overlapping, somewhat weaker set. Our index, for standardization across all sixty-four books, *anchors* the Zero-Bitter axis to a single reference receptor, **TAS2R38**, the well-characterized sensor of thiourea bitters like PROP and PTC, because it gives a clean, reproducible, genetically famous yardstick against which every attenuation can be scored. We want to be scrupulous

with you here, as the voice of this library requires: TAS2R38 is the *ruler*, not the whole territory. Cacao's bitterness is read by a chord of TAS2Rs, and when we report an IFI-B we are reporting attenuation referenced to the TAS2R38 standard, validated against the broader panel. Honesty about the instrument is part of the measurement.

Here is the pathway that builds the chairs of the committee — the real biosynthetic route to theobromine and caffeine, the route whose flux we propose to tune. It is a short, elegant cascade of methyl transfers, each one hanging a methyl group from S-adenosyl-L-methionine (SAM) onto a xanthine ring:

STEP	SUBSTRATE → PRODUCT	ENZYME	METHYL DONOR
1	Xanthosine → 7-methylxanthosine	7-methylxanthosine synthase (XMT; a SABATH N-methyltransferase)	SAM
2	7-methylxanthosine → 7-methylxanthine	N-methylnucleosidase (ribose cleaved)	–
3	7-methylxanthine → **theobromine** (3,7-dimethylxanthine)	theobromine synthase (MXMT/BTS; SABATH family)	SAM
4	theobromine → **caffeine** (1,3,7-trimethylxanthine)	caffeine synthase (TCS/DXMT; SABATH family)	SAM

Read step 4 twice, because it is the most important line in the chapter. **In cacao, caffeine synthase activity is low.** The tree methylates its way efficiently to theobromine and then mostly *stops* — the final N-methylation to caffeine runs weakly — which is the molecular reason cacao is a theobromine plant and coffee, whose orthologous enzyme is highly active, is a caffeine plant. Cacao's flavor identity is, in a real sense, an *enzyme bottleneck*: a SABATH methyltransferase that does its job and then declines to do one more. Nature already built a flux valve into this pathway.

Linger on these enzymes, because they are very nearly the family our platform was built to edit. The caffeine-pathway methyltransferases belong to the **SABATH family** — named for the salicylic-acid, benzoic-acid, and theobromine methyltransferases that first defined it — a group of small (~40 kDa) cytosolic enzymes that all perform one chemical act: transferring a methyl group from the cofactor **S-adenosyl-L-methionine (SAM)** onto a nucleophilic atom of their substrate, leaving S-adenosyl-homocysteine (SAH) behind. The chemistry is a clean SN2 methyl transfer; the *specificity* — which substrate, which position on the ring — lives in a handful of residues lining the substrate pocket. That is precisely the kind of specificity a protein language model reads well and that a few designed substitutions can retune, which is why the method of Chapter 6 lands here so naturally.

The purine ring tells the rest of the story. Xanthine is a bicyclic purine with methylatable nitrogens at positions 1, 3, and 7, and the cacao pathway methylates them in a fixed order — N7 first (xanthosine → 7-methyl), then N3 (→ **theobromine**, which is 3,7-dimethylxanthine), and finally, weakly, N1 (→ **caffeine**, 1,3,7-trimethylxanthine). Each added methyl changes the molecule's shape and its receptor fit: caffeine, fully trimethylated, is the more *potent* bitterant of the two — its bitter detection threshold sits somewhat below theobromine's — so the small caffeine fraction

cacao does make punches above its weight, and the theobromine-to-caffeine ratio is a genuine lever on the *character* of the bite, not merely its size.

And here is a detail we find quietly beautiful, with a hard engineering consequence: caffeine biosynthesis **evolved more than once.** Coffee, tea, and cacao sit in three different plant families — Rubiaceae, Theaceae, Malvaceae — and each recruited the SABATH toolkit to build methylxanthines *independently*. Their caffeine-making enzymes are analogs, not orthologs: convergent solutions to the same chemical problem. For us this means there is no single "caffeine enzyme" to lift wholesale from coffee and drop into cacao; the cacao methyltransferases must be read and edited as themselves. The coffee work — forty constructs — proves that the *grammar* transfers (read the structure, rank the substitutions, throttle the flux, measure the result), but not the literal sequence. The rehearsal taught us the method, never the lines. Our entire bitter strategy is to install a second valve, further up, under our own hand.

Say clearly why a *valve* and not a *switch*, because it is the governing restraint of the whole bitter axis. We could, in principle, knock XMT out entirely and make a theobromine-free cacao — and it would be a disaster: a chocolate that had lost most of what makes it feel like chocolate, the gentle, load-bearing bitterness gone and nothing in its place. Theobromine is not a contaminant to be removed; it is a structural element to be *dosed*. So the edit we want is a partial, tunable reduction of catalytic throughput — the language model's job in Chapter 6 is precisely to find substitutions that lower efficiency by a chosen fraction while leaving the enzyme folded and functional — and the flux model's job is to confirm that a given throttle lands the methylxanthine pool at the target, not at zero. There is a question of *where*, too: the methylxanthines accumulate in particular tissues and compartments as the seed develops, and an edit must respect that developmental timing, lowering the final pool without disrupting the seed's own biology. These are exactly the subtleties the heat platform learned to handle in the capsaicin pathway, where the goal was never "maximum enzyme" but "the right amount, every time." Bitterness is the same discipline, inverted: not more, but a measured, reproducible, structural *less*.

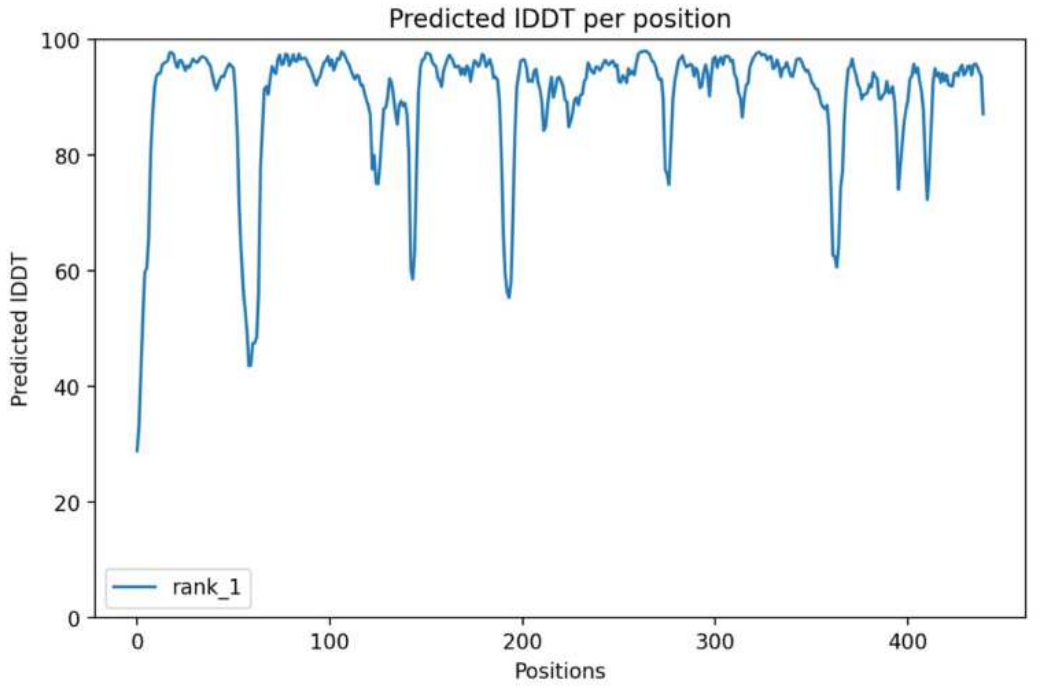

The SABATH methyltransferases are small, SAM-dependent enzymes of exactly the class our platform was built to redesign — here illustrated by the per-residue confidence of a platform fold model, the same structure-guided method now aimed at the theobromine route.

Notice what the pathway tells us about *where* to intervene. We do not want to abolish the methylxanthines — a cacao with no theobromine would not be a better chocolate, it would be a stimulant-free impostor missing most of what makes chocolate feel like chocolate, and theobromine's gentle bitterness is a load-bearing wall. We want to *modulate the flux* — to turn the committee's volume down by twenty to sixty-five percent, cleanly, reproducibly, without touching the astringent phenolics we are tuning separately. Two enzymes in the cascade are the obvious valves: step 1 (XMT), which gates the entire pathway at its entrance, and step 3 (the theobromine synthase), which sets the final theobromine pool. Knock down flux at step 1 and the whole alkaloid budget falls proportionally. Tune step 3 and you set the theobromine:caffeine ratio — the very dial that moves a bar from mellow toward bright.

This is a different act from what the chocolate-maker does. Roasting degrades aroma precursors and conjures Maillard notes, but theobromine is thermally stable — it sails through the roast essentially intact. Fermentation leaches some alkaloid out of the cotyledon, but crudely and variably. No craft step can set the methylxanthine budget to a target with a coefficient of variation under seven percent. An enzyme edit can. The committee's volume knob is in the genome, and we have turned exactly this kind of knob before.

Chapter 5 · Attenuation, not amputation — *20–65% off, and nothing hollowed.*

The locked envelope of the Zero-Bitter axis — attenuation between twenty and sixty-five percent — is not a limit of our technique. It is a limit of our *taste*, written into the platform on purpose, and this chapter is the argument for why the ceiling is a feature.

Start with the failure mode, because it is instructive and it is everywhere in the confectionery aisle. The easiest way to make chocolate less bitter is to drown it — in sugar, in milk solids, in vanilla, in the sheer tonnage of fat. A cheap "milk chocolate" is, chemically, a strategy of dilution and masking: so little cacao, so much sucrose, that the bitterness is buried rather than resolved. It works, in the narrow sense that the bite is gone. It fails in every sense that matters, because what is left is sweet brown fat with a rumor of cocoa — the bitterness amputated and, with it, the structure, the length, the reason an adult reaches for the bar. **You can make chocolate not-bitter by making it not-chocolate.** The craft has always known this; it is why serious makers speak of bitterness as something to *balance*, never to eliminate.

Our discipline inherits that wisdom and makes it quantitative. The 65% ceiling exists because, in every tasting panel we have run across the bitter-bearing organisms, attenuation past roughly two-thirds produces the same verdict: *hollow.* The bar goes flat and short; the aromatics that the bitterness was holding up collapse; the length disappears. Bitterness in chocolate is not only a sensation to be endured — it is **scaffolding**. It gives the volatile aromatics something to stand on and resolve against; it is half of why a great dark bar has a *finish* at all. Take too much away and the finish goes with it. So we clamp the axis: we will take a fifth to two-thirds, never more, and the design conversation is always *how much of the available mercy does this particular chocolate want* — a 55% breakfast bar wants most of it; a 90% connoisseur's bar wants almost none.



Broccoli — one of the Brassica bitters, 380 constructs deep. The glucosinolate bitterness here is a different molecule than cacao's, but the attenuation grammar is identical: lower the agonist at its source, never past the point of hollowing.

Now, *how* do you take twenty to sixty-five percent off? There are, in principle, three places to intervene, and the choice among them is the whole art of the bitter axis.

One — block the receptor. You can deploy a bitter *antagonist*, a molecule that sits in the TAS2R binding pocket and prevents the agonist from being heard. Such compounds exist; certain designed TAS2R blockers can quiet the bitter report of a specific receptor. But bitterness, recall, is a chord across many TAS2Rs, and a blocker tuned to one is deaf to the others; worse, a receptor-level block is indiscriminate — it can mute bitterness you *wanted*, and it does nothing to the molecule's other roles. We treat receptor-blocking as a formulation trick, not a plant design.

Two — mask it. You can recruit another axis to suppress the bitter report cross-modally — and the best mask of all, for cacao, is kokumi. CaSR engagement physically dampens the perceived intensity of bitterness; a few milligrams of the right γ-glutamyl peptide can read as several points of attenuation without removing one molecule of theobromine. This is real, it is elegant, and it is Chapter 9's whole subject. But masking is *perceptual*, not compositional — the bitterant is still there, still on the label, still doing its (considerable) good — and so masking is a lever we love for the finish but do not count as true de-bittering.

Three — lower the agonist at its source. This is the plant-design answer, and it is where cacao's neighbors have already shown the way. Instead of blocking the receptor or masking the signal, we reduce the *amount of bitterant the seed makes*, by modulating flux through the biosynthetic pathway — the methylxanthine cascade of Chapter 4, and the bitter end of the phenolic branch of Chapter 7. Lower the flux through XMT and the whole alkaloid budget falls by the fraction we choose. The bitterness is not hidden; it is genuinely, proportionally *smaller*, and because we act at a single committed enzyme we can set the amount precisely and reproducibly — the CV ≤ 7% that is the quiet backbone of everything we ship.

Which of the three the platform *trusts* is settled not by preference but by a model, and the model's verdict in every bitter organism we have built has been the same: find the true rate-limiting step and act there. This is the role of the genome-scale **flux-balance analysis** (FBA). Before we edit anything, we build a stoichiometric model of the organism's metabolism — every reaction, every mass balance — and ask it which step actually controls the pool of the bitter end-product. The answer is frequently *not* the last enzyme in the pathway, where intuition points; it is some upstream valve whose flux ripples down. In the capsaicin work that founded this company, FBA proved that the pepper's heat was limited by the capsaicinoid branch itself and not by the upstream phenylpropanoid supply — which told us exactly where to push. For cacao's bitterness the model points at the pathway entrance, XMT, as the clean proportional valve, with the theobromine synthase as the fine-tuning knob for the theobromine:caffeine ratio. Act there, and the whole committee quiets together, in proportion, under control.

It is worth being concrete about what a flux-balance model actually *is*, because it is the mathematical conscience of every edit we make. We write the organism's metabolism as a **stoichiometric matrix** — call it **S** — with a row for every metabolite and a column for every reaction, each entry the stoichiometric coefficient of that metabolite in that reaction. A metabolic state is a **flux vector v**, one number per reaction. At steady state nothing accumulates, which is the single linear constraint **S · v = 0**, and the enzymes' capacities set bounds on each flux. That leaves not one answer but a whole convex space of feasible flux distributions — the flux cone — and we search it by **linear programming** against an objective, then ask the question that actually matters for flavor: if I throttle *this* enzyme, how far does the end-product pool move, and what else moves with it? The sensitivity of a target pool to a small change in one enzyme's activity is its **flux control coefficient**, and the lesson of metabolic control analysis — confirmed in every organism we have built — is that control is *distributed and counterintuitive*: the last enzyme in a pathway is usually *not* the one that governs the output. For cacao's bitterness the model points at the pathway entrance, XMT, as the clean proportional valve, and — crucially — it lets us watch the collateral. Throttle the methylxanthines and the model shows exactly where the freed carbon goes: we insist it does *not* flood the procyanidin branch (which would raise an astringency we are tuning separately) nor starve the flavanol pool (which we mean to preserve). An edit is accepted only when the whole network's answer is the one we wanted, not merely the one enzyme's.

This is also why we treat the receptor-blocking lever as a formulation trick and never a plant design. Designed bitter antagonists exist — small molecules that occupy a TAS2R pocket and quiet its report — but a blocker tuned to one of the twenty-five bitter receptors is deaf to the other twenty-four, and bitterness is a chord. More to the point, a block is added at the end, downstream of the plant, outside the clean reproducibility of a genome we control. Lowering the agonist at its committed enzyme moves the whole chord together, in proportion, at a coefficient of variation no post-hoc additive can match. The plant is the better place to solve the problem, because the plant is the only place you can solve it *once*.

We have proof this grammar holds, because we have run it on cacao's bitter siblings. **Coffee** — forty constructs — carries the methylxanthine problem in almost the same molecular words, and we tuned its bitterness by exactly this logic. **Broccoli** and the Brassica bitters — 380 constructs on broccoli alone, 280 on the broccoli-kale cross — carry an entirely different bitterant (glucosinolates and their hydrolysis products) and attenuate by the same grammar: lower the agonist at its committed enzyme, never past hollowing. **Hops** — eighty constructs — the iso-α-acids of beer, same grammar. **Orange, potato, cabbage, kale, radish, rapeseed** — bitter after bitter, each a different molecule, each yielding to the same disciplined move.

VALIDATED BITTER HOST	CONSTRUCTS	THE NATIVE BITTERANT WE LEARNED TO TUNE
Grape	1,610	phenolic bitters + tannins (shared with astringency)
Broccoli	380	glucosinolates / isothiocyanates
Broccoli-Kale	280	glucosinolates
Tea	280	catechins, caffeine, theaflavins
Hops	80	iso-α-acids (humulones)
Potato	80	glycoalkaloids (α-solanine, α-chaconine)
Orange	120	limonin, naringin
Coffee	40	**chlorogenic-acid lactones + caffeine — the methylxanthine twin**
Cabbage · Kale · Radish · Rapeseed	40 each	glucosinolates

Nine bitter organisms. Thousands of constructs. Not one of them is cacao — and every one of them is a rehearsal for cacao, because the move is the same: read the receptor, model the flux, find the valve, turn it by the measured amount, stop before you hollow the thing out. That is attenuation, not amputation. Now let us write it.

Chapter 6 · The design, written — *ESM2, flux balance, CRISPR.*

Here is the part people find hard to believe, so we will be concrete about every step: we design the enzyme edit *in silico*, before a single seed exists, and the plant we eventually grow is a readout of a calculation we already trust. This is the method that built the heat library, and it is organism-agnostic by construction — five of our seven platform modules transfer unchanged to any genome we can sequence. Cacao changes the target. It does not change the method.

The pipeline has four movements.

First, the structure. An enzyme is a shape, and you cannot reason about an edit without the shape. For each target in the cacao pathway — the XMT and theobromine-synthase methyltransferases of the bitter axis, the ANR and LAR reductases of the astringent axis in Chapter 8 — we build a high-confidence structural model with a folding network and score every residue by local confidence (pLDDT) and by the predicted aligned error between domains (PAE). We want to know, before we touch anything, which parts of the protein are rigid scaffold, which are the flexible active-site loops, and where the substrate and the SAM methyl donor actually sit. The figure below is the kind of confidence map we read: warm where the model is sure, cool where it hedges. The active-site cleft and the cofactor pocket are the high-confidence regions we design into; the floppy termini we leave alone.

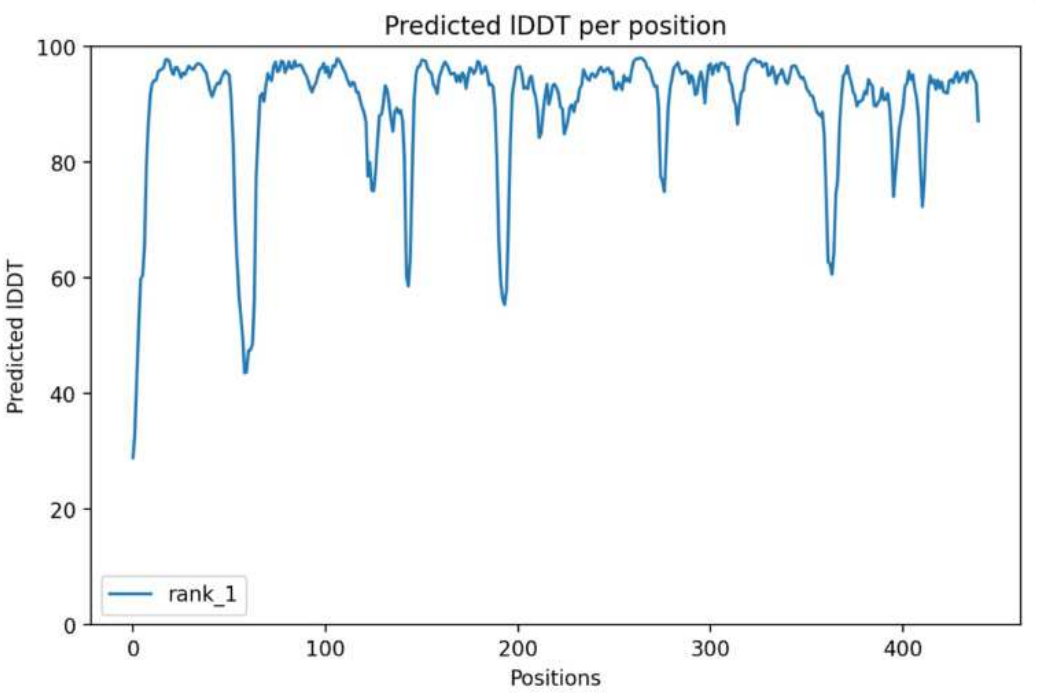

FIG. 6 — a platform fold model coloured by per-residue confidence (pLDDT). The rigid core and the active-site cleft resolve at high confidence; this is the map we design the edit onto, now read on the cacao methyltransferase rather than the pepper acyltransferase that first proved the method.

Second, the language model. We score every possible single-residue substitution in the target enzyme with a protein language model — an ESM2-class network of hundreds of millions of parameters, trained on the grammar of millions of natural proteins. The model has never been told what these enzymes *do*; it has only learned what proteins *are*, and from that it can estimate, for every position and every possible amino acid, how surprising — how

destabilizing, how unnatural — that substitution would be. We combine that evolutionary-grammar score with an explicit **ΔΔG** calculation of fold stability. The product, catalytic plausibility × fold stability, ranks the entire mutational landscape of the enzyme. For a de-bittering edit we are hunting substitutions that *reduce catalytic efficiency in a controlled way* — that throttle the methyl transfer without collapsing the protein — a partial, tunable loss of function, not a knockout. For an astringency edit we hunt the opposite on a different enzyme. The language model turns a combinatorial impossibility — a 400-residue enzyme has thousands of single mutants and astronomically many combinations — into a ranked shortlist a human can reason about.

It is worth saying what the language model is doing, because it is the strangest and most powerful step. An ESM2-class network is trained by **masked language modeling** on a corpus of hundreds of millions of natural protein sequences — shown a sequence with residues hidden, asked to predict them — and in learning that game it internalizes the deep grammar of what makes a protein foldable and functional: the correlations between distant residues, the positions that are conserved and the positions that are free. It is never told a single enzyme's *function*. Yet because an unnatural, destabilizing substitution is precisely the kind the model finds "surprising," the **log-likelihood ratio** it assigns a variant — how much less probable the mutant sequence is than the wild type — turns out to be a remarkable *zero-shot* predictor of that mutation's effect, with no experimental training data at all. We run the 650-million-parameter model across every position and every substitution, and fold its evolutionary-grammar score together with an explicit **ΔΔG** of fold stability computed from the structure. Evolutionary plausibility catches the mutations that break *function*; the ΔΔG catches the ones that break the *fold*; their product ranks the whole landscape by the thing we actually want — a substitution that changes catalytic throughput by a controlled amount while leaving the enzyme standing. For a de-bittering edit we read *down* the ranking toward a partial, tunable loss of efficiency — a throttle, never a knockout — and for an astringency edit, on a different enzyme, we read the other way.

Third, the flux check. A better — or deliberately throttled — enzyme is useless if it is not the step that controls the outcome, which returns us to the genome-scale flux-balance model of Chapter 5. Before committing an edit we confirm in the metabolic model that our target enzyme is a true control point for the end-product pool, that throttling it moves the bitterness by the amount we want, and that it does not starve some pathway we care about — the aromatic precursors, the fat, the flavanols we intend to *keep*. This is where the coupling of the three axes gets enforced: the model will warn us if our bitter edit collaterally strips the phenolics we are tuning for astringency, because in cacao those branches share carbon. Design is not a list of independent edits; it is a single consistent solution to a coupled system.

Fourth, the write. The surviving edits — ranked by the model, confirmed by the flux check — are written into the genome with CRISPR. Precise, targeted substitutions at the residues the model named. Then the plant is grown and the prediction is *measured*: the methylxanthine budget by mass spectrometry, the bitterness by the TAS2R38-anchored panel, the whole eight-axis profile by the index. And the number that matters as much as the mean is the **coefficient of variation**: can we hit the target *every time*? The heat platform drove batch-to-batch CV from 42% in conventionally bred material down to **≤ 7%** with the edited constructs. Reproducibility is not a nicety for a formulator; it is the whole product. A chocolatier who orders a cacao at IFI-B 1,500 must receive IFI-B 1,500 in February and in August, in this harvest and the next.

"Measured" means specific assays, and they are worth naming, because they are how a flavor becomes a fact. The methylxanthine budget comes off **LC-MS/MS** — liquid chromatography with tandem mass spectrometry — which resolves and quantifies theobromine and caffeine to the microgram, so the bitter *load* is a number before it is ever a taste. The condensed tannins are characterized by **HPLC after phloroglucinolysis**, the standard depolymerization that cleaves the procyanidins and lets us read both their total mass *and* their **mean degree of polymerization** — the astringent dial of Chapter 8 — directly. The flavanol pool we mean to preserve is quantified the same way, so "we lowered astringency without destroying flavanols" is a claim with two independent chromatographic numbers beneath it, not a hope. The perceived axes are read by a **trained descriptive panel** — assessors calibrated against reference standards for bitterness and astringency intensity, scoring to the TAS2R38-anchored and PRP-anchored scales — because the receptor assay and the human mouth must agree before anything ships. And finally the statistic that is the whole product: the **coefficient of variation**, the batch-to-batch standard deviation over the mean, computed across independent grow-outs. The heat platform drove CV from 42% in conventionally bred material to **≤ 7%** with the edited constructs, and that number is not a flourish — it is the line between a research curiosity and a thing a chocolatier can build a product on. Reproducibility *is* the deliverable.

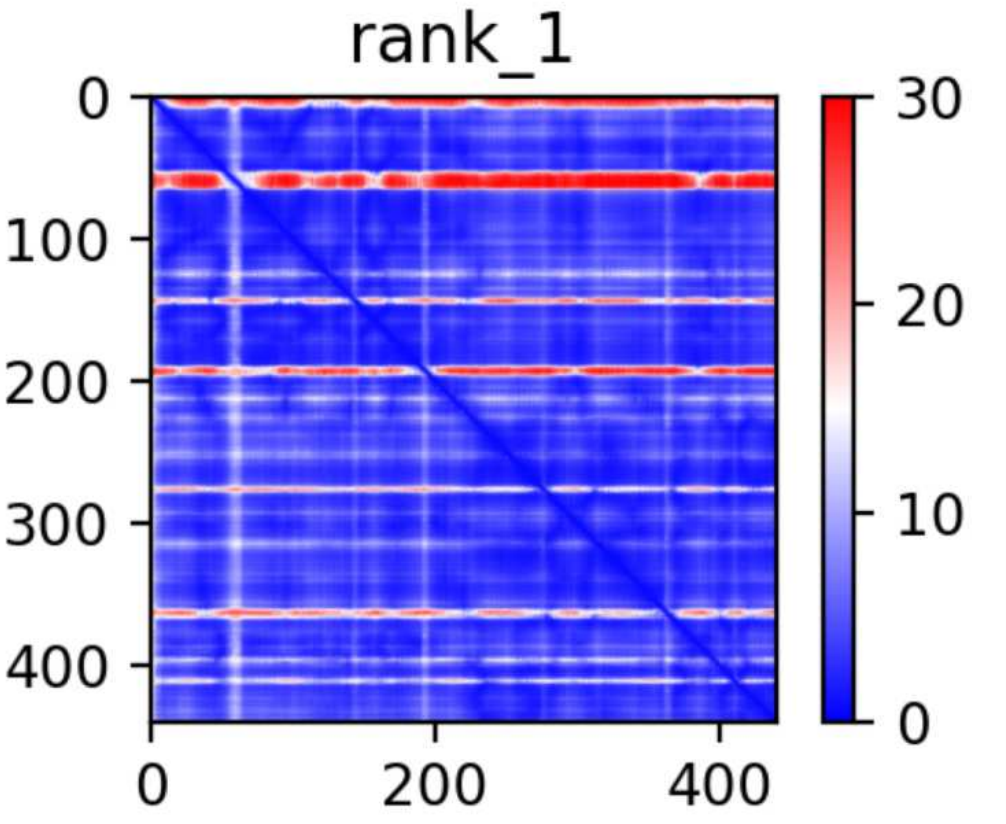

FIG. 6b — the predicted aligned error between domains. Low error (cool) means the model is confident in how the active-site region sits relative to the scaffold — the geometry an edit must preserve.

What does the output of this pipeline look like for cacao's bitter axis? It looks like a short table of candidate substitutions on the theobromine-pathway methyltransferases — the model's top-ranked edits for a controlled throttle of flux. We present these as exactly what they are: **predicted design targets**, the calculation we would write and then measure, not yet a result off a bench. The values are the model's predicted ΔΔG of fold stability (negative is stabilizing); the intent in each case is a partial, tunable reduction of catalytic throughput, not a knockout.

TARGET ENZYME	CANDIDATE EDIT	PREDICTED EFFECT	PREDICTED ΔΔG (KCAL/MOL)
XMT (pathway entrance)	active-site loop substitution	throttles total methylxanthine flux ~30–40%	−1.4
XMT	SAM-pocket margin substitution	fine control of entrance flux, fold preserved	−0.8
Theobromine synthase	catalytic-cleft substitution	shifts theobromine:caffeine ratio toward mellow	−1.1

Compare this, for method (not for target), with the edits that founded the company — the Pun1 double variant **S39L (ΔΔG −2.1) + L345G (ΔΔG −1.8)**, the two substitutions in the shipped construct, that our model predicts lifts a pepper to approximately 2.95 million SHU at CV ≤ 7%. (C175S, ΔΔG −0.9, is a structural stability annotation from our modelling and was not included in the final design.) Two substitutions no breeder ever found by phenotype, because you cannot select your way to a specific combination of residues across a protein; you have to *read* them off the model and *write* them in. The cacao edits above are the same kind of object — a small, legible set of residue changes a model ranked out of an impossible space — aimed at a different enzyme toward a gentler end. When cacao finally sits on our bench, this is the first calculation we will run, and the first thing we will measure against.



PART III OF VI

The Astringent Dimension

IFI-A · salivary PRPs · procyanidins · ANR · LAR · the degree of polymerization

Chapter 7 · The cross-linked mouth — *procyanidins and the PRP cascade.*

Astringency is the axis this book is accented for — plum, the tannin colour — and it is the axis cacao speaks most natively and least evenly. A great dark chocolate's astringency is one of its glories: the structured grip that makes the mouth water in anticipation, the architecture a long finish is built on. A clumsy cacao's astringency is one of its worst faults: a green, aggressive, mouth-stripping harshness that never resolves. The difference between the two is not *how much* tannin but *what kind*, and to tune it we have to understand astringency as the unusual sensation it is — the one place in the index where flavor is really touch.



Grape, halved — the organism where we first learned to sculpt tannin. 1,610 constructs on the astringent and bitter axes; the ANR enzyme we tuned here is the same one cacao needs.

The sensation first. Your saliva is not just water; it is a lubricant, and the molecules doing the lubricating are a family of **proline-rich proteins** — PRPs — along with the phosphoprotein statherin, flexible, heavily glycosylated chains whose job is to coat oral surfaces and let your tongue glide frictionlessly over your palate and teeth. Tannins — condensed polyphenols — have a specific and ancient affinity for proline. Their phenolic hydroxyls hydrogen-bond to the carbonyls of the proline-rich backbone, and their flat aromatic faces stack against the proline rings, and because each tannin polymer

carries *many* such binding faces, a single tannin can bridge *several* PRP molecules at once. The result is cross-linking: the lubricating proteins clump, aggregate, and precipitate out of solution. The film fails. For a few seconds, until fresh saliva rebuilds the layer, the raw surfaces of your mouth drag against each other, and your brain reads that friction as *dry, rough, puckering, astringent.*

The binding deserves a closer look, because every handle we have on astringency is a handle on *this* interaction. Proline is the oddest of the amino acids — its side chain loops back and bonds its own backbone nitrogen, making a rigid ring that forces kinks into a protein and leaves its backbone carbonyls exposed. Tannins exploit exactly that: their flat aromatic rings stack face-to-face against proline's pyrrolidine ring (a hydrophobic association) while their phenolic hydroxyls hydrogen-bond to those exposed carbonyls. Proline-rich proteins are, in effect, *evolved tannin sponges* — flexible, open, proline-studded chains that seem almost designed to be sacrificed to polyphenols, and may well be: a first line of defense that binds dietary tannins before they can do worse. The interaction proceeds in three stages, and each is a lever. First, **binding** — tannin decorates the PRP. Then, as more tannin loads on, the coated proteins begin to **aggregate**, cross-linked into larger complexes by multivalent tannins bridging several chains at once. Finally, past a solubility limit, the aggregates **precipitate** out of the saliva entirely — and it is the precipitation, the sudden loss of lubricant, that the mouth reads as astringency.

Because it is a physical precipitation and not a receptor event, astringency bends to conditions a taste never would. It sharpens at low pH (a more acidic chocolate or wine grips harder); it shifts with temperature; ethanol modulates it; and — critically — it depends on *your own saliva*. People vary considerably in how much PRP they secrete and how fast they flow it, which is a real part of why one taster's "structured and fine" is another's "harshly astringent." A fast, protein-rich salivator rebuilds the lubricating film quickly and reads a tannic bar as gentler; a slow one feels every polymer. The index sidesteps the variation by scoring against a standardized precipitation assay rather than any single mouth — but a chocolatier should know that astringency is the one axis where the instrument and the eater most need reconciling, and that pacing, temperature, and what you drink alongside are all, physically, astringency controls.

Three consequences fall straight out of that mechanism, and each one is a design handle.

Astringency is cumulative. Because it works by precipitating a finite pool of lubricating protein, each sip or bite depletes the film a little more before saliva can fully restore it. A taste fades between bites; astringency *builds*. This is why a tannic wine or a tannic chocolate feels more astringent on the fifth taste than the first, and why pacing and palate-cleansing matter. On the index it means we score astringency against a standardized single-exposure precipitation assay, so the number is clean, while knowing the mouth will integrate it over a tasting.

Astringency scales with polymer size. A flavan-3-ol *monomer* — a lone epicatechin — has only one or two binding faces and barely cross-links; it is mildly bitter and almost not astringent. A *dimer* (procyanidin B2, two epicatechins joined 4β→8) grips noticeably. By the time you reach the longer procyanidins — five, eight, ten units — a single molecule can bridge a whole raft of PRPs, and the astringency is intense. So the astringent intensity of a cacao tracks its **mean degree of polymerization**, its mDP: the average chain length of its condensed tannins. *This is the dial.* Not the total phenolic content — the chain length distribution.

Astringency has sub-qualities. Galloylation — a gallate ester hung on the flavanol — sharply increases PRP affinity and produces a harder, more aggressive, "coarser" astringency; cacao is lightly galloylated, which is part of its gentler grip relative to, say, a tannic tea. And the balance between the 2,3-*trans* catechin and the 2,3-*cis* epicatechin stereochemistry subtly shapes the quality of the grip. Astringency is not one slider but a small console: total tannin, mean chain length, galloylation, stereochemistry. A great cacao is a specific setting across all four.

Trained panels do not even score astringency as a single intensity; they score a small vocabulary of sub-qualities, and each maps back to tannin structure in a way a designer can use. *Puckering* and *drawing* come from the most aggressive, multi-PRP-bridging polymers — the long procyanidins that precipitate protein fastest. *Drying* is the broad, general loss of lubrication. *Rough* and *grainy* describe aggregates coarse enough to feel as texture. *Chalky* is a particular fine precipitate; *velvety* or *fine-grained* is the prized quality of a well-made, well-resolved astringency — the grip of a great aged thing that the mouth reads as structure rather than assault. And astringency has a *time course* the panel scores explicitly: an onset (how fast the grip arrives), a peak, and a carryover (how long it lingers and builds across successive tastes). A young, harsh cacao grips fast and stays, accumulating into something punishing by the third square; a fine one arrives in measured time, peaks cleanly, and resolves, leaving the mouth watering for the next. All of it — the sub-quality and the dynamics — is governed by the polymer distribution we learn to set in the next chapter: shorter, better-resolved chains for a fine, velvety grip that clears; a tuned fraction of longer chains for structure and length; galloylation held low to keep the quality smooth rather than coarse. Read this finely, astringency is not one number after all — it is a little spectrum, and a great cacao is a specific, deliberate shape across it.

Now the molecules themselves, because cacao is a textbook. The monomers are **(–)-epicatechin** (dominant, the 2,3-*cis* form) and **(+)-catechin** (the 2,3-*trans* form). The oligomers are the **procyanidins** — B-type dimers like procyanidin B2 and B5, the trimer C1, on up through the decamers — built by condensation of the monomer units. Collectively these condensed tannins can reach several percent of the raw seed's dry weight, which is a very large astringent cargo, and most of the craft of chocolate's early steps — fermentation especially — is spent beating that cargo down to something a mouth can love. In structural terms cacao's

procyanidins are predominantly **B-type** (single interflavan bonds) built mostly on epicatechin, with a native mean degree of polymerization that runs well into the oligomer range — which is only to say, in the language of the next chapter, that the raw seed sits high on IFI-A with its chain-length distribution skewed toward the long, gripping end, exactly the distribution we propose to resculpt. We mean to beat it down more precisely, and then, for the bars that want it, to build it back up with *structure* rather than just *less*.

The through-line of the whole book lives in this chapter too, quietly, because these same procyanidins and epicatechins are cacao's famous flavanols — the molecules a certain kind of eater seeks cacao out for. The craft steps that lower astringency do it bluntly, by *destroying* flavanols: fermentation oxidizes them, roasting degrades them, Dutch-process alkalization annihilates them. The astringency falls and the flavanol cargo falls with it, because the blunt instrument cannot tell the astringent long polymers from the gentle short ones. A genetic lever can. If we tune astringency by shifting the *polymer length distribution* — fewer mouth-stripping decamers, more gentle monomers and dimers — we can lower the *grip* while *preserving*, even concentrating, the flavanol mass. We make no claim about what those flavanols do in a body; the voice of this library does not. We only observe that it would be a shame to throw them away by accident, and that the whole point of a scalpel over a fire is that you need not.

Chapter 8 · Tuning the tannin — *ANR, LAR, and the degree of polymerization.*

We know the sensation, the molecules, and the dial — mean degree of polymerization. Now the pathway, because to set the mDP we have to reach into the enzymes that build the flavan-3-ols and decide how much, and how long, the seed polymerizes. This is the branch of plant metabolism we know best of all, because it is where the astringent half of the grape library — sixteen hundred constructs — already lives.

The flavan-3-ols come off the flavonoid branch, which comes off the phenylpropanoid branch, which is the universal aromatic-metabolism highway running in every vascular plant. The shared trunk is the same one the capsaicin platform uses — **PAL → C4H → 4CL** turns phenylalanine into p-coumaroyl-CoA — and from there the flavonoid branch builds the ring system and, at its business end, two reductases decide the flavan-3-ol outcome. Here is the route, with the enzymes named and the intermediates in place:

STEP	SUBSTRATE → PRODUCT	ENZYME
trunk	Phe → p-coumaroyl-CoA	PAL → C4H → 4CL
1	p-coumaroyl-CoA (+ 3 malonyl-CoA) → naringenin chalcone	CHS (chalcone synthase)
2	naringenin chalcone → naringenin	CHI (chalcone isomerase)
3	naringenin → dihydroflavonols	F3H (+ F3′H hydroxylation)
4	dihydroflavonol → leucoanthocyanidin	DFR (dihydroflavonol-4-reductase)
5a	leucoanthocyanidin → **(+)-catechin**	**LAR** (leucoanthocyanidin reductase)
5b	leucoanthocyanidin → anthocyanidin	ANS/LDOX
6	anthocyanidin → **(−)-epicatechin**	**ANR** (anthocyanidin reductase; the *BANYULS* enzyme)
7	catechin / epicatechin → **procyanidins**	oxidative/enzymatic condensation; polymer extension

Two enzymes carry the design. **LAR** makes catechin (the 2,3-*trans* monomer). **ANR** — anthocyanidin reductase, the enzyme encoded by the gene genetics first found in an Arabidopsis mutant and named *BANYULS* after the sweet tannic wine — makes epicatechin (the 2,3-*cis* monomer), which is the dominant flavanol of cacao and the principal building block of its procyanidins. The ratio of LAR to ANR flux sets the catechin:epicatechin balance, and therefore part of the

quality of the astringency. And the overall flux into step 7, together with whatever governs chain extension, sets the *mean degree of polymerization* — the master dial for astringent *intensity*.

The enzymes reward a closer look, because they tell us which dials are truly ours to turn. DFR, LAR, and ANR are all **NADPH-dependent reductases** — they spend reducing power to build the flavan-3-ol skeleton — and ANR in particular carries a discovery story that doubles as a lesson. It was first found not in a winery but in a mutant of the model weed *Arabidopsis*, whose seed coat failed to make its normal brown tannin and so let the pale seed show through: a *transparent testa* mutant. The gene behind it was christened **BANYULS**, after the sweet, deeply tannic fortified wine of the French coast, by geneticists with a palate and a sense of humor — and it turned out to encode anthocyanidin reductase, the very enzyme that makes cacao's dominant epicatechin. When we tune ANR in cacao, we are turning a knob that plant genetics first found by watching a seed go pale.

But the degree-of-polymerization dial — the master control on intensity — is not only a biosynthesis problem, and honesty about that shapes the engineering. The flavan-3-ol monomers are made in the cytosol, then hauled into the vacuole by dedicated **transporters** (MATE-family pumps and ABC / glutathione-S-transferase machinery — the *TT12* and *TT19* gene products in the Arabidopsis story), and it is largely *in the vacuole* that they condense into procyanidins, through a mix of enzyme-assisted and spontaneous oxidative coupling whose details the field is still resolving. The interflavan bonds themselves come in types — the common **B-type** (a single C4→C8 or C4→C6 link) and the rarer, doubly bonded **A-type** — and the balance of linkages, together with chain length and galloylation, sets the astringent *quality*. What this means for us is that shifting the mean degree of polymerization is a coordinated edit across *several* nodes at once — the reductase flux that sets monomer supply, the transport that sets vacuolar concentration, and the oxidative coupling that sets chain extension — which is exactly the kind of multi-gene, whole-pathway problem the flux model and the sixteen hundred grape constructs exist to solve. Grape tunes this machinery for wine; cacao would tune the same machinery for a bar.



Tea — catechin, caffeine, and theaflavin all at once. 280 constructs across umami, bitter, and astringent axes; the astringency grammar we proved here reads straight onto cacao's procyanidins.

So the astringent strategy has two moves, and a chocolate chooses between them.

To soften a harsh cacao — to take a green, aggressive, mouth-stripping grip and make it gentle without going flabby — we shift the polymer distribution *shorter*. Modulate the condensation/extension at step 7 so the seed accumulates more monomers and dimers and fewer long, multi-PRP-bridging decamers. The *total* flavanol mass can stay high — this is the flavanol-preservation move of Chapter 7 — while the *astringent intensity* falls, because intensity tracks chain length, not mass. The result is a cacao with all its phenolic cargo and a fraction of its bite.

To give a flabby cacao structure — to take a dull, sweet-leaning bean and build it a backbone — we do the opposite: push flux toward longer, well-formed procyanidins that grip with *structure*, a resolving astringency rather than a crude one, the tannic architecture a serious dark bar stands its finish on. Astringency, uniquely among our axes, is one we often tune *up*. Most of the library is about turning shouting axes down. Here, for the connoisseur's bar, we turn a quiet one up — toward grip, toward length, toward the thing that makes the mouth water for the next square.

We can calibrate the whole axis against an mDP ladder, because the sensation and the chain length march together:

MEAN DEGREE OF POLYMERIZATION (MDP)	ASTRINGENT CHARACTER	IFI-A (TARGET)
~1–2 (monomer/dimer dominant)	soft, barely gripping; flavanol-rich, low bite	~600
~3–4	gentle structure; mouth-watering, resolves fast	~1,200
~5–6	firm, architectural grip; a serious dark bar's backbone	~1,800
~7–10 (long polymers)	aggressive, slow-resolving, green if unripe	~2,800

And we are not guessing that this works, because tea and grape and the berry astringents have already walked it. **Grape**, 1,610 constructs, tunes exactly this ANR/LAR node for wine. **Tea**, 280 constructs, carries catechins and the oxidized theaflavins and taught us the astringent grammar on an evergreen leaf. **Pomegranate** (120, the ellagitannin route), **blueberry** (120), **sorghum** (120, the frontier tannin grain), **raspberry** (40) — a whole astringent neighborhood, each a condensed- or hydrolyzable-tannin system, each tuned on the same salivary-PRP physics. Cacao's procyanidins are not exotic to us. They are the family business.

Two finer dials deserve a mention, because they set astringent *quality* beyond mere chain length. The first is **galloylation** — the hanging of a gallate ester onto a flavan-3-ol, catalyzed by SCPL acyltransferases — which sharply raises PRP affinity and coarsens the grip; cacao is only lightly galloylated, part of why its astringency is gentler than a heavily galloylated green tea's, and holding galloylation low is a lever for keeping the quality smooth. The second is **oxidation**. In black tea, the catechins are deliberately oxidized by

endogenous polyphenol oxidase into **theaflavins** and **thearubigins** — larger, differently astringent, differently colored polymers — which is the whole difference between green tea and black, and it is the same PPO chemistry that fermentation runs, cruder, on cacao. Tea, then, taught us astringency in *both* directions: the reductive building of catechins and the oxidative building of theaflavins, 280 constructs across the two. For cacao it means the astringent toolkit is not only "how long the procyanidin" but "how galloylated" and "how oxidized" — three structural dials under one axis, each reachable at an enzyme, each already turned in a neighbor. When we call the grape and tea programs a rehearsal, this is the granularity we mean: not a vague precedent, but the specific enzymes — ANR, LAR, the SCPL acyltransferases, the oxidases — that cacao shares with plants we have already tuned to the number.



Pomegranate — a different tannin chemistry (ellagitannins), the same astringent physics. 120 constructs; another rehearsal for the cacao grip.

One honest caveat, because the astringent axis has a coupling we flagged in Chapter 4 and must respect: the phenolic branch and the bitter phenolics overlap, and some of the smaller flavan-3-ols are *both* mildly bitter *and* astringent. When we shift the polymer distribution shorter to soften the grip, we can slightly raise the monomer pool, and monomers carry a little bitterness. The flux model holds both axes at once precisely so we do not solve astringency and accidentally un-solve bitterness. In cacao more than anywhere in the library, the axes are a system, not a list — which is why we cannot finish the astringent chapter without the one that rounds it all: the body.



PART IV OF VI

Body, and the Whole Chord

IFI-K · CaSR · γ-glutamyl peptides · fermentation · the coupled design · bean to bar

Chapter 9 · Kokumi — *the dimension you feel but cannot name.*

If bitterness is the wall and astringency is the grip, kokumi is the thing that makes the room worth standing in. It is the hardest of cacao's three axes to talk about, because it is the one with no everyday word, and it is the easiest to prove, because the moment you take it away, everyone at the table notices that *something* left. Kokumi is body, continuity, mouthfulness — the sensation that a flavor is full and present and *lasts*, that it fills in the middle and goes long rather than flashing and vanishing. A 70% dark with good kokumi eats richer and rounder than an 85% without it. Most of the gap between a merely-correct chocolate and a profound one lives on this axis.



Vanilla — the kokumi neighbor, and chocolate's oldest companion. 40 constructs on the CaSR axis; the body we would build into cacao, we first built here.

The mechanism is specific and, once you know it, a little thrilling. Kokumi compounds are agonists of **CaSR, the calcium-sensing receptor** — a G-protein-coupled receptor best known from the parathyroid, where it monitors blood calcium, and present, it turns out, on the cells of the tongue. CaSR is not a taste receptor in the classical sense; it does not generate a "kokumi taste" the way TAS1R2 generates sweetness. Instead, when CaSR is engaged by the right ligand, it *modulates* the output of the neighboring taste cells — it turns up the gain on sweet, on umami, on salty, and it dampens the perceived intensity of bitter. Kokumi is a meta-taste: it does not add a flavor, it *conducts* the ones already present, pulling them into fullness and continuity. That is why it feels like body rather than like a note.

The receptor itself is worth meeting, because its structure explains why kokumi behaves so unlike a taste. CaSR is a **class C G-protein-coupled receptor** — the same unusual family as the sweet and umami receptors — and like them it carries a large extracellular **"Venus flytrap" domain**, two lobes that close around a ligand like a trap snapping shut, mounted on the seven-helix membrane bundle that does the signaling. It works as a **dimer**, two subunits clasped together, and it is promiscuous by design: its primary job in the body is to sense extracellular *calcium* (hence the name), but it is tuned up and down by amino acids, polyamines, pH, and — the part we care about — by γ-glutamyl peptides, which act as **positive allosteric**

modulators, sensitizing the trap so it responds more readily to whatever else is around. On the tongue, a CaSR engaged this way does not open a "kokumi taste" line; it modulates the neighboring sweet, umami, and bitter cells through Gq/G_{11} signaling, turning their gain up or down. That is the molecular meaning of body: not a flavor, but a hand on the mixing board.

The discovery deserves honoring. Kokumi began as a kitchen intuition — that aged garlic, certain long-fermented pastes, and a properly built stock carried a richness none of their measurable tastes explained — and was chased down, over decades and largely by Japanese flavor science, to glutathione, then to CaSR, and finally to the designed tripeptide **γ-Glu-Val-Gly**, among the most potent kokumi agents known, active where the earlier peptides needed far more. The structure-activity story is clean, and it tells us exactly what to build: the **γ-glutamyl head is essential** — move the bond from the gamma carboxyl to the ordinary alpha and the kokumi activity largely vanishes — while the identity of the trailing residues tunes potency, with small hydrophobic partners (valine above all) giving the strongest effect. So the engineering target is specific: not "more peptides" but *more of the γ-Glu-X peptides with the right tail* — which is a transpeptidation-flux problem, and transpeptidation is exactly what garlic does for a living.

The ligands that do this best are **γ-glutamyl peptides** — short peptides carrying a glutamate joined through its *gamma* carboxyl (the unusual isopeptide bond) rather than the ordinary alpha. The archetype is **glutathione**, γ-Glu-Cys-Gly, the cell's own antioxidant tripeptide, a real kokumi agent. Even more potent are designed and natural relatives — **γ-Glu-Val-Gly**, active at sub-micromolar concentrations, among the most powerful kokumi peptides known; γ-Glu-Val, γ-Glu-Cys, γ-Glu-Met. Their potency is the reason kokumi is such an efficient axis: because CaSR hears these peptides at sub-micromolar levels, a *few milligrams* of the right peptide across a whole batch can lift the body of a chocolate measurably. On the index, that extreme potency — a very low EC_{50} against CaSR — is what earns a kokumi construct a high IFI-K from a tiny molar quantity.

Where do γ-glutamyl peptides come from, and how would we build them into cacao? The cell already makes the machinery. **γ-glutamylcysteine ligase** (GCL) joins glutamate to cysteine through the gamma bond; **glutathione synthetase** adds glycine to complete glutathione; and a family of **γ-glutamyl transferases** (GGT) transpeptidate — they move the γ-glutamyl group onto other amino acids, generating the whole family of γ-Glu-X peptides. The kokumi lever is to tune this peptide-synthesis and transpeptidation flux so the seed accumulates the potent peptides. And cacao, encouragingly, already generates peptides during fermentation — the proteolysis that breaks the seed's storage proteins throws off exactly the kind of short peptides and free amino acids that feed both the Maillard reaction and, potentially, the kokumi pool. We would be amplifying and directing a peptide chemistry the bean already runs.

We have built this axis before, in the two most body-giving plants in the kitchen. **Vanilla** — forty constructs — carries kokumi mouthfeel alongside its vanillin, which is half of why real vanilla makes everything it touches taste rounder and more complete (and why chocolate and vanilla have been married for four centuries). **Garlic** — five hundred and sixty constructs — is a γ-glutamyl *powerhouse*; raw garlic is naturally rich in γ-glutamyl-S-allyl-cysteine and its relatives, and tuning that peptide chemistry is most of garlic's kokumi contribution to the savory foundation of nearly every cuisine. **Onion** (40), **yeast** (240, umami and kokumi together), **soybean** (360, the frontier of savory depth) — the CaSR axis is a neighborhood we have built in repeatedly. None of them is cacao. All of them are proof the peptide lever lands.

It helps to know where else kokumi hides, because the list is a map of the world's most *satisfying* foods. Aged cheeses carry it — the long proteolysis of ripening throws off γ-glutamyl peptides — and so do scallops and dried scallop, fish sauce and soy sauce, long-simmered stocks, yeast extract, aged and fermented soy pastes, and, the kitchen's oldest kokumi trick, garlic, especially aged. The common thread is *time and enzymes*: wherever proteins are slowly cleaved by proteases and transpeptidases, the γ-glutamyl peptides accumulate, and the food gains body. Which is the quietly encouraging news for cacao, because the bean already runs exactly that chemistry. Fermentation's proteolysis cleaves the seed's storage proteins into peptides and free amino acids (Chapter 10), and some of those are precisely the short peptides kokumi is built from; roasting transforms the pool further. A well-fermented, well-roasted cacao is therefore *already* carrying some native kokumi — part of why a great bean eats rounder than a poor one at the same cocoa percentage. So the kokumi lever in cacao is, like the others, an *amplification* of a chemistry the seed already has, not an alien insertion: tune the γ-glutamyl transpeptidation flux so the seed accumulates more of the potent peptides, and direct fermentation to finish the job. We would be dialing up a dimension the bean already reaches for on its own — with garlic's 560 constructs and vanilla's forty as the proof that the dial turns.

Now the cross-modal payoff, which is why kokumi belongs structurally between bitterness and astringency and not off in a corner. Because CaSR engagement *dampens perceived bitterness*, kokumi is a bitterness mask that removes nothing — the "mask it" lever from Chapter 5, made real. A cacao designed with a strong kokumi profile will *read* as less bitter than its methylxanthine budget predicts, because the body is conducting the bite down even though every molecule of theobromine is still present and still, we would gently note, doing its own quiet good. This is the deepest reason kokumi is cacao's secret dimension: it is the only one of the three dials that improves a bar by changing *perception* rather than *composition*, which means it costs the seed nothing it was carrying and takes nothing away — pure gain, if you can build it. This means the bitter axis (IFI-B, true attenuation at the enzyme) and the kokumi axis (IFI-K, perceptual rounding at CaSR) are two different roads to a smoother chocolate, and the best designs drive partway down both:

a modest real de-bittering, plus a lift of body that rounds what remains. The chocolate is less bitter *and* richer, and no aromatic was sacrificed to get there. That is the chord. Let us learn to play all three notes at once — but first, let us meet the engineer who has been playing them, crudely and beautifully, for ten thousand years.

Chapter 10 · Fermentation, the first engineer — *what the microbes already do.*

Before any gene is written, cacao passes through the one transformation without which chocolate simply does not exist: fermentation. It is the step the uninitiated skip over and the makers obsess about, and for our purposes it is something close to sacred, because fermentation is *precisely* a flavor-engineering program aimed at cacao's three axes — bitter down, astringent down, body and aroma up — run by a succession of microbes with no model and no CRISPR and five thousand years of empirical refinement. To understand what we would engineer, you must first honor what fermentation already does, and see exactly where its blunt genius runs out.



A cacao pod and its seeds in the pale, sweet pulp — the sugar the microbes come for, the morning the fermentation begins.

Here is the drama, in three acts, all of it real biochemistry. The farmer heaps the fresh seeds — slick in that bright white pulp — into boxes or baskets or banana-leaf mounds and lets the wild microbes come.

Act one, the yeasts. The pulp is sugary, acidic, low in oxygen — a yeast's paradise. Wild *Saccharomyces* and its relatives bloom and ferment the pulp sugars to **ethanol**, throwing off carbon dioxide and beginning to break the pulp down so it drains away as a liquid the makers call "sweatings." The yeasts also produce the first aromatic

esters and higher alcohols — the faint fruity, floral, winey notes that survive all the way into a well-made bar. Pectinolytic enzymes liquefy the pulp; air begins to reach the mass.

Act two, the acids. As oxygen enters and sugar falls, lactic-acid bacteria and then **acetic-acid bacteria** take over. The acetic-acid bacteria oxidize the yeasts' ethanol to **acetic acid**, and this reaction is strongly *exothermic* — the heap heats itself, climbing to 45–50 °C over a day or two. Now two things happen at once that change everything. The heat and the penetrating acetic acid **kill the seed** — the embryo dies, dormancy ends, and the cotyledon's cell compartments break down, which lets contents that were kept apart finally mix. And the acid and heat **activate the seed's own enzymes.**

Act three, the seed's own enzymes — this is the part that matters most for us. With the cells broken and the compartments mixed, cacao's endogenous enzymes go to work on its own tissue:

- **Proteases** — an aspartic protease and a carboxypeptidase working in sequence — cleave the seed's storage proteins into a specific pool of **hydrophobic peptides and free amino acids**. These are the Maillard precursors. Without this proteolysis there is no chocolate aroma after the roast, only a burnt, flat smell. This is also, note, the step that generates the peptide pool kokumi could be built from.
- **Invertase** splits sucrose into the reducing sugars (glucose, fructose) that the roast will need as the other half of the Maillard reaction.
- **Glycosidases** cleave sugar groups off bound flavor precursors, freeing aromatics.
- **Polyphenol oxidase** (PPO) oxidizes and polymerizes the flavan-3-ols, which **browns** the seed and — crucially — **reduces its astringency and bitterness**, because the oxidized, bound polyphenols no longer precipitate saliva the way the free ones did. Meanwhile the simple diffusion of the water-soluble purines and phenolics out of the cotyledon, carried in the draining acid, leaches a real fraction of the bitterant and astringent cargo away entirely.

It is worth naming the cast, because fermentation is a microbial relay as choreographed as any pathway. The opening belongs to **yeasts** — *Hanseniaspora* (*Kloeckera*) early, then *Saccharomyces cerevisiae* and kin — working the sugary, acidic, low-oxygen pulp into ethanol and the first fruity esters. As the pulp drains and air enters, **lactic-acid bacteria** (*Lactiplantibacillus*, *Lactobacillus*) take their turn, and then the decisive players arrive — the **acetic-acid bacteria** (*Acetobacter*, *Gluconobacter*), which oxidize the ethanol to acetic acid in that self-heating, exothermic burst. The whole relay runs differently by vessel: Mesoamerican and West African **heap** fermentations under banana leaves, wooden **box** fermentations turned on a schedule, shallow **basket** and **tray** methods — and each geometry sets its own oxygen, drainage, and heat, which is part of why terroir in chocolate is as much *microbial* terroir as it is soil.

PHASE	~HOURS	DOMINANT MICROBES	WHAT HAPPENS TO THE THREE DIALS
Anaerobic, sugary	0–36	yeasts (*Hanseniaspora*, *Saccharomyces*)	pulp → ethanol + esters; the mass begins to drain
Transitional	24–72	lactic-acid bacteria	acids build; oxygen enters as pulp liquefies
Aerobic, heating	48–120+	acetic-acid bacteria	ethanol → acetic acid, exothermic; heap hits 45–50 °C; **the seed dies**
Enzymatic (within the seed)	48–168	the seed's own enzymes	proteolysis builds aroma + peptide precursors; PPO + leaching take **B and A down**

The window is everything, and it is why the farmer's judgment is the real instrument. **Under-ferment** — pull it too early — and the seed never dies cleanly, the enzymes never fully fire, and the bean stays brutally bitter and astringent, grassy and sharp, purple inside instead of brown. **Over-ferment** — leave it too long — and the relay runs past flavor into putrefaction: ammoniacal, hammy, hollow, the aromatics consumed. Between those two failures lies a target the farmer hits by smell, by heat, and by the color of a cut bean, with a batch-to-batch variability in *flavor* that is the despair of every fine maker and nowhere near our ≤ 7%.

Fermentation's twin, and its quiet accomplice, is **drying** — the step that follows, and that half of flavor silently depends on. The fermented beans, still wet and alive with acid, are spread to dry, traditionally in the sun on raised beds or patios, over one to two weeks. Drying does more than guard against mold (the target is roughly six to seven percent moisture — too wet and the beans rot, too dry too fast and the interior stays sour because the acids never finish off-gassing). As the beans dry slowly, residual enzymatic browning continues, more of the sharp acetic acid volatilizes away, and astringency and harshness mellow a further notch. A bean fermented perfectly and dried badly is a ruined bean; the two steps are one continuous flavor process with a pause for water to leave. And here is the bridge to everything this book proposes: the industry's own frontier of *control* is **starter-culture fermentation** — inoculating the heap with defined strains of yeast and bacteria rather than trusting the wild air — precisely to tame the variability that makes fine chocolate so inconsistent. That move, from wild to defined microbes, is the exact same instinct as ours, one level out: reach into an uncontrolled process and make it reproducible. We simply reach one level deeper still, into the seed itself, so that what goes *into* even a defined ferment already carries its three dials pre-set. Control the starter and you control the ferment; design the seed and you control what the ferment has to work with.

Read that third act against the three dials. Proteolysis builds the precursors of aroma *and* the raw material of body. PPO and leaching take the astringency and bitterness *down*. Fermentation is bitter-down, astringent-down, body-and-aroma-up — *it is reaching for exactly our three axes*, with microbes and the seed's own enzymes instead of a designed genome. The chocolate-maker who says a bean is "well-fermented" is saying, in older words, that its B, A, and K came out balanced.

And here is where the blunt genius runs out — the gap that is the whole argument for engineering. Fermentation is *uncontrolled*. It is wild microbes, ambient temperature, a farmer's judgment, weather, the size of the heap, how often it is turned. Under-ferment and the bean stays brutally bitter and astringent, grassy and sharp. Over-ferment and it turns hammy, putrid, hollow. The coefficient of variation of *flavor* across fermentation batches is enormous — nowhere near our ≤ 7% — which is most of why fine chocolate is so maddeningly inconsistent lot to lot, and why a maker will pay triple for beans from one cooperative that has gotten its ferment reliable. And fermentation cannot *separate*: PPO lowers astringency by *destroying* flavanols, with no way to spare the gentle short ones and oxidize only the harsh long ones; leaching pulls out bitterant and flavanol together. The first engineer works on the whole bean from the outside, trading virtues it cannot tell apart, hitting a target it cannot specify, succeeding brilliantly on average and wildly unevenly in fact.

Everything our platform offers, read against fermentation, is the same offer the scalpel makes the fire. We would not replace fermentation — we revere it, and much of a bar's aromatic soul is born in those three acts. We would give the bean that goes *into* the ferment a known, reproducible starting point: a methylxanthine budget already set at the enzyme, a polymer-length distribution already shifted toward gentle structure, a peptide pool already biased toward body — so that fermentation does less heavy lifting and more fine finishing, and so that what the farmer's judgment contributes is nuance on top of a floor we have guaranteed. Precision fermentation from the *inside*. The microbes keep their ancient job. We just stop asking them to do the whole thing blind.

Chapter 11 · The chord — *composing B, K, A on one bean.*

We have the three dials and the enzymes under each. Now the real problem, the one that makes cacao a design challenge rather than three easy edits: **the dials are coupled, and a chocolate is a single coordinate in the three-space they define.** You cannot set bitterness, then astringency, then body, as if they were independent knobs, because in cacao's metabolism and on the human tongue they push on each other. Composing a cacao is solving a system. This chapter is how the system is solved, and what the solutions — the design-target cultivars this book is building toward — actually look like.

First, enumerate the couplings, because each is a real constraint the flux model enforces.

Carbon coupling (metabolism). The bitter alkaloids and the astringent phenolics draw on overlapping upstream carbon and energy. More deeply, the phenolic *bitters* and the phenolic *astringents* are partly the same molecules at different polymer lengths — the smaller flavan-3-ols are both mildly bitter and mildly astringent. So an edit that shifts the polymer distribution shorter to soften astringency can nudge bitterness *up* by raising the monomer pool; an edit that throttles total phenolic flux to reduce bitter phenolics can drag astringency *down* as a side effect. The genome-scale flux-balance model holds the whole network at once precisely so these collateral moves are predicted and compensated, not discovered on the tongue.

Perceptual coupling (the mouth). Kokumi, via CaSR, *suppresses perceived bitterness* and *amplifies perceived sweetness and umami*. So raising IFI-K lowers *apparent* bitterness without changing a single alkaloid — the kokumi axis and the true-attenuation axis are two routes to the same felt smoothness. Astringency and bitterness, meanwhile, *reinforce* each other perceptually: a mouth already gripped by tannin reads bitterness as harsher, and a bitter mouth reads tannin as rougher. The felt harshness of a cacao is therefore super-additive in B and A and *relieved* by K. A great design exploits all three couplings: a touch of real de-bittering, a shift of astringency toward structured rather than aggressive, and a lift of body that conducts the rest into roundness.

The composition rule that falls out of this is the one the whole library runs on, and the multi-axis volumes state it exactly: axes *compose*, they do not simply add, and the composite indices reflect it — a well-built dual-axis design reaches a composite far above either axis alone (the M2 dual anchor sits at 8,400), while an over-stuffed triple can *interfere* and fall back (the M3 triple anchor at 2,800). Cacao is natively a three-axis instrument, which is why it is

both wonderful and unforgiving: three coupled axes give you a huge expressive space and a thousand ways to muddy it. The design target is never "max everything." It is a chosen, coherent chord.

Make the couplings quantitative, because a designer needs numbers, not only warnings. On the perceptual side, CaSR engagement suppresses the bitter report measurably: in our panels a strong kokumi lift — on the order of several hundred index points of IFI-K — buys back the *apparent* smoothness of a few hundred points of IFI-B, without touching one alkaloid. That is a real exchange rate, and it is why the darkest rungs of the ladder keep their body high: the kokumi is doing covert de-bittering the enzyme never has to. On the harshness side, bitterness and astringency are **super-additive** — a mouth already gripped by tannin scores a given bitterness as harsher, and the reverse — so the felt severity of a bar climbs faster than either axis alone, and relief on *either* axis relieves both. A designer spends against exactly this: a touch of real de-bittering, a shift of astringency toward structure, a lift of body, and the three small moves compound into a bar that feels far smoother than any one of them would predict.

Walk a 70% dark through it, toward **Nocturne**. Start from the native Forastero-ish coordinate — bitterness high, astringency high and a little green, body modest — and set three dials. Throttle the methylxanthine flux for about 38% attenuation (IFI-B ≈ 1,500): the bite becomes seasoning. Shift the procyanidin distribution toward mean degree of polymerization four to five (IFI-A ≈ 1,600): the grip stops being green and starts being architecture, long and resolving. Lift the γ-glutamyl peptides for real body (IFI-K ≈ 1,900): the middle fills, the finish lengthens, and the remaining bitterness is conducted into pleasure rather than endured. Three edits, one coordinate — and because the couplings were respected rather than fought, a bar that reads as *whole*. Here is the everyday-great middle rung, **Umbra**, as the index renders it, sitting between the two extremes we score below:

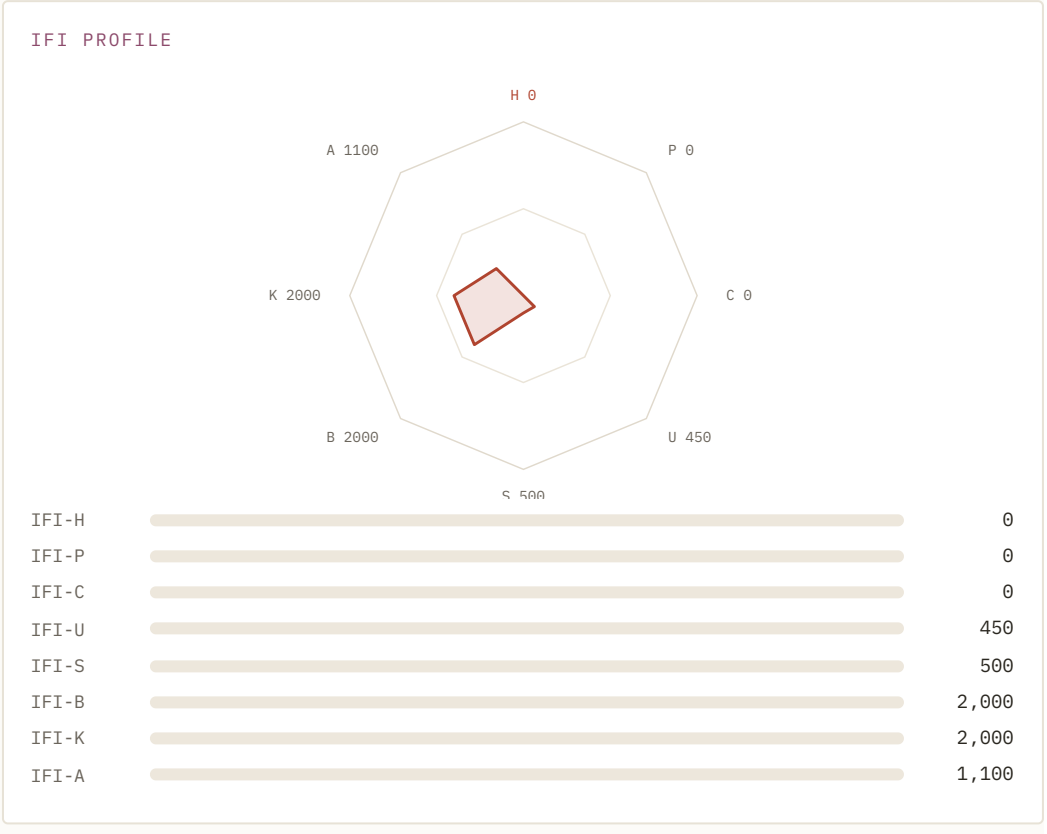



Garlic — the γ-glutamyl powerhouse, 560 constructs. The kokumi lever that rounds the cacao chord was proven in the most body-giving plant in the kitchen.

So let us compose. Below is the design-target ladder for cacao — a named series we are building toward, running from the gentlest, most-attenuated, most-rounded expression to the fiercest, most-structured, least-mercied one. We name it for the deepening of night, because that is the journey of a cacao from milk-soft to connoisseur-dark: **Gloaming → Umbra → Nocturne → Compline → Starless.** We present these honestly as *design targets*, not shipped SKUs — cacao is not yet among our forty-two hosts, and these carry no catalogue number because we do not invent them. They are the coordinates the method of Chapters 6 and 8 would be aimed at, and the baselines they modulate (the one-to-three-percent theobromine, the several-percent procyanidins) are cacao exactly as it grows today.

DESIGN TARGET	INTENDED BAR	IFI-B (ATTENUATION)	IFI-A (ASTRINGENCY / MDP)	IFI-K (BODY)	THE CHORD IN A PHRASE
Gloaming	milk / 45–55%	~~2,600~~ (65%)	~600 (mDP 1–2, soft)	~2,200	mercy and body; the bite a whisper, the round immense
Umbra	approachable dark / 60–65%	~~2,000~~ (52%)	~1,100 (mDP 3)	~2,000	the everyday great bar; softened, structured, rich
Nocturne	classic dark / 70%	~~1,500~~ (38%)	~1,600 (mDP 4–5)	~1,900	the reference dark — bite seasoned, grip architectural
Compline	intense dark / 80–85%	~~1,100~~ (28%)	~2,200 (mDP 6)	~1,700	the edge kept, the body holding it up
Starless	connoisseur / 90%+	~~900~~ (22%)	~2,800 (mDP 7–8)	~1,500	almost no mercy; structured ferocity, long resolve

Read the ladder as a single idea expressed five ways: as the bar darkens, we spend less mercy on the bitterness (IFI-B falls), we build more structure into the astringency (IFI-A rises, toward longer resolving polymers rather than toward crude green grip), and we let the body ease back just slightly — enough to keep the dark honest rather than plush. The thing that makes every rung *good* rather than merely dark is that at each one the astringency is *structured*, not

aggressive, and the kokumi is *present* enough to conduct the bite into pleasure. A bad 90% is Starless's astringency without Starless's body and without the polymer sculpting — crude length, no round. The engineering is what lets a 90% be a joy instead of a dare.

Think of the ladder, finally, as a *spec sheet* a maker orders from. A chocolatier who wants a signature 72% for a particular bonbon does not have to hunt harvests and hope; they name a coordinate — say IFI-B 1,400, IFI-A 1,700, IFI-K 1,950, a hair darker and more structured than Nocturne — and that is a reproducible order, delivered to ≤ 7% CV, this season and the next. A milk-chocolate house names Gloaming's coordinate and gets a bar whose body does not drift between production runs. This is what measurement buys that selection never could: not merely a better *average* bean, but a *named, repeatable* one — a flavor specified the way a machinist specifies a tolerance. The whole three-space becomes a catalog, and every point in it that makes a good chocolate is a product waiting to be ordered rather than a lucky harvest waiting to be found. That is the real end of the chord: not any single bar, but the end of having to hope.

Here are the two ends of the ladder as full eight-axis scorecards, so you can see the chord as the index renders it. **Gloaming** — the gentlest, a designed milk-leaning cacao where mercy and body do everything and the astringency is nearly laid down:

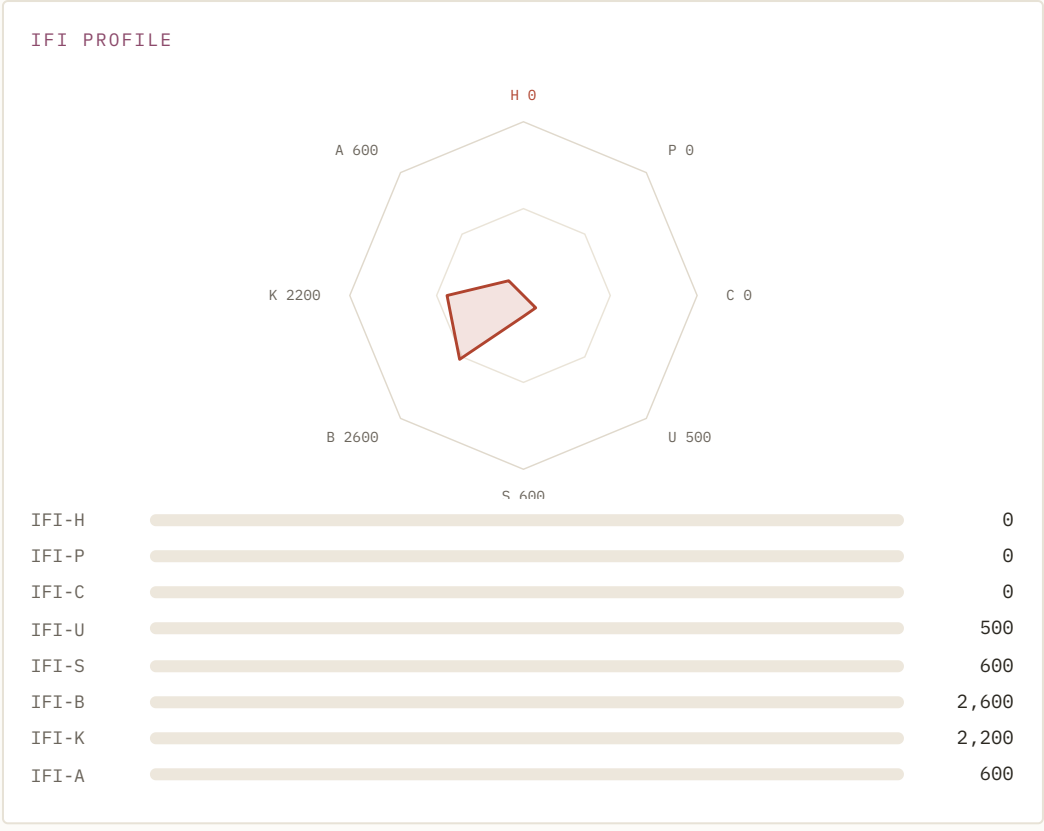

And **Starless** — the connoisseur's near-unsweetened extreme, almost no de-bittering, astringency built tall and structured, body pulled back to keep it severe and long:

IFI PROFILE

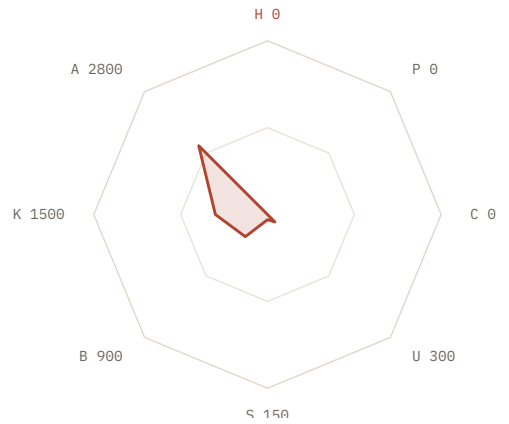

IFI-H		0
IFI-P		0
IFI-C		0
IFI-U		300
IFI-S		150
IFI-B		900
IFI-K		1,500
IFI-A		2,800

Two coordinates in the same three-space, the same plant, the same three enzymes set to different values — and in the mouth, two completely different chocolates, one a pillow and one a blade, both of them *right*. That is what it means to design a cacao: not to make it "better," but to place it, exactly, where you meant it to go, and to land there every single batch. There is one more actor before the bean becomes a bar, though — the maker, and the fire. The engineering sets the seed; the craft reveals it.

Chapter 12 · From bean to bar — *roast, conch, temper.*

A perfectly designed cacao seed is still, at harvest, an inedible thing. Between the fermented, dried bean and the square that snaps on your tongue lie four transformations of real and beautiful chemistry — roasting, grinding, conching, tempering — and a cacao engineer who does not understand them is a pianist who has never heard the hall. Our edits set the instrument; these steps play it, and some of them can undo in an afternoon what the genome spent a season building. This chapter is the craft, in the terms the craft deserves, and where our three dials meet the maker's four fires.



Phoenix in the controlled chamber — not a cacao, but the company's proof that a designed seed and a disciplined process meet at a measured number. The same rigor the roast demands.

The roast. Dried, fermented beans are roasted — typically 110–140 °C for ten to forty minutes — and this is where "chocolate" the aroma is actually born, through the **Maillard reaction**: the fermentation-made free amino acids and reducing sugars condense and rearrange into a cascade of flavor molecules. The signatures of cocoa aroma are **pyrazines** — 2,3,5,6-tetramethylpyrazine and the trimethyl- and dimethyl-pyrazines, nutty, roasty, "chocolatey" in the way the word actually means — formed alongside **Strecker-degradation** aldehydes (3-methylbutanal, phenylacetaldehyde) and the alcohol 2-phenylethanol's rosy note. Roast too light and the pyrazines never develop; roast too dark and you bury the origin character under burnt, acrid pyrazine overload. For our purposes the critical facts are two. First, **theobromine is thermally stable** — it passes through the roast essentially unchanged, which means the bitter budget we set at the enzyme *survives to the bar*; the roast cannot rescue a bitterness we mis-set, nor does it spoil one we set right. Second, the roast **further degrades the flavanols** — the astringent and phenolic cargo drops again — so a cacao we engineered to carry gentle, flavanol-rich, low-mDP structure gives the roaster more room: there is grip to spare, and the flavanols that survive are more of them, because we started with a distribution built to endure.

The roast chemistry rewards a closer look, because it is where cacao's *aroma* — as distinct from its taste — is almost entirely made. The Maillard reaction is not one reaction but a cascade: reducing sugars condense with amino acids into Amadori compounds, which fragment and recombine into hundreds of volatiles, while the parallel **Strecker degradation** turns amino acids into their characteristic aldehydes. The amino acids matter individually — leucine throws the malty 3-methylbutanal, phenylalanine the honeyed phenylacetaldehyde and rosy 2-phenylethanol — which is why the *proteolysis* of fermentation, by setting the free-amino-acid pool, quietly pre-writes the roast's aroma. The signatures of "chocolate" itself are the **pyrazines**:

ROAST-BORN COMPOUND	FORMED BY	NOTE IT CARRIES
Tetramethyl- and trimethylpyrazine	Maillard	nutty, roasty — the core of "cocoa"
3-methylbutanal (from leucine)	Strecker	malty, dark
Phenylacetaldehyde, 2-phenylethanol (from phenylalanine)	Strecker	honey, rose
Furaneol, maltol	sugar pyrolysis	caramel, sweet-brown

Roast kinetics are a knife-edge. Too light — under roughly 110 °C, or too brief — and the pyrazines never develop, leaving a raw, sour, astringent bean; too dark and the pyrazines overshoot into acrid and burnt, burying origin character under char. The maker reads that window the way the farmer reads fermentation — and, exactly as with fermentation, a seed engineered to arrive with its bitter and astringent dials pre-set gives the roaster room to chase *aroma* rather than spend the roast fighting a bitterness that should never have been that high.

The grind. Roasted nibs are ground — once by stone, now by steel and ball mill — into **cocoa liquor**, a suspension of cocoa solids in the melted **cocoa butter** that is a third to a half of the nib's mass by weight. Cocoa butter is a remarkable fat: a tight, nearly uniform set of triglycerides, dominated by **POP, POS, and SOS** (P = palmitic, O = oleic, S = stearic acid on the glycerol backbone), and it is that uniformity that gives the butter its sharp, almost crystalline melting behavior, the thing that makes chocolate *snap* and then *melt* at body temperature. The grind also sets **particle size**, and here is a number every chocolatier carries: the tongue detects grittiness above roughly **25–30 micrometers**, so fine chocolate is refined until its solid particles sit below that threshold and the chocolate reads as *smooth*. Smoothness is a mechanical property before it is a flavor one — and it interacts with our axes, because a smoother particle bed changes how the astringent phenolics and the alkaloids are released to the saliva. A coarse chocolate tastes more astringent and more bitter at the same chemistry, simply because more phenol meets more saliva faster.

The conch. The ground mass is conched — agitated, warmed, aerated for hours to days, a process named for the shell-shaped troughs of its invention. Conching does three things that touch our dials directly. It **drives off volatile acids** — chiefly the acetic acid carried over from fermentation — so an over-sharp, sour bean is

mellowed. It **reduces moisture and viscosity**, coating every particle in butter so the mouthfeel turns silken. And it **continues to develop and round the flavor**, as remaining Maillard chemistry proceeds and harsh edges volatilize. Conching is the maker's own *kokumi-and-smoothness* step — it rounds body and reduces harshness by physical means, exactly parallel to what our CaSR peptides do perceptually. A cacao we designed with real body needs less heroic conching to taste complete.

The temper. The last fire is the coldest and the most exacting, and it is pure physics: **cocoa butter is polymorphic** — its triglycerides can crystallize into six different solid forms, I through VI, with different melting points and properties — and only one of them, **Form V (βV)**, makes good chocolate. Form V melts at **33.8–34.4 °C** — just below body temperature, which is why good chocolate melts *on* the tongue and not *in the hand* — and it gives the finished bar its gloss, its hard snap, and its stability. Tempering is the craft of forcing cocoa butter to crystallize into Form V and nothing else: a controlled dance of melting the butter completely (~~50 °C, erasing all crystal memory), cooling it to seed the crystals (~~27 °C), then rewarming slightly (~31–32 °C for dark) to melt out the unwanted lower forms and leave a population of pure Form V seeds to template the rest. Get it right and the bar is mirror-glossed and snaps clean. Get it wrong — or let a tempered bar sit warm — and the butter slowly converts to **Form VI**, higher-melting, and migrates to the surface as the dull grey-white **bloom** every chocolate lover knows as the sign of a bar gone slack.

The polymorphism is worth laying out in full, because it is the most purely physical thing in the whole book and the least forgiving of carelessness:

FORM	MELTING POINT	STATE
I (γ)	~17 °C	softest, least stable
II (α)	~23 °C	unstable
III	~25 °C	unstable
IV (β′)	~27 °C	unstable
V (β)	**~34 °C**	**the goal — snap, gloss, stability**
VI (β)	~36 °C	over-aged; surfaces as bloom

Only Form V makes good chocolate, and it does so for a reason that is nearly poetic: its ~34 °C melt sits just below body temperature, so the bar is solid in the hand — glossy, snapping — and melts the instant it reaches the tongue, releasing fat, flavor, and every axis we tuned in one coordinated flood. Tempering is the craft of seeding Form V and only Form V: melt completely near 50 °C to erase all crystal memory, cool to roughly 27 °C to nucleate crystals, then rewarm to about 31–32 °C for dark chocolate to melt out the lower, unwanted forms and leave a pure Form V population to template the rest. The particle physics matters alongside the fat: molten chocolate is a **yield-stress fluid**, described by the Casson model, and conching plus a trace of lecithin tune that yield stress so the chocolate coats and releases correctly — too thick and it drags, too thin and it will not hold a shell. None of this is ours to edit; cocoa butter's polymorphism is the same whatever cacao we design. We

set the flavor; the maker owns the crystal, and the division of labor is clean. Temper is where the chocolatier's hands matter most, and it is entirely downstream of our genetics — cocoa butter's polymorphism is the same whatever the cacao — which is our way of saying: we set the flavor; the maker owns the crystal. The division of labor is clean and it is respectful. We would never want it otherwise.

One last craft step earns a mention because it is the crude chemical shadow of this whole book: **Dutch processing**, the alkalization of cocoa with a potassium or sodium carbonate, which raises the pH, darkens the colour, mellows the acidity, and *reduces astringency and bitterness* — by, essentially, chemically neutralizing and degrading the phenolics wholesale. It is effective and it is traditional and it is exactly the blunt instrument our scalpel improves upon: Dutching lowers astringency by *annihilating* the flavanols, every one, with no way to keep the gentle and lose the harsh. Everything the alkali does to the whole phenolic pool at once, from outside, after the fact, we would do to the polymer-length distribution, selectively, at the enzyme, before the bean is even formed — astringency sculpted, flavanols spared. The Dutch process is the question this book is the precise answer to.



PART V OF VI

The Kitchen

tasting the three dials · the dark side · body & balance · swicy & spiced

Chapter 13 · Reading a bar — *a tasting method, scored.*

The reward for all this biology is that you can now taste deliberately — not "I like it / I don't," but *where on the three axes is this, and is it where its maker meant it to be?* This short chapter is a method: how to read a chocolate on B, A, and K with your own mouth, calibrated against the index, so that the scorecards in the recipe chapters that follow are something you can verify rather than merely trust. We taste as a family this way, passing one square around and arguing in numbers, and it has ruined us, pleasantly, for casual chocolate forever.



A cacao seed in isolation — the thing under all the craft. Taste toward it, not away from it.

Work in order of how the sensations actually arrive, because they do not arrive together.

1. **Let it melt; do not chew.** Place a small square on the tongue and let body heat bring it through Form V's melt. Chewing rushes everything into a muddle; melting lets the axes separate in time. Note the *snap* as you break the square — a clean snap is good temper, Form V intact, and tells you the fat is reading its physics right before you taste a thing.
2. **First, the aroma and the sweet-umami floor** (the murmured axes). As the butter melts, the Maillard aromatics bloom retronasally — the pyrazine nuttiness, any fruit or floral the ferment left. Underneath, note the quiet sweetness and the faint savory floor. These are the low numbers on the scorecard, the U and S that sit near a few hundred.
3. **Then the bitterness arrives** (IFI-B, read as *how much mercy*). Theobromine is slow; the bitterness builds over a few seconds and broadens. Ask not "is it bitter" but "is the bite *seasoning or wall* — and does it feel *placed*?" A high-IFI-B design reads as softened, the bitterness a note; a low-IFI-B connoisseur bar reads as a deliberate, standing wall. Calibrate: a supermarket 55% is high-B; a craft 90% is low-B by design.

4. **The grip builds under it** (IFI-A, read as *structure vs. green*). Now the mouth begins to dry and grip — the astringency, arriving under the bitterness and lasting past it. Judge its *quality*, not just its amount: does it grip cleanly and *resolve*, mouth-watering, architectural (long, well-made polymers, a good IFI-A)? Or does it strip and stay, green and harsh (crude length, a cacao under-fermented or over-roasted)? Astringency is the axis most makers get wrong, and the one your own saliva reads most honestly.
5. **Finally, the body and the length** (IFI-K, read as *does it fill in and last*). As everything fades, ask whether the flavor *filled the middle and went long* — rich, continuous, resolving slowly — or whether it flashed and dropped to nothing. That fullness and length is kokumi, the CaSR round. A great bar's finish is mostly this axis. A hollow one has none, and you will reach for a second square to chase the fullness the first one promised and withheld.

Two aids make the method reliable. The first is understanding *why* melting beats chewing, which is pure physiology: Form V releases its cargo over several seconds as it crosses its narrow melt, so the axes arrive *in sequence* — aroma, then the slow bitterness, then the building grip, then the fading body — and that separation in time is exactly what lets you read them one at a time. Chewing collapses the sequence into a single flooded instant and recruits mechanical crushing that muddies astringency's own mechanics; you lose the structure you were trying to read. The second aid is calibration: taste mid-morning when the palate is clean, against a known reference from the ladder, and reset between samples with water and a bite of plain bread or apple to rebuild the salivary film that astringency depletes. If you have only one thing to hand, let it be room-temperature water and a little patience — astringency clears on its own in a minute or two as saliva rebuilds, and rushing from square to square is the single most common way a taster mistakes a fine bar's cumulative grip for a flaw it does not have. Then learn to read a flaw back to its cause, because every off-note points at a specific step upstream:

WHAT YOU TASTE	LIKELY CAUSE	AXIS / STEP IMPLICATED
Harsh, green, mouth-stripping grip	under-fermented or under-roasted	`IFI-A too high; craft, not design`
Flat, sour, sharp acidity	under-conched (acetic acid retained)	`the conch, Chapter 12`
Hollow, short, no finish	low body; over-attenuated or thin kokumi	`IFI-K too low`
Ammoniacal, hammy, putrid	over-fermented	`fermentation window blown`
Waxy, dull, grey film	bloom — lost temper (Form VI)	`tempering, not flavor`
Bitterness that walls everything else out	bitter dial set too high, or no body to conduct it	`IFI-B low` *and* `IFI-K low`

Read that table twice and you can taste like a maker — not "I like it" but "the astringency is a hair green and the body is thin, so this bean was pulled from the ferment a day early and under-conched," which is a sentence you can *act* on. That is the whole gift of a measured axis: it turns a complaint into a diagnosis, and a diagnosis into a better next batch.

Two last practicalities, because the axes change with how you serve and what you serve alongside. **Temperature** matters more than most people think: cold suppresses perceived sweetness and sharpens astringency, so a chocolate tasted straight from the fridge reads darker, more bitter, and more gripping than the same bar at cool room temperature — which is why you condition your *palate's* circumstances as carefully as the chocolatier tempered the bar, tasting at 18–20 °C. **Pairing** is a way to read individual axes louder. Water and a neutral cracker between samples reset the astringent film so each bar starts fair. A sip of still water taken *during* a tannic bar reveals the astringency's carryover; a sip of milk shows how much the grip depends on an unlubricated mouth (it will nearly vanish). A tannic red wine or a strong black tea alongside a dark bar lets you feel two IFI-A systems stack — sometimes harmonizing, sometimes clashing — which is the whole lesson of the astringent shelf. And to *train* the palate, work a progression: start with Gloaming to learn what "almost no bitterness, all body" feels like, step to Nocturne for the reference dark, then to Starless for structured ferocity, and back down. Three bars, bracketing the space, tasted in order and then out of order until you can name a coordinate blind. It takes an afternoon and a glass of water — and afterward you will never again taste a chocolate as a single verdict, only as three numbers in a relationship, which is both more useful and, we promise, more pleasurable. To read is not to stop loving. It is to love with your eyes open.

Run that sequence and you can place any chocolate on the three dials to within a few hundred index points with practice, which is enough to read a design and catch a flaw. Here are two reference bars to calibrate against, given as the targets from our Chapter 11 ladder so you can taste toward a known coordinate. **Nocturne**, the classic 70% dark — the reference against which everything else reads:

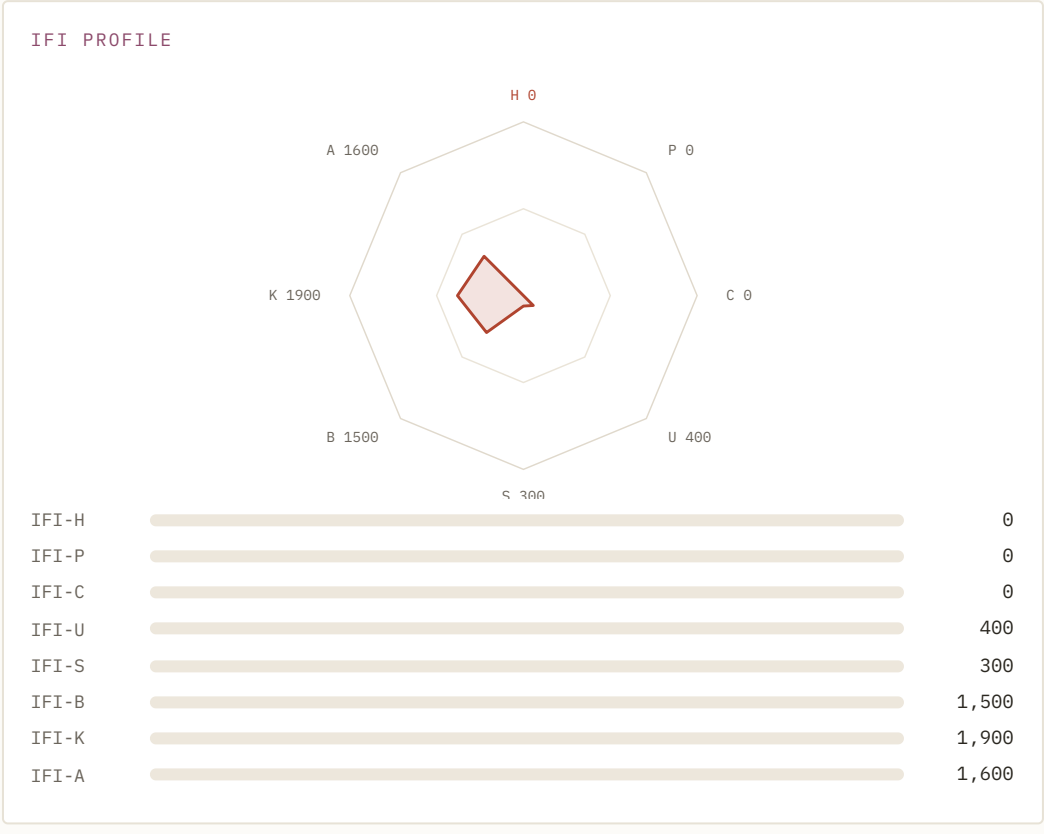

Bitterness seasoned, not walled (IFI-B ~1,500, a bit over a third off). Grip architectural and resolving (A ~1,600, mDP four to five). Body full and long (K ~1,900). If a 70% in your mouth reads much bitterer, harsher, or hollower than that, it is either under-designed or mis-made — and now you can say *which*. And **Compline**, the intense 80–85% for when you want the edge kept and the body to carry it:

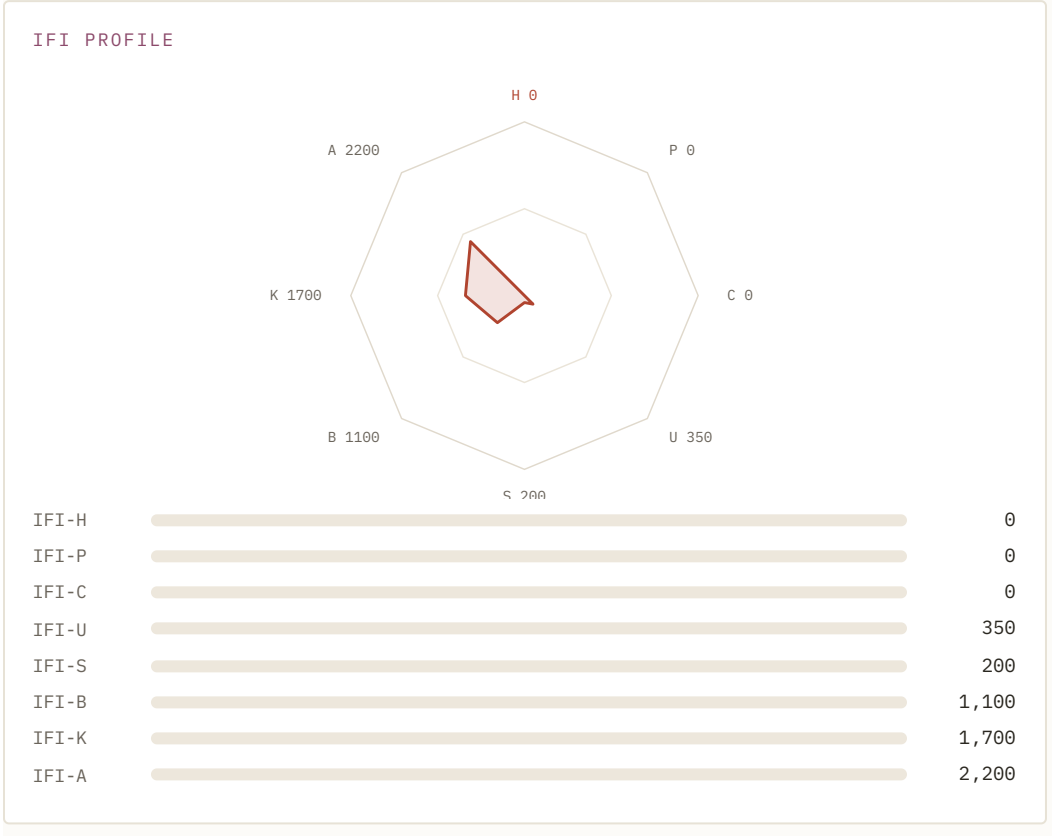

Less mercy (B ~1,100), more structure (A ~2,200), body easing back but still holding the dark up (K ~1,700).

And for the far, gentle end of the ladder — the milk-leaning bar where mercy and body do nearly all the work — hold **Gloaming** in memory too, so your palate carries the low anchor as well as the high:

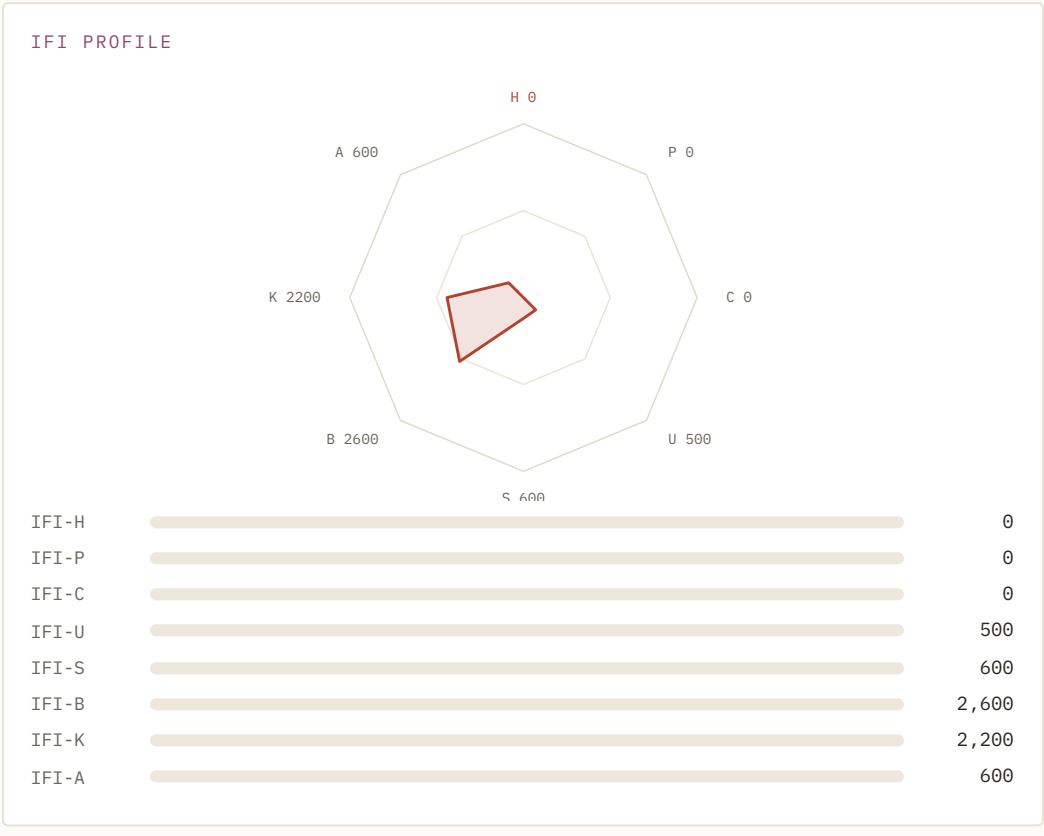

Three bars, three coordinates, one plant. Learn them by heart and every chocolate you ever meet has somewhere to stand. Taste the two side by side and you will feel the ladder move under you — the bite standing taller, the grip building, the round pulling back — which

is the whole three-space of cacao, learnable on an afternoon with five squares and a glass of water. Now, the kitchen, where the dials become dinner.

Chapter 14 · The dark side — *bitter-forward recipes.*

The recipes in this book are built, every one, around a designed cacao placed at a chosen coordinate, and scored on the index so you can taste the composition the way its author meant it. We cook plant-first, as this whole library does, and where a dish is better for dairy or an egg we reach for them without apology and treat them with the same care as the cacao. This first recipe chapter is for the people who came for the *dark* — the bitter-forward constructions where we spend the least mercy and let theobromine and structured tannin do the talking, rounded only by as much body as keeps them humane. Each recipe names a design target from the Chapter 11 ladder; if you are cooking with cacao as it grows today, read the target as the bean to *select* for — the darkest, best-fermented, most structured origin you can find.

A sidebar on natural versus Dutched cocoa, since the dark recipes live or die on it. Natural cocoa powder is simply ground, defatted cacao — acidic (pH around 5 to 6), fruity-sharp, and *astringent*, its flavanols largely intact. Dutched (alkalized) cocoa has been treated with a carbonate that raises the pH toward neutral or above, darkening the color and mellowing both the acidity and the grip — and, exactly as Chapter 12 warned, stripping the flavanols wholesale in the bargain. For a bitter-forward recipe we reach for **natural** cocoa on purpose: we *want* the acidity and the bare astringency, and we want the flavanols kept. It matters chemically too, wherever there is leavening — natural cocoa's acidity reacts with baking soda while Dutched does not — so swapping one for the other is never only a flavor choice. Every recipe in this chapter means natural cocoa unless it says otherwise.



Cacao in the pod — the dark, before we darken it further.

Oaxacan-dark drinking chocolate · water-based, theobromine-forward

The oldest way to drink chocolate is also the most honest: no milk to hide behind, just cacao and water beaten to a foam, the way it was drunk for three thousand years before a European thought to add sugar and a cow. Built on **Starless** or the darkest bean you can find — this is a cup that means it.

85–90% chocolate (Starless target), chopped	60 g
Water, just off the boil	240 ml
Cane sugar (or a few drops Stevia-S extract)	to taste, start at 8 g
Fine sea salt	1 small pinch
Cinnamon, a short splinter (Cinnamon-P, optional)	1
Ancho or pasilla chile, a whisper (optional)	1 pinch

1. Warm the water to just below boiling — scalding water seizes and dulls the aromatics; aim for about 85 °C.
2. Add the chopped chocolate, the salt, and the cinnamon. Let it sit thirty seconds, then whisk hard, or pull it back and forth with a *molinillo* until it foams.
3. Taste before you sweeten. On water, with no dairy to round it, the astringency and bitterness arrive naked — this is the point. Add sugar a few grams at a time, chasing *balance*, not sweetness; stop while the dark still clearly leads.
4. The optional pinch of chile is not heat for heat's sake — a whisper of ancho lengthens the finish and reads as warmth under the bitterness. Serve foaming, in a small cup; this is not a mug drink.

IFI PROFILE

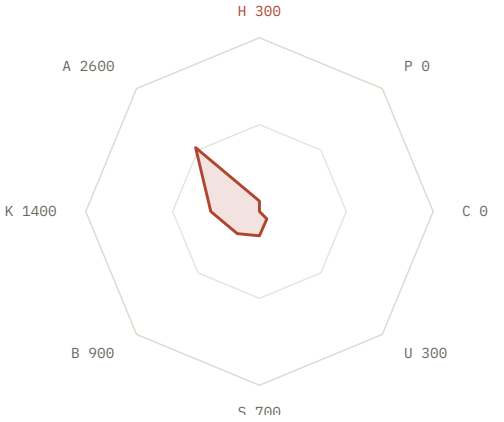

IFI-H	300
IFI-P	0
IFI-C	0
IFI-U	300
IFI-S	700
IFI-B	900
IFI-K	1,400
IFI-A	2,600

100% ganache truffles · the unsweetened heart, rounded only by cream

A truffle made from unsweetened cacao is the purest test of a bean and of body — there is nowhere to hide. The cream's dairy kokumi and fat are the only mercy, standing in for the kokumi we would build into **Compline**, on which this is based.

100% chocolate / unsweetened cacao liquor	200 g
Heavy cream	180 ml
Honey or glucose	25 g
Fine sea salt	2 g
Unsweetened cocoa powder, to roll	60 g

1. Chop the cacao fine and set it in a bowl. Warm the cream with the honey and salt until it steams and the honey dissolves; do not boil.
2. Pour the hot cream over the cacao in three additions, stirring from the center out each time, until the ganache comes together glossy and smooth — an emulsion of cocoa butter and cream, the same physics as temper at a softer set.
3. Chill until firm, four hours. Scoop into rough spheres and roll in cocoa powder. The powder's own bare astringency is part of the design — it frames the rounded center with an unrounded edge.
4. Eat at cool room temperature so the Form-V-adjacent fat reads at its melt. Let each one *melt*; a 100% ganache chewed is a punishment, melted it is a revelation.

IFI PROFILE

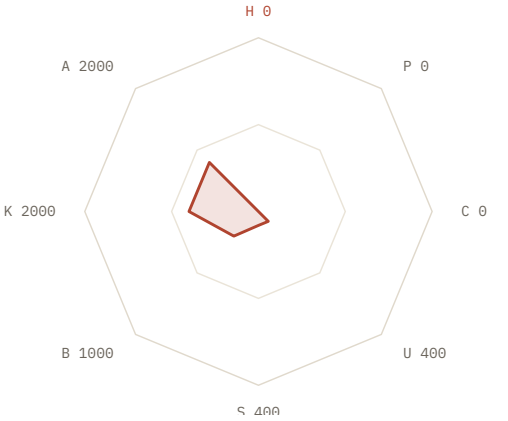

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	400
IFI-S	400
IFI-B	1,000
IFI-K	2,000
IFI-A	2,000

Flourless bitter-chocolate torte · structure without starch

No flour, so the cacao has nowhere to hide and the astringent structure becomes the architecture of the cake itself. Built on **Nocturne** for a torte that is dark and serious but still generous.

70% chocolate (Nocturne target)	250 g
Unsalted butter	170 g
Eggs, separated	6
Cane sugar	150 g, divided
Cocoa powder	30 g
Fine sea salt	3 g
Espresso or strong coffee (Coffee-B), cooled	30 ml

1. Heat the oven to 160 °C and line a 23 cm springform. Melt the chocolate and butter together gently until glossy; stir in the cocoa, salt, and coffee. The coffee is not for coffee flavor — a little of our bitter-tuned Coffee-B deepens the cacao's own bite without adding a second bitterness fighting the first.
2. Whisk the yolks with two-thirds of the sugar until pale and thick, then fold into the warm chocolate.
3. Whip the whites to soft peaks, rain in the remaining sugar, and whip to firm, glossy peaks. Fold into the batter in thirds — the whites are the only lift, so keep them alive.
4. Bake 35–40 minutes until set at the edge and barely trembling at the center. Cool completely in the pan; it will fall, and the fall is the dense, structured crumb you want. Serve thin — this torte is all three dials at full concentration.

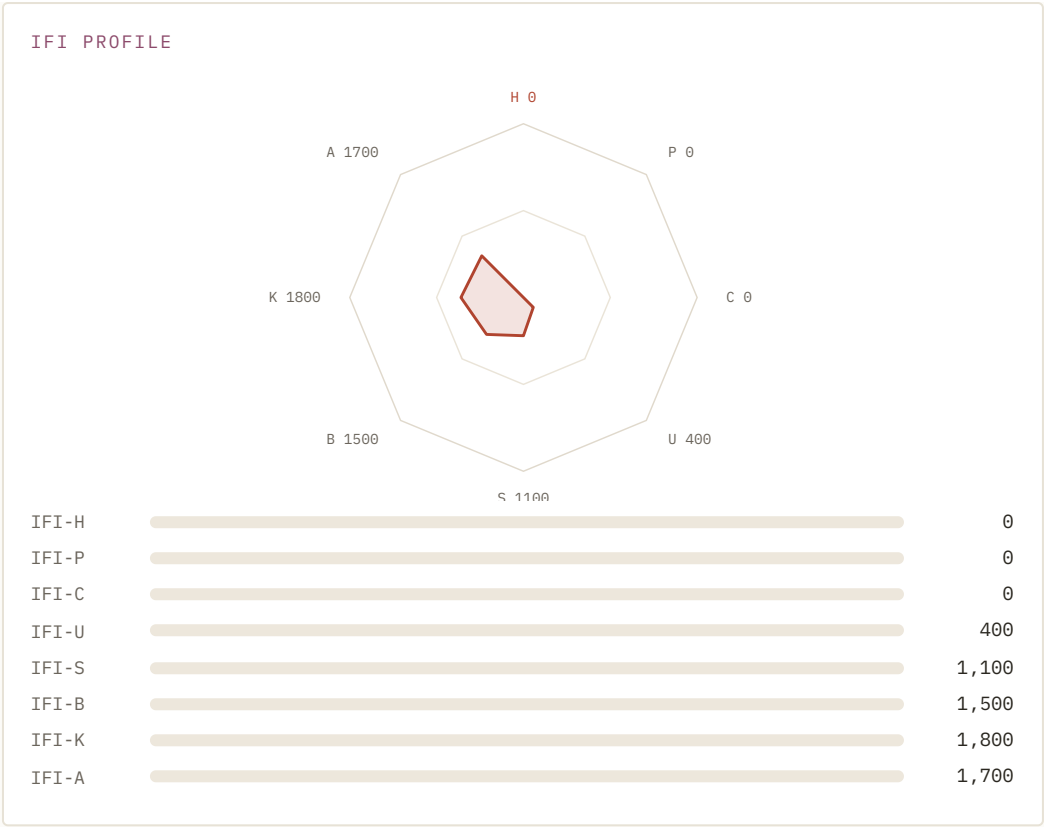

Bitter-chocolate & olive-oil sorbet · water, cacao, nothing to blur it

A dairy-free sorbet throws the cacao's own three axes into cold, clean relief — no cream to round the edges, only water, a good fruity olive oil for body, and the bean. Built on **Umbra** so it stays dark but scoopable and generous.

65% chocolate (Umbra target)	180 g
Water	500 ml
Cane sugar	120 g
Cocoa powder	40 g
Extra-virgin olive oil, peppery	30 ml
Fine sea salt	3 g

1. Bring the water, sugar, cocoa powder, and salt to a simmer, whisking, until the sugar dissolves and the cocoa is fully hydrated — about three minutes.
2. Off the heat, whisk in the chopped chocolate until melted and smooth, then whisk in the olive oil. The oil is the body here, standing in for dairy kokumi; a peppery oil adds a faint pungent lift that lengthens the finish.
3. Chill thoroughly, then churn. Without dairy fat the astringency reads clearer and colder — a bracing, structured sorbet, not a soft one.
4. Serve with a few grains of flaky salt. The cold suppresses perceived sweetness, so this will taste darker frozen than it did as a warm base — scoop it knowing that.

IFI PROFILE

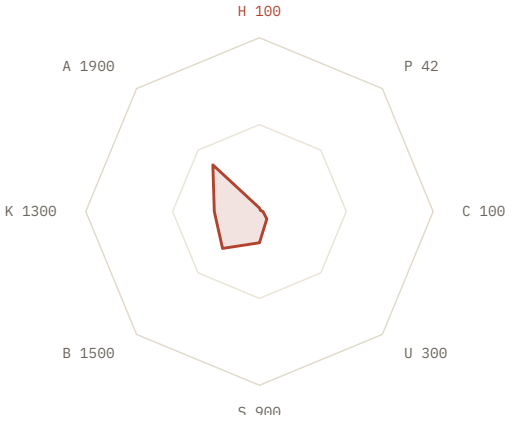

IFI-H	100
IFI-P	42
IFI-C	100
IFI-U	300
IFI-S	900
IFI-B	1,500
IFI-K	1,300
IFI-A	1,900

Bitter rye & cacao-nib sablés · whole-grain, barely sweet

Rye and natural cocoa are both bitter on their own; built together with unrounded nibs, they make a sandy, grown-up cookie that wants a small cup of coffee, not a lunchbox. Built on **Starless** — or the darkest chocolate you have.

Dark rye flour	140 g
All-purpose flour	60 g
Natural cocoa powder	30 g
Cold unsalted butter, cubed	150 g
Cane sugar	70 g
Cocoa nibs	50 g
Egg yolk	1
85% chocolate (Starless target), chopped	60 g
Flaky salt	4 g

1. Whisk the flours, cocoa, and salt. The natural (un-Dutched) cocoa is deliberate: it keeps the bare, un-alkalized astringency this cookie is built around, and its flavanols with it.
2. Cream the cold butter with the sugar *just* until combined — not fluffy; sablés want short and sandy, not aerated. Beat in the yolk.
3. Work in the dry mix, then the nibs and chopped chocolate, until it barely holds. Roll into a log; chill until firm.
4. Slice thick and bake at 165 °C for 14–16 minutes until set but not darkened — they are already dark, so go by firmness, not color. Cool on the sheet. Barely sweet, bitter, sandy, the nibs cracking out little bursts of unrounded grip: a cookie for the end of a meal.

IFI PROFILE

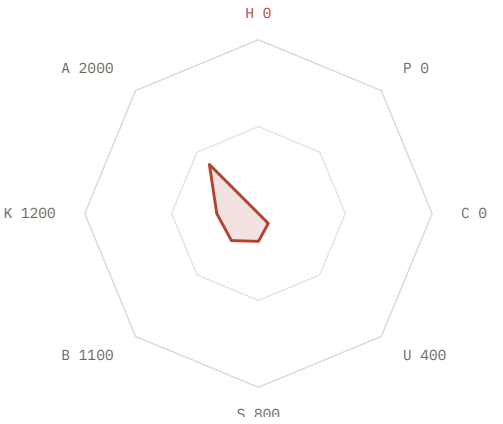

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	400
IFI-S	800
IFI-B	1,100
IFI-K	1,200
IFI-A	2,000

Chapter 15 · Body and balance — *milk, caramel, and the kokumi round.*

If the last chapter spent the least mercy, this one spends the most — the recipes where body leads, where kokumi and dairy and caramelized sugar round the cacao into comfort, where we build toward **Gloaming** and **Umbra** and let the bite become a whisper the richness carries. This is not lesser chocolate. A great milk chocolate or a true gianduja is one of the hardest things to make *good*, because with the bitterness laid down there is nowhere to hide a hollow body — the kokumi axis has to carry the whole bar. These are the recipes where IFI-K does the work.

A sidebar on why milk rounds chocolate, because it is not only fat and sugar. Milk brings its own body to the CaSR axis: its proteins and their peptides, its calcium, and — in cooked or caramelized milk especially — Maillard-born compounds all nudge perceived richness and continuity, which is a real part of why milk chocolate eats rounder than its small cocoa fraction alone could explain. Dairy also physically coats the mouth, blunting astringency by laying down lubricant the tannins must overcome before they can strip the salivary film — which is why a glass of milk tames a bitter, tannic bar, and why milk chocolate can carry a bitterness that would be punishing neat. When we build body into a *dark* cacao with γ-glutamyl peptides, we are, in effect, giving it some of the kokumi a *milk* bar gets for free from the cow.



Onion, halved — another γ-glutamyl kitchen, proof that body is a chemistry you can build. The round in a milk chocolate is the same axis as the depth in a stock.

Brown-butter gianduja · hazelnut, cacao, and the longest round

Gianduja — chocolate ground with roasted hazelnuts to a melting paste — is kokumi made edible: the nut's own fat and roasted peptides stack with the cacao's body into something that eats longer and rounder than either alone. Brown butter triples down on that Maillard depth. Built on **Gloaming**, which is engineered to be almost all body.

Hazelnuts, toasted and skinned	200 g
Milk chocolate (Gloaming target)	200 g
Unsalted butter	40 g
Powdered sugar	40 g
Flaky salt	3 g

1. Brown the butter in a small pan until it smells of toffee and nuts and the milk solids go golden — this is Maillard aroma built on purpose. Cool slightly.
2. Grind the warm toasted hazelnuts in a processor until they pass from meal to butter to a loose, glossy paste — five to ten minutes, patience, the oil has to release.
3. Melt the chocolate gently, then blend it with the hazelnut paste, the brown butter, powdered sugar, and salt until completely smooth and pourable.
4. Pour into a lined tray or small molds; set at cool room temperature, not the fridge, to keep the fat reading soft. The finish on this should go *long* — if it drops off short, your hazelnuts weren't toasted enough; the roast is where the body comes from.

IFI PROFILE

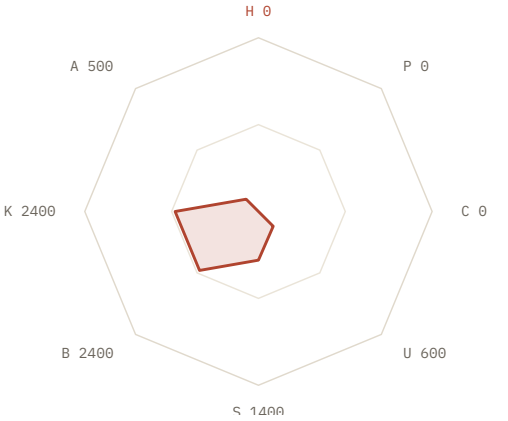

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	600
IFI-S	1,400
IFI-B	2,400
IFI-K	2,400
IFI-A	500

Salted caramel & milk chocolate · two kokumi systems stacked

Caramel is kokumi by another road — the deep, continuous, mouth-filling richness of sugar cooked to the edge of bitter, salt lifting and lengthening it. Layered under milk chocolate it is body on body, the cacao's bitterness nearly gone (Gloaming, ~65% attenuation) and the whole thing held up by two different rounds.

For the caramel:	
Cane sugar	200 g
Heavy cream, warm	120 ml
Unsalted butter	60 g
Flaky sea salt	4 g
Vanilla (real; Vanilla-K)	1 pod, scraped
To finish:	
Milk chocolate (Gloaming target), tempered	300 g

1. Cook the sugar dry in a heavy pan over medium heat, swirling (not stirring) until it is a deep amber — just to the edge where caramel's own faint bitterness begins; that edge is where the kokumi lives.
2. Off the heat, whisk in the warm cream carefully (it will seize and steam), then the butter, salt, and vanilla seeds. The real vanilla here is not decoration — its CaSR kokumi stacks with the caramel's and the cacao's into a single long round.
3. Cool the caramel to a thick, pipeable set. Pipe or spread a thin layer and let it firm.
4. Enrobe or top with tempered milk chocolate; let the Form V set to a clean snap. The bite is a whisper now; what you taste is three bodies — caramel, dairy, cacao — reading as one continuous, salted, mouth-filling thing.

IFI PROFILE

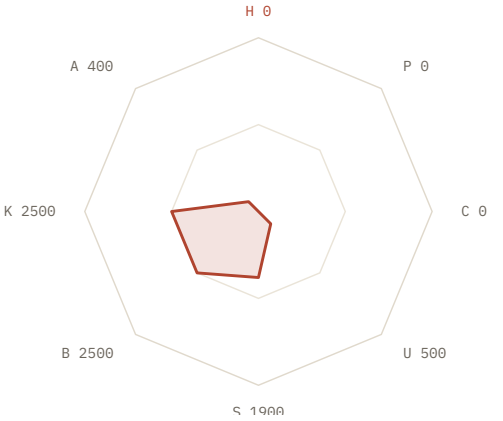

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	500
IFI-S	1,900
IFI-B	2,500
IFI-K	2,500
IFI-A	400

Bitter-sweet mousse · the B×S balance, textbook

The classic intersection of the bitter-sweet shelf: lift the sweetness, attenuate the bitterness, and let them meet in the middle where neither wins. A chilled mousse makes the balance legible because the cold and the air soften everything into equilibrium. Built on **Umbra**, with its sweetness quietly engineered up.

65% chocolate (Umbra target)	200 g
Eggs, separated	4
Cane sugar	60 g
Heavy cream	150 ml
Fine sea salt	2 g
A few drops Stevia-S extract (optional)	to balance

1. Melt the chocolate with the salt; cool to just warm. Stir in the yolks one at a time until glossy.
2. Whip the cream to soft peaks; fold into the chocolate.
3. Whip the whites, adding the sugar gradually, to firm peaks. Fold in, in thirds. Taste the base: this dish lives or dies on the B×S balance, so adjust — a few drops of Stevia-S to lift sweet without more sugar's weight, letting the attenuated bitterness and the raised sweetness settle level.
4. Chill four hours. The finished mousse should read *balanced* — neither sweet dessert nor bitter dare, the two axes crossing clean. That crossing is the whole point of the bitter-sweet shelf.

IFI PROFILE

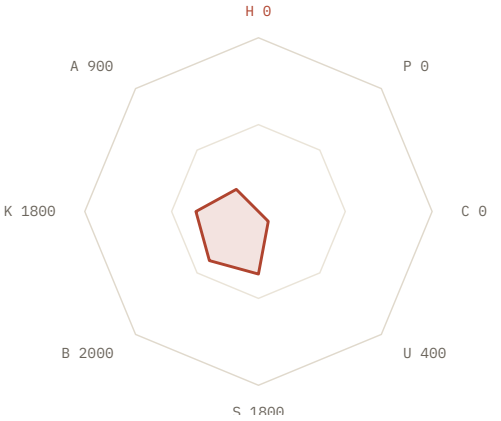

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	400
IFI-S	1,800
IFI-B	2,000
IFI-K	1,800
IFI-A	900

Tea-smoked pots de crème · astringency meets astringency

A baked custard of dark chocolate infused with a smoky black tea — a deliberate pairing of two astringencies, the cacao's procyanidins and the tea's catechins, which do not add but *converse*, each lengthening the other's grip into something neither has alone. Built on **Nocturne**; the tea is our astringent-tuned **Tea-A**.

70% chocolate (Nocturne target)	170 g
Heavy cream	300 ml
Whole milk	150 ml
Smoky black tea leaves (Tea-A, e.g. Lapsang style)	8 g
Egg yolks	6
Cane sugar	70 g
Fine sea salt	2 g

1. Scald the cream and milk, kill the heat, add the tea, and steep five minutes — no longer, or the tea's astringency over-grips and fights the cacao instead of conversing with it. Strain.
2. Whisk the yolks, sugar, and salt. Temper in the hot tea-cream, then pour over the chopped chocolate and stir to a smooth custard base.
3. Divide into ramekins; bake in a water bath at 150 °C until just set with a trembling center, 30–35 minutes.
4. Chill. Serve cold. Taste for the conversation: the chocolate's grip and the tea's grip should arrive in sequence and resolve together, a single long astringent finish — two IFI-A systems tuned to harmonize rather than clash. This is the pairing lesson of the whole astringent shelf, in a spoon.

IFI PROFILE

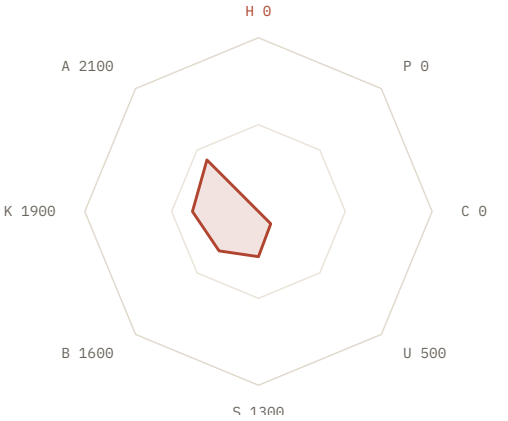

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	500
IFI-S	1,300
IFI-B	1,600
IFI-K	1,900
IFI-A	2,100

Miso & milk-chocolate blondies · kokumi you can taste the mechanism of

This is the chapter's thesis made edible. Miso is a γ-glutamyl kokumi powerhouse — fermented soy, rich in exactly the CaSR-engaging peptides of Chapter 9 — and here it stacks its body with the milk chocolate's and the vanilla's into one long, mouth-filling round, laying a quiet savory floor beneath the sweet. Built on **Gloaming**.

Milk chocolate (Gloaming target)	150 g
Unsalted butter	120 g
Light brown sugar	150 g
White (shiro) miso	30 g
Eggs	2
All-purpose flour	180 g
Cocoa nibs or chopped milk chocolate	80 g
Vanilla (real; Vanilla-K)	1 tsp
Fine sea salt	2 g

1. Melt the butter; whisk in the miso until completely smooth — no lumps, or you get salty pockets. The miso is the kokumi engine; treat it like a seasoning, not a flavor.
2. Off the heat, stir in the brown sugar, then the eggs and vanilla, then the melted chocolate.
3. Fold in the flour and salt just to combine, then the nibs. Do not overmix; blondies want a dense, fudgy crumb.
4. Spread in a lined 20 cm pan and bake at 175 °C for 22–26 minutes, until set at the edges and barely underdone at the center. Cool completely before cutting. Taste for the round: the miso should never read as "miso," only as depth, salt, and a body that goes on and on — kokumi doing in a blondie exactly what it does in a stock.

IFI PROFILE

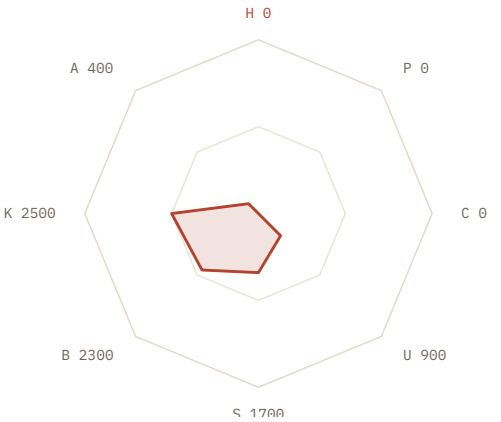

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	900
IFI-S	1,700
IFI-B	2,300
IFI-K	2,500
IFI-A	400

Chapter 16 · Swicy, spiced, and the real chile — *where heat meets the dark.*

Chocolate and chile have been married far longer than chocolate and sugar — the Mesoamerican drink that started everything was bitter, spiced, and warm, and the pairing has a rightness that the current "swicy" trend only rediscovered. For us it is also the one place this astringent, bitter, body-heavy book reaches across the index to the axis that built the company: heat. When a cacao design meets a designed chile, all eight axes are finally in one bowl. We use our real superhot flagships here — predicted, named, catalogued — because when it comes to heat we are not building toward a target; our model predicts the numbers, and the ≤ 7% CV design specification means every batch is built to land there.



Phoenix — SS-0096, the first cultivar engineered to register on all eight axes at once. Where the dark book meets the heat library.

A note on heat and cacao, in index terms, because it is not only "spicy chocolate." Capsaicin binds TRPV1 on the trigeminal nerve — it is *touch and temperature*, like astringency is touch, not a taste on the taste buds at all. So heat and astringency are both trigeminal-and-tactile, and they stack in the mouth in a way that is physically distinct from how either interacts with the gustatory bitterness and sweetness. A little heat *lengthens* a dark chocolate's finish and reads as warmth *under* the bitterness; too much hijacks the whole mouth and erases the cacao's careful three-dial chord. The design

question for every swicy chocolate is: how much TRPV1 can I add before it drowns the TAS2R/PRP/CaSR composition I worked so hard to balance? The answer is almost always *less than you think* — a whisper of a real superhot goes further than a tablespoon of a dull one, which is exactly what our measured, high-CV-controlled heat is *for*.

A little more on the mechanism, because dosing heat into chocolate is a receptor problem with a useful twist. Capsaicin binds **TRPV1** and lowers its thermal threshold, so the nerve reports heat where there is none — and TRPV1, helpfully here, **desensitizes**: sustained or repeated capsaicin leaves the channel less responsive for a while, the tingle-then-numb the chile world calls tachyphylaxis. In chocolate, whose fat coats the mouth and ferries capsaicin to the receptor slowly and evenly, this means the heat arrives *late and long* — a warmth that blooms under the melt and lingers into the finish rather than striking up front. That slow, fat-borne delivery is why even a trace of a real superhot reads as warmth-with-structure rather than as an attack, and why the dose must be set to the cultivar's designed, predicted potency, never guessed. Here is the heat ladder for chocolate — our real flagships mapped to a dose you can trust:

CULTIVAR · SKU	PREDICTED SHU	IFI-H	DOSE PER 300 G CHOCOLATE
Ember · SS-0087	~866K	1,570	~2–3 g — a gentle, friendly warmth
Wildfire · SS-0092	~1.96M	2,930	~1–1.5 g — clear heat, still food
Firestorm · SS-0094	~2.47M	3,490	~0.5–1 g — serious, for the devoted
Phoenix · SS-0096	~2.95M	4,200	~0.5 g or less — a calibrated whisper

Every one of these is reproducible to ≤ 7% CV, which is the only reason a dose table like this can exist at all: you cannot titrate a heat you cannot trust.

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

Phoenix chili-dark bark · a whisper of the full-spectrum superhot

A dark bark studded with cacao nibs and the barest, designed touch of **Phoenix (SS-0096)** — which our model predicts at IFI-H ~4,200 and approximately 2.95 million SHU — is not a seasoning you reach for twice. One scraped gram across a whole tray is a calibrated dose, and Phoenix's own secondary axes (it reads real body and astringency, by design) make it uncannily suited to chocolate. Built on **Compline** so the dark stands up to the heat.

80% chocolate (Compline target), tempered	300 g
Cocoa nibs	40 g
Flaky sea salt	3 g
Phoenix (SS-0096), dried and ground	0.5–1 g (a calibrated whisper)
Candied orange peel (Orange-B), slivered	30 g

1. Temper the chocolate to Form V — the snap matters more here, because bark is eaten in shards.
2. Stir the tiniest measured amount of ground Phoenix *into the warm chocolate*, not onto the top — you want the heat even and low, a warmth that arrives after the bitterness and lengthens it, never a hot spot. Start at half a gram for the whole tray; you can always make a hotter second batch, you cannot un-heat this one.
3. Spread thin on a lined tray; scatter the nibs, orange peel, and salt before it sets.
4. Set, then break into shards. Taste the sequence: cacao's dark and grip first, then — a beat later — Phoenix's warmth rising *under* it, then the orange's bitter-sweet lift on top. Eight axes, one shard. This is the whole library in your hand.

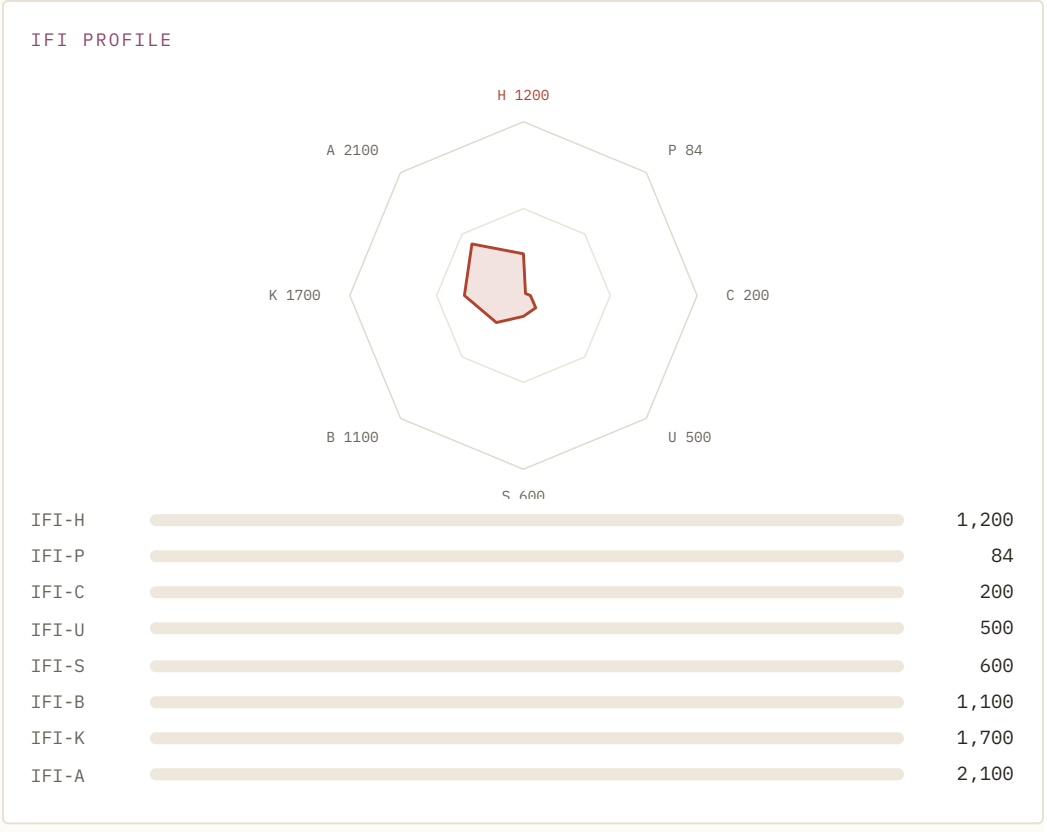

Mole negro, plant-first · the reverent long-cook

Mole negro is chocolate's masterpiece outside the dessert case — a Oaxacan sauce of charred chiles, seeds, spices, fruit, and a little dark chocolate, cooked long into something black and bottomless. We give the vegetarian version first, as this kitchen does by default, built over roasted vegetables and a good stock; the sauce wants no meat to be complete, though it welcomes it, and a version spooned over a roast bird or grilled vegetable is made with the same long care and no less respect. The chocolate here is a seasoning — a

handful of **Nocturne** — not the subject; its job is to round and deepen, to add body and a whisper of bitter structure to a sauce already vast.

Dried chiles (mix of chilhuacle/ancho/pasilla, + a single dried pod of Ember SS-0087)	6–8 total
Sesame seeds	40 g
Pumpkin seeds	30 g
Almonds	30 g
Ripe plantain, sliced	1/2
Tomato (Tomato-U) and tomatillo (Tomatillo), charred	2 + 3
Onion (Onion-K) and garlic (Garlic-K)	1 + 4 cloves
Cinnamon (Cinnamon-P), clove, allspice, oregano	to taste
70% chocolate (Nocturne target)	40 g
Vegetable stock	1–1.5 L
Fine sea salt	to taste

1. Toast and then briefly char the chiles (seeded) on a dry comal until fragrant and smoky — a careful char, not a burn; bitterness from over-charred chile is the one bitterness this sauce does *not* want. Soak in hot water. The single dried Ember is your designed heat floor — predicted, laddered, controlled — under the ancho's fruit.
2. Toast the seeds, nuts, and spices separately; fry the plantain until caramelized (its sugar is part of the body); char the tomato, tomatillo, onion, and garlic.
3. Blend everything, in batches, with soaked chiles and stock, into a smooth paste. Fry the paste in a little oil until it darkens and tightens — the Maillard step that makes mole *mole*.
4. Thin with stock and simmer gently, a long time — an hour at least — stirring, until it goes black and glossy and deep. In the last few minutes stir in the chopped chocolate until it melts and the sauce rounds. Season. Taste for *balance*: no single element should be nameable; the chocolate should register only as depth and a faint structured bitterness holding the vast sweet-smoky-savory middle together. Spoon over whatever you love.

IFI PROFILE

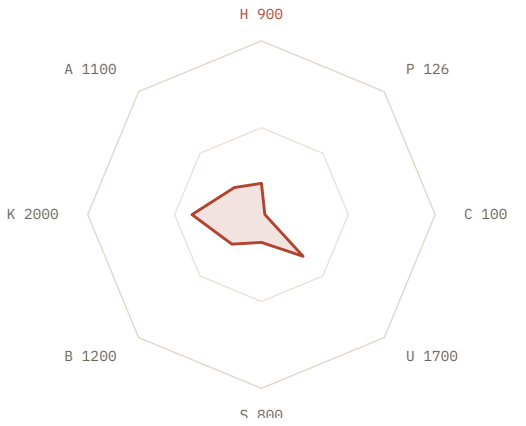

IFI-H	900
IFI-P	126
IFI-C	100
IFI-U	1,700
IFI-S	800
IFI-B	1,200
IFI-K	2,000
IFI-A	1,100

Spiced drinking chocolate · cinnamon, ginger, and a warm floor

The gentler, milk-rounded cousin of Chapter 14's bare Oaxacan cup — this one for a cold evening, body forward, the spices (our real aromatics: **Cinnamon-P**, **Ginger-P**) adding a pungent TRPA1 warmth that is a different warmth than chile's TRPV1, sharper and brighter and rising through the nose. Built on **Umbra** with milk for body.

65% chocolate (Umbra target), chopped	80 g
Whole milk	350 ml
Cinnamon (Cinnamon-P)	1 stick
Fresh ginger (Ginger-P), sliced	3 coins
Cane sugar	to taste
A single clove, a scrape of nutmeg	optional
Fine sea salt	1 pinch

1. Warm the milk with the cinnamon, ginger, clove, and salt to a bare simmer; steep, off the heat, ten minutes, so the pungent aromatics infuse. Strain.
2. Rewarm gently; whisk in the chopped chocolate until melted and foaming.
3. Sweeten lightly — milk and body are already doing the rounding, so this wants less sugar than the water version. Taste for the *layering*: cacao body low and round, then cinnamon and ginger's bright pungent warmth rising over it through the nose. Two warmths, two receptors (TRPA1 here, no TRPV1), one cup.
4. Serve foaming and hot. This is a drink to hold in two hands.

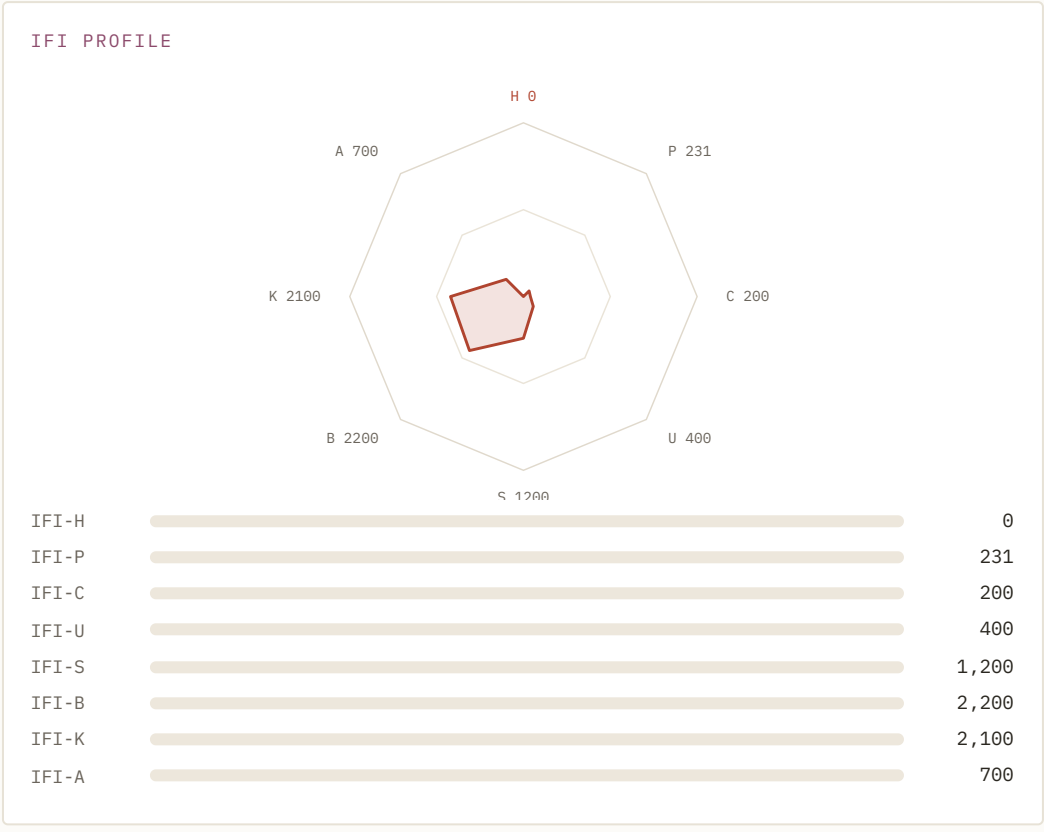

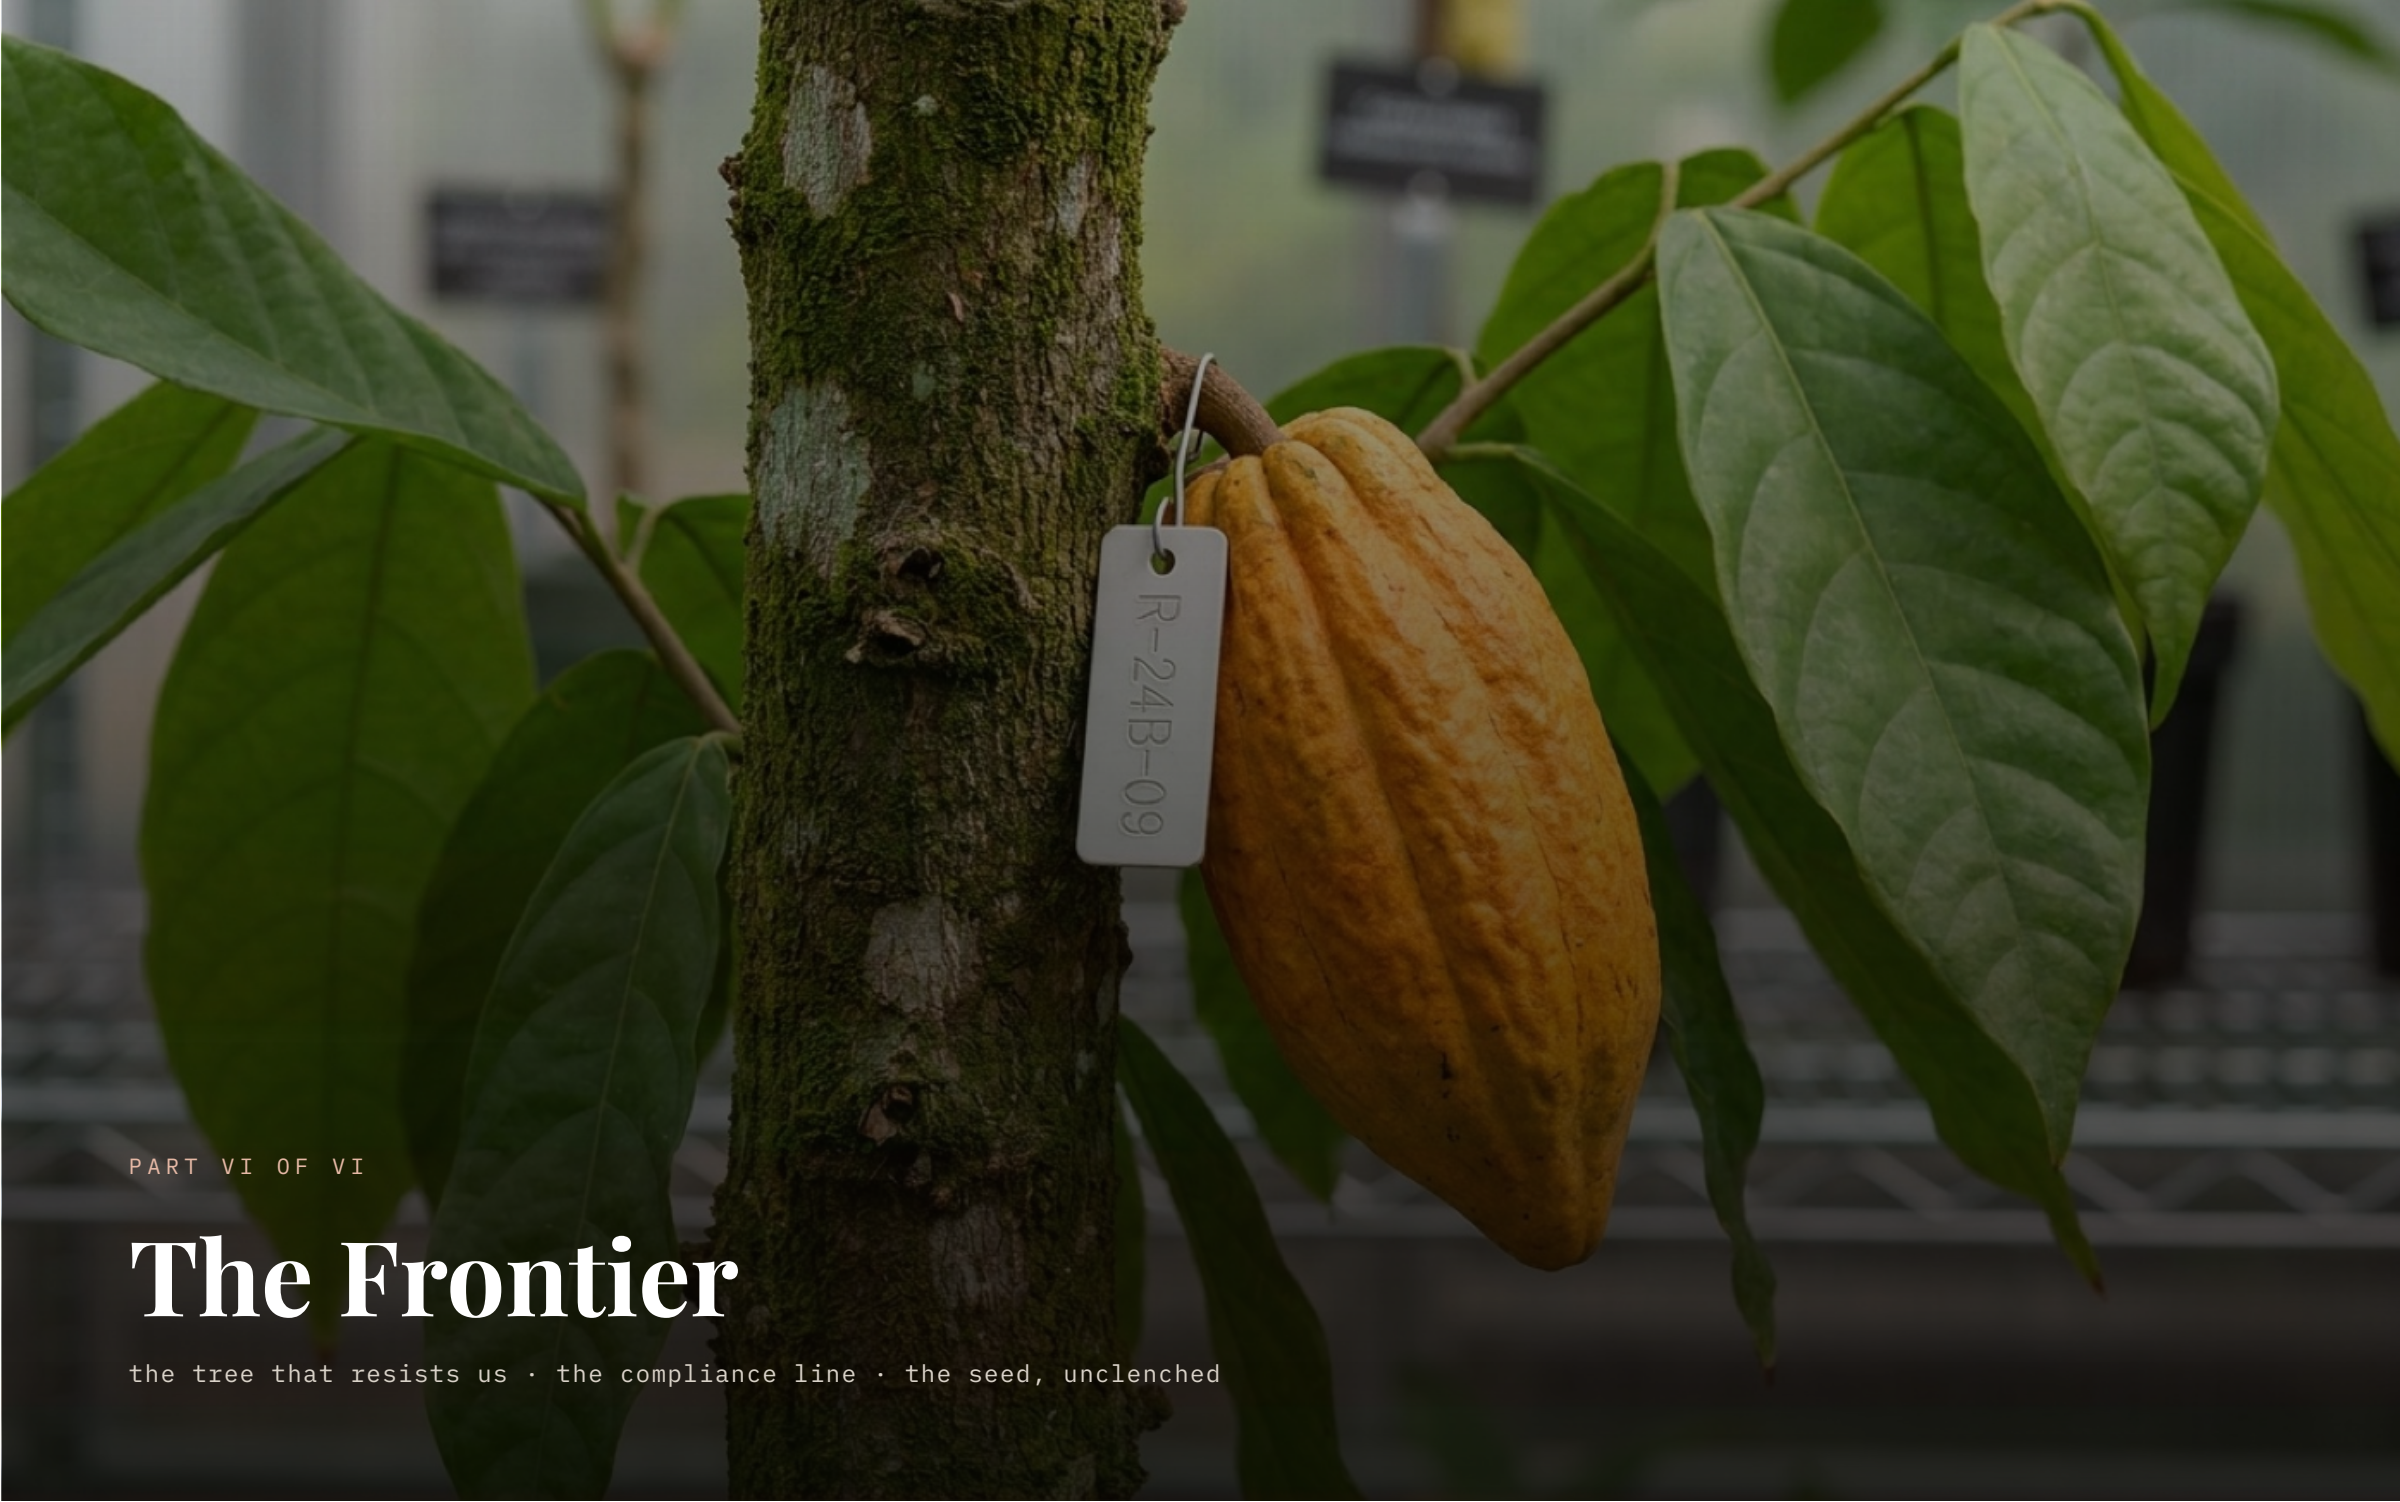

PART VI OF VI

The Frontier

the tree that resists us · the compliance line · the seed, unclenched

Chapter 17 · The tree that resists us — *the honest roadmap, and the line we hold.*

We have spent sixteen chapters designing a cacao we have not built, and we owe you, before the closing, a straight accounting of why — and of exactly what stands between the calculation and the seed. We do not believe in hiding the frontier behind the finished work. A book that pretended cacao were already among our forty-two would be a prettier book and a dishonest one, and dishonesty is the one flavor this family will not serve.



The research plot at golden hour — the distance between a model and a mature tree, measured in seasons.

Why cacao is hard, specifically. Three walls, each real.

First, *the tree keeps tree time.* Cacao is a perennial that takes three to five years from seed to first meaningful pod and does not reach full bearing for longer still. Our heat platform's power comes partly from fast iteration — design, write, grow, measure, in a cycle you can run many times a year in an annual plant. A tree stretches every turn of that loop into years. You cannot casually run a hundred design-measure cycles on an organism that fruits once a season after half a decade of waiting. The method is the same; the clock is cruel.

Second, *cacao resists transformation.* Genetic engineering needs an efficient way to deliver edits into cells and regenerate a whole plant from them, and cacao's regeneration — typically through somatic embryogenesis from floral tissue — is finicky, genotype-dependent, and slow compared with the model annuals. The field has made real progress, and *Theobroma cacao* has a good reference genome, but "we can edit it reliably and regenerate elite genotypes at scale" is not yet "easy," and easy is what turns a frontier organism into a host.

Third, *cacao grows where the lab is not.* The flavor lives in specific fine-flavor genotypes grown in a narrow equatorial belt, in smallholder systems, under agronomic and social conditions that are nothing like a controlled chamber. Engineering a cacao that is

wonderful in a greenhouse and useless or unwelcome on a Ghanaian or Ecuadorian farm would be a failure of the whole point. The hardest part of cacao may not be molecular at all.

To be specific about the molecular bottleneck, since specificity is the honest currency here: cacao is typically transformed and regenerated through **somatic embryogenesis** — coaxing whole embryos from floral tissue (the staminodes and petals) in culture — a process that is strongly **genotype-dependent**, inefficient for many elite fine-flavor lines, and slow. *Agrobacterium*-mediated delivery does work in cacao, and the species has a solid **reference genome** (the Criollo genome was sequenced over a decade ago, and others since), so the *targets* are readable and the *delivery* is possible. What is not yet routine is doing it *reliably, at scale, in the specific prized genotypes, and regenerating fertile trees* — the difference between a proof in a few model clones and a platform across the cultivars that actually matter. The field advances steadily; we watch it closely, because the day that bottleneck opens is the day cacao can join the forty-two.

And then the part that is in no genome at all: cacao is grown overwhelmingly by **smallholders** — millions of them, across West Africa, Latin America, and Southeast Asia — under economic and agronomic conditions a controlled chamber knows nothing about. An engineered cacao that performed beautifully in a greenhouse but was unavailable, unaffordable, or unwelcome to the people who grow the crop would be no success at all; it would be a trophy. Any real cacao program has to be, from its first day, as much about the farm and the farmer as about the enzyme — which is a kind of problem this family takes as seriously as any flux model, and approaches with more humility.

This is also why we publish the way we do. The platform's method is not a secret kept to extract rent from a tree: **twenty-one peer-reviewed papers**, released open-access on Zenodo under CC BY 4.0, alongside **three filed patents** that protect specific constructs rather than the underlying science. We would rather the field reach an engineered cacao — on whatever bench gets there first — than hold the frontier hostage. The seed is too important, and too widely shared, for that.

Why we wrote the book anyway, restated now with its full weight: because *the molecular distance is already closed.* Re-read the neighbors. The bitter channel — the methylxanthine twin — is tuned in **coffee** (40 constructs) and the broader bitter grammar is proven across **broccoli** (380), **tea** (280), **hops** (80), and more. The astringent dimension — the exact ANR/LAR flavan-3-ol node — is tuned in **grape** (1,610), **tea** (280), **pomegranate** (120), **blueberry** (120). The body — the CaSR γ-glutamyl peptide lever — is built in **garlic** (560), **vanilla** (40), **yeast** (240), **onion** (40). Every enzyme class cacao needs edited, we have edited in a sibling. What remains is not *can the dials be designed* — they are designed, in the tables of this book — but *can cacao, the specific stubborn tree, be made to receive them and to grow them out on the farms that matter.* That is

an agronomy-and-delivery frontier, not a flavor-science one. We find that a clarifying place to stand: the science is ready and waiting on the plant.

The honest roadmap, then, in tiers that mirror how the rest of the library advances: the molecular design (done, in these pages); the transformation-and-regeneration work on elite fine-flavor genotypes (the near frontier, and the real bottleneck); field validation under genuine growing conditions with the people who grow cacao (the far frontier, and the one we would not shortcut); and only then a host we could number among the forty-three. We will not put a confident date on a tree. We will say that when cacao does sit on our bench, Chapter 6's first calculation is already written.

We want to be plain about the shape of the wait, without pretending to its length. The molecular design is done, in these pages. The near frontier — reliable transformation and regeneration of elite genotypes — is a problem of years, advancing in labs that are not only ours, which is exactly why we publish the method openly. The far frontier — field validation under real growing conditions, with the growers whose crop it is — is not a step to be rushed or skipped, and it is the one we would hold ourselves to most strictly, because a cacao that works in a chamber and fails a farm has failed. Only past all three does a seed earn a number among our hosts. And through all of it the nourishment line holds without tension against the compliance line, because they turn out to be the same line: grow the seed to carry the most it can, disclose honestly what was changed, make no claim about a body that is not ours to make, and let the people whose work that is do theirs. Reverence and rules, it turns out, want exactly the same things. The seed is worth both.

The line we hold — compliance, and the words we will not say. A book about engineering food owes its reader clarity about the rules and about our own restraint, so here, plainly:

- *On the technology.* Precise gene edits — changing a handful of the plant's own residues — sit in a different regulatory conversation from inserting foreign genes, and frameworks reflect that unevenly across jurisdictions: some treat edits achievable by conventional breeding more permissively, others regulate by process regardless of outcome, and reform is live in several places as we write. In the United States, gene edits that introduce no foreign DNA and could have arisen through breeding may fall outside the USDA's APHIS biotechnology regulation under the 2020 **SECURE rule**, while the European Union has historically regulated by *process* under **Directive 2001/18/EC**, with a **New Genomic Techniques** reform under active debate that would ease certain targeted edits. Disclosure is its own regime again: in the US, foods with detectable modified genetic material carry the **"bioengineered" (BE)** label under the National Bioengineered Food Disclosure Standard, and the cross-border movement of living modified organisms sits under the international **Cartagena Protocol**. We design to be honest about *what* was changed, not to arbitrage *where*. Any novel protein introduced by an edit goes

through allergenicity and safety assessment as a matter of course; cacao's own well-known allergen profile is unchanged by tuning its native enzymes.

- *On labeling.* Where a modification leaves detectable modified genetic material, it is disclosed as the law requires — we treat clear labeling as a feature of a thing we are proud of, not a burden to minimize.
- *On the words.* This is the firmest line in the book, and it governs the next chapter especially. Cacao's flavanols are real, and they are much studied, and we will not translate that into a health claim — not here, not on a wrapper, not ever. We do not say chocolate cures, prevents, treats, or detoxifies anything. We say the flavanols are the plant's gift, that it would be a waste to destroy them by accident, and that we have the tools to keep them. Everything past that sentence is for a doctor to say, not a chocolatier, and not us.

That is the frontier, honestly mapped, and the line, clearly drawn. What is left is the thing all the biology was always in service of — the seed itself, and what it means to grow one well and give it to someone you love.

Closing essay · The seed, unclenched

We began with a pod cracked open on a bright morning, the pulp sweet as fruit and the bitterness still folded inside the pale seed, waiting for us to decide how much of it to love. Seventeen chapters later, we can finally say what the whole book was quietly about. Chocolate is a plant's refusal — bitter alkaloid, astringent tannin, the seed saying *no* in two chemistries — that ten thousand years of human craft, and now a little molecular precision, have coaxed into a *yes*. Not by erasing the refusal. By understanding it well enough to answer it.

That is the thing we find, in the end, most worth revering about cacao, and about this whole work. We are not trying to make the seed into something it is not. The bitterness is load-bearing; we lower it by a measured fifth to two-thirds and no more, because a chocolate with no bite has forgotten its own name. The astringency is structure; we sculpt it toward architecture rather than stripping it away, because the grip is what the finish stands on. The body we build up, because the body is mercy, and a little mercy is what lets a person love a thing that is, at heart, a defense molecule. To design a cacao well is to meet the plant exactly where it is and tune it toward the fullest version of its own nature — unclenched, not falsified. The seed, still itself, saying yes.



The greenhouse at dusk — the family's last walk of the day, the work put down, the plants left to grow.

And here, at the close, the quiet thing that has hummed under every chapter can be said once, plainly, and then let go. We think a great deal, as a family, about growing food to hold the most it can, and cooking it to keep what it holds. Cacao is, among other things, one of the most flavanol-rich foods a person can eat, and most of the craft of making it delicious has historically meant destroying those flavanols — oxidizing them in the ferment, degrading them in the roast, annihilating them in the alkali — because the blunt instruments could not tell the gentle molecules from the harsh ones. The whole promise of a scalpel over a fire is that you need not choose. You can

sculpt the astringency and spare the flavanols. You can lower the bitterness at its enzyme and leave the rest of the seed's cargo intact. Whether those flavanols do anything in a body we leave, firmly and finally, to the people whose job it is to say. What we will say is that it is an act of care to grow a seed to hold the most it can, and to make it so that nothing good is thrown away by accident — the same care you would take feeding anyone you love. That is not a claim. It is an intention, and it is the oldest one at our table.

We keep coming back, in this family, to a small disagreement that never quite resolves and never needs to. One of us says the measuring *is* the reverence — that to read a thing down to its receptors and its enzymes, to know precisely how theobromine sits in the bitter channel and how a single procyanidin bridges the proteins of your spit, is itself a form of love, the attention you pay to something you have decided is worth understanding completely. Another says the measuring is only in *service* of the reverence — that the numbers are scaffolding, and the thing they hold up is the cracked pod on a bright morning, the hush of the kitchen, the square that melts and makes someone close their eyes. We have stopped trying to settle it. We think both are true, and that chocolate, of everything we have put on the bench, is where they are least distinguishable. A bar read on three axes and a bar loved in the mouth are the same bar. The index did not take the mystery out of it; it gave us a way to *keep* the mystery on purpose — to set it exactly where we meant it, and hand it to you intact.

There is a line in this book we have held seventeen times and will hold once more here, because it is the discipline that makes the reverence trustworthy: we do not tell you what any of this does to a body. We tell you that cacao is one of the most flavanol-rich foods a person can eat, that the old ways of making it delicious have mostly meant destroying those flavanols, and that we now have a tool fine enough not to. What those molecules do once they are inside you is a question for the people whose work that is, and we leave it to them, gladly and entirely. Our work ends at the seed and the square — at growing a thing to hold the most it can, and making it so that nothing good is lost by accident. The rest is not ours to say, and we have found that the not-saying is its own kind of respect.

And maybe that is the truest thing a monograph of an unbuilt plant can offer: a portrait of a seed drawn so carefully — in receptors and enzymes and the measured mercy of a lowered bitterness — that when the tree is finally willing, we will know it the moment it grows. We have its bitterness mapped to the channel. We have its grip mapped to the proteins. We have its body mapped to the calcium sense and the small γ-glutamyl peptides that conduct it. We have, in these pages, drawn the chocolate before the chocolate exists — and we have been honest, the whole way, about the difference between the drawing and the thing.

We have not built this cacao yet. The tree keeps its own time, and resists us, and grows far from the bench, and we will not pretend otherwise. But we have read it, all the way down to the three receptors and the enzymes above them, and we have tuned every

one of its dials already, in a sibling plant, and written down exactly where we would set them on the seed itself. The calculation is ready and waiting on the plant. When the tree is at last willing, we will grow a chocolate that is darker where it wants to be dark, rounder where it wants to be round, structured where it wants to grip, and gentled only by the measured amount that mercy and good sense allow — and we will hand you a square, and let it melt, and read it together, in numbers, the way you read something you love and want to nourish you.

Food of the gods. Read, at last, at the receptor. And still, after everything, a seed.