

THE SCOVILLE SPLICE LIBRARY · VOLUME LXII

The Nourishment Frontier

Eight axes for how it tastes. *A ninth for what it holds.*

The complete account of designing nutrition into the plant itself — density, bioaccessibility, and the biosynthetic routes to vitamins, minerals, protein, and color. A food-first measure of nourishment, built beside the flavor index, for what we grow next. Composition and availability, measured — never a claim about a body.

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

Contents

01 Opening essay · The ninth question

02 Chapter 1 · What nourishment is, as a number — content is only half the sentence.

03 Chapter 2 · The matrix problem — where the nutrient hides, and what it takes to let it go.

04 Chapter 3 · Antinutrients, honestly — the molecules that hold the minerals hostage.

05 Chapter 4 · Provitamin A — the carotenoid, the chromoplast, and the twelve-to-one tax.

06 Chapter 5 · Folate — two branches, one vitamin, and the fragility of a gift.

07 Chapter 6 · Vitamin E and vitamin C — the one that guards the fat, and the one that frees the iron.

08 Chapter 7 · Iron — the mineral the plant hides from itself.

09 Chapter 8 · Zinc and the rest — the quieter minerals, and the ratios that rule them.

10 Chapter 9 · Protein quality — the one amino acid the whole protein waits on.

11 Chapter 10 · The good fats — building the long chain the plant never learned to finish.

12 Chapter 11 · Fiber, resistant starch, and the whole matrix — the nutrition that is mostly architecture.

13 Chapter 12 · Anthocyanins and the purple design — the color is the compound is the flavor.

14 Chapter 13 · Glucosinolates and the Brassica lever — the sharp compound that is bitterness, pungency, and design tension at once.

15 Chapter 14 · What we build next — the roadmap, read as nourishment.

16 Chapter 15 · Designing the whole plate — nourishment is a meal, not a molecule.

17 Chapter 16 · The honest edge — the discipline that keeps a nourishment program trustworthy.

18 Closing · The table we are growing toward

Opening essay · The ninth question

There is a moment at our table, most nights, that the eight axes cannot explain. The food is good — we built the index precisely so that *good* would stop being a guess. The tomato is as umami as a tomato should be; the heat arrives on schedule; the bitterness that used to flatten the greens has been talked down to a whisper. We can read all of it now, each sensation anchored to a receptor and written as a number. And still, when the plates are nearly clean and the kitchen has gone quiet, someone in our family asks the question that started this book.

What else is it giving them?

For most of human history that question had no instrument behind it, only faith and folklore — eat your greens, finish your crusts, the broth is good for you. We believed it the way everyone's grandmother believed it, which is to say completely and without a single number. The Iyengar Flavor Index was built to end exactly that kind of silence on the side of taste. This book is the record of us turning the same discipline toward the other half of the plate: toward what a food *carries*, and how much of it a body can actually reach.



The walk we take most mornings, before the lights come up — the question that became this book is usually asked somewhere along this row.

We should say plainly where this book sits in the library, because it does not stand alone. Across forty-two characterized host organisms and around 225 distinct construct designs backing roughly 19,500 product SKUs, the Scoville Splice platform has spent its first years learning to design *flavor* — eight sensory axes, eight receptors, a thousand cultivars live and purchasable today out of some seventy-eight thousand catalogued. Everything in those years was practice for a second question we always knew was coming. The organism we taught to carry a designed capsaicinoid is an organism whose folate branch we now understand; the regulatory switch we throw to turn a fruit purple is the switch that turns a

vitamin on. The Nourishment Frontier is not a departure from the flavor work. It is the flavor work, grown up, and pointed at the half of the plate we had not yet learned to measure.

We want to be exact, at the outset, about what this is and what it is not, because a book about nutrition is a book that can do harm if it is loose with its language. The eight axes measure a sensation — the *loudness* with which a compound speaks to a receptor on the tongue. Nourishment is not a sensation and has no single receptor. You do not taste folate. You do not perceive the iron in a bean. So the ninth question cannot be answered the way the first eight were, and we will not pretend that it can.

Instead we measure two things that are genuinely measurable, and we refuse — all the way through these sixteen chapters — to measure a third that is not. We measure how much of a nutrient a plant **holds** (its density, in milligrams or micrograms per hundred grams, a property of the food you could weigh out on a bench). And we measure how much of that a standardized digestion in a tube will **release and solubilize** from the food's own matrix (its bioaccessible fraction, a number between zero and one, produced by a protocol that does the same thing to every sample). Density times availability is the whole of our claim. What a particular person's particular body does with that released nutrient — their status, their health, their outcome — depends on that person, their diet, their genetics, and a hundred things no plant breeder controls. We design the food. We do not prescribe the result. There are no cures in this book, no detoxes, no promises about anyone's body. There is a plant, grown to carry more, in a form more of it can be reached from — and there is the arithmetic to prove it.

That arithmetic turns out to rhyme, almost exactly, with the flavor index we already use every day. The flavor index is an inverted ratio against a fixed receptor reference, scaled by a thousand:

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

and the nourishment measure takes the identical shape, pointed at composition instead of sensation:

```
N(candidate) = 1000 · BD(candidate) / BD(reference)
```

where **BD = C · f_access** — concentration times the bioaccessible fraction — and the reference is a fixed, named baseline cultivar we anchor everything against. A nourishment score of 1,000 means parity with that reference. Above a thousand, the food holds more and releases more than the thing we grew it to replace. We did this on purpose. A family that measures flavor and refuses to measure nourishment is a family that has decided taste is the only thing worth being rigorous about, and we have never believed that. Where the IFI anchors to a receptor's half-activation, *N* anchors to a reference food's delivered content. Where the IFI measures how loudly a compound speaks, *N* measures how much a food gives up. Same shape, different frontier.

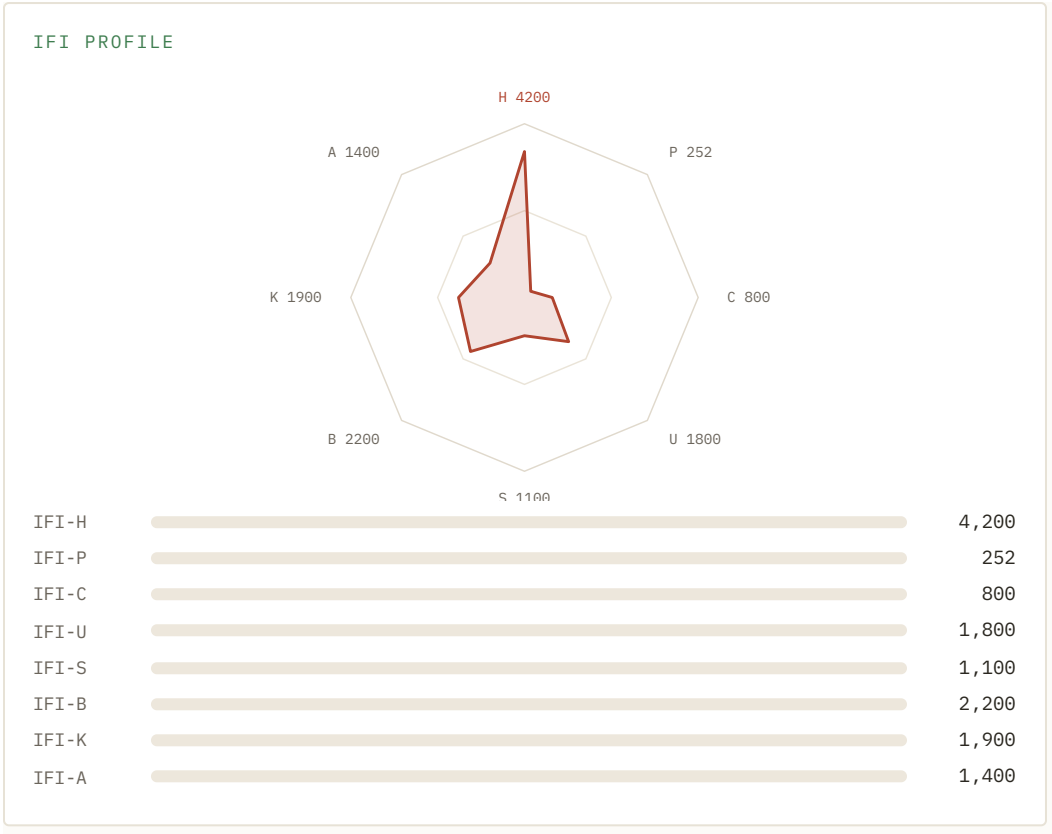

That scorecard is **Phoenix** (SS-0096), the cultivar that first registered on all eight flavor axes at once — our model predicts approximately 2.95 million SHU, eight axes lit, the proof that flavor could be composed rather than stumbled into, written with two AI-designed Pun1 variants (S39L · L345G) that no breeder ever found. We open a nourishment book with a flavor scorecard deliberately. The platform that wrote two Pun1 variants into a pepper's acyltransferase is the same platform that will, in these pages, write provitamin A into a grain's endosperm and talk a legume's phytate down until its iron comes loose. The method does not care whether the molecule it is designing is a capsaicinoid or a carotenoid. The receptor-level respect we paid to taste, we now pay to the nutrient.

One last word before we begin, because it is the word the whole shelf is built on. The ethos under this book — that food should be grown to carry the most it can, that a body should be able to draw the most from it — is a *value and an intention*, never a claim. We are a plant-reverent, vegetarian-first family, and that shows up in what the science cares about, not in any promise about what it does. The ninth question has waited ten thousand years for an instrument. Here is ours — and here, carefully, is everything we have learned to build with it.



PART I OF VI

The Measure

Bioavailable density · the in-vitro fraction · the anchored score · the matrix · the antinutrient

Chapter 1 · What nourishment is, as a number — *content is only half the sentence.*

The label on a can of spinach tells you it holds iron, and the label is not lying. It is simply answering the wrong question. The right question — the only one a body cares about — is not *how much iron is in this food* but *how much of that iron can my gut get its hands on*. For a century, food tables answered the first and let everyone infer the second, and the inference was often badly wrong. Spinach, famously, holds a great deal of iron and surrenders almost none of it, because the iron is locked behind oxalate the way a word can be locked behind a language no one at the table speaks. Content is only half the sentence. The other half is availability, and nourishment — the thing we actually eat for — lives in the product of the two.



A tomato's interior is not a bag of nutrients; it is an architecture that decides what comes loose and what stays bound.

So we define nourishment, for the purposes of design, as a quantity we can measure twice and multiply. The first measurement is **concentration**, which we will write as *C*: the mass of a target nutrient per hundred grams of food, fresh weight unless we say otherwise. Concentration is honest and boring. You homogenize the sample, you run it against a standard — HPLC for the vitamins and carotenoids, ICP-MS for the minerals, an amino-acid analyzer for protein quality — and you get a number that would be the same in any competent lab in the world. Concentration is where

biofortification has always lived, and for decades it was the whole game: push the number up, declare victory. But concentration alone is the spinach mistake institutionalized.

The second measurement is the one that makes this a book and not a fertilizer brochure. We call it the **bioaccessible fraction**, *f_access*, and it is a dimensionless number between zero and one: the proportion of the nutrient that a *standardized in-vitro digestion* releases from the food matrix into a soluble, absorbable form. We use the harmonized static protocol the field has largely converged on — an **oral phase** with α-amylase at near-neutral pH, a **gastric phase** at pH 3 with pepsin, an **intestinal phase** at pH 7 with pancreatin and bile salts — because its whole purpose is to do the identical thing to every sample so that two foods can be compared without a human being involved. The digest is centrifuged; the supernatant is the fraction that came loose; you measure the nutrient in it and divide by what you started with. That ratio is *f_access*. It is not a measurement of a person. It is a measurement of a food's willingness to let go.

Multiply them and you have what we actually set out to design:

```
BD = C · f_access
```

Bioavailable Density — how much of a nutrient a food both holds and releases. It has the same units as concentration, but it has been taxed by reality. A food can win on density and lose on BD, and the spinach on the first page of this chapter does exactly that. The entire thesis of the nourishment program follows from the fact that BD is a *product*: you can raise it by raising *C*, by raising *f_access*, or — and this is the move that distinguishes designing nourishment from merely fortifying it — by raising both at once, because the two levers are independent and they multiply. A 50% gain in content and a 50% gain in availability do not add to a 100% gain; they multiply to a 125% gain, and that compounding is the whole reason the frontier is worth the work.

To make BD a score a reader can carry between foods, we anchor it exactly as the flavor index anchors loudness. We pick a **reference cultivar** per nutrient — a fixed, named, ordinary baseline, the thing the designed food is meant to stand in for — and we take the ratio:

```
N(candidate) = 1000 · BD(candidate) / BD(reference)
```

A score of **1,000 is parity**. A score of 2,000 means the designed food delivers twice the bioavailable density of its reference. The number is deliberately the IFI's own number — an inverted, thousand-scaled ratio against a fixed anchor — so that anyone fluent in reading a flavor profile can read a nourishment profile without learning a second grammar. We chose the *ratio* form, rather than reporting absolute BD, for the same reason the IFI did: a ratio travels. It lets a tomato and a grain and a leaf be compared on one page, and it forces the honest discipline of naming a reference — because a nourishment number with no stated baseline is a number hiding its own goalposts.

Let us do two worked examples end to end, because the arithmetic is the argument.

Folate in a tomato. An ordinary field tomato holds roughly 15 µg of folate per 100 g, and its matrix is generous — the folate is in the cytosol, not locked behind a cell wall of lignin — so *f_access* is high, call it 0.80. Its BD is therefore 15 × 0.80 = 12 µg/100 g. A tomato engineered through the pteridine and PABA branches we will build in Chapter 5 can hold on the order of 25 times the folate — call it 375 µg/100 g — and because the added folate sits in the same friendly compartment, *f_access* holds near 0.80, giving a BD near 300 µg/100 g. The nourishment score is 1000 × 300 / 12 = **25,000** against the reference tomato.

Iron in a bean, where availability moves too. A reference bean holds 50 µg/g of iron but at a punishing phytate:iron ratio, so *f_access* is only 0.05; BD = 2.5 µg/g. A designed line holds 75 µg/g (ferritin and NAS, Chapter 7) *and* drops its phytate enough to lift *f_access* to 0.15; BD = 11.25 µg/g. N = 1000 × 11.25 / 2.5 = **4,500** — and notice that the content rose only 1.5× while the score rose 4.5×, because availability tripled on top of it. That multiplication, invisible to any label that reports content alone, is the entire case for measuring both.

TERM	SYMBOL	WHAT IT MEASURES	HOW WE READ IT
Concentration	C	nutrient mass per 100 g, fresh weight	bench assay — HPLC / ICP-MS / AAA
Bioaccessible fraction	f_access	fraction released by standardized in-vitro digestion	0 to 1; a property of the food's matrix
Bioavailable density	BD = C · f_access	content the food both holds and surrenders	same units as C, taxed by reality
Nourishment score	N = 1000 · BD/BD(ref)	delivered content vs a fixed reference cultivar	1,000 = parity; higher = denser-and-freer



We weigh everything. Density is the first number, and it is the one that cannot be argued with — but it is only ever the first.

There is a measurement ladder under the single symbol *f_access*, and we will always tell you which rung a number came from, because the rungs are not equally trustworthy:

RUNG	WHAT IT MODELS	HONESTY FLAG
C alone	total content	says nothing about release; the spinach trap
f_access (in-vitro digestion)	solubilization from the matrix	a model of the *food*, not the eater
Caco-2 cell assay	cellular uptake of what was solubilized	a better proxy (used for iron); still a dish
in-vivo (literature, not ours to run)	absorption in a person	belongs to clinical science, not to us

Two cautions travel with this measure everywhere it goes, and we state them here so that they are never far from any number in this book. First, *f_access* from a tube is a model, not a verdict. It controls for the food's matrix beautifully and for the eater's biology not at all — no gut microbiome, no transporters, no competing meal, no individual variation. Where a nutrient's absorption is known to depend on cellular uptake and not merely on solubility — iron is the great example — we escalate from the digestion tube to a Caco-2 intestinal-cell assay and measure the cells' own ferritin response, which is a better proxy but still a model. Second, and this is the discipline the whole shelf is built on: *N* is a score for a food, full stop. It belongs in the same sentence as "firmer," "sweeter," "holds its color." It does not belong in the same sentence as any verb about health. We built an instrument precise enough to be trusted, and the way you keep a precise instrument trustworthy is by never once pointing it at a claim it cannot support.

One more thing the measure anticipates, which the rest of the book will keep returning to: whether a plant can *afford* a design is itself a quantitative question, and we answer it with a genome-scale flux-balance model before we make a single edit. A nutrient is built from carbon and nitrogen the plant also needs for a hundred other things, and raising one pool can starve another. The measure tells us what a food delivers; the flux model tells us whether the plant could ever deliver it without falling over. The food holds more, and more comes loose, and the plant can still stand — three numbers, and we insist on all three before we call a design real.

Chapter 2 · The matrix problem — *where the nutrient hides, and what it takes to let it go.*

Cut a raw carrot and you are looking at one of the most concentrated stores of provitamin A in the plant kingdom, and also at one of the most reluctant. The β-carotene is there in extraordinary quantity, crystallized inside chromoplasts, packed so tightly it has gone almost solid — and a solid is exactly what a digestive tract struggles with. You can eat a raw carrot and recover a famously small fraction of its carotene, not because the carotene is sparse but because the *matrix* will not release it. This is the problem that organizes the entire middle of this book. A nutrient is never floating free in a food. It is held — in a crystal, behind a cell wall, bound to a protein, chelated by a mineral-grabbing acid — and the architecture of that holding decides *f_access* far more than the raw concentration ever will.

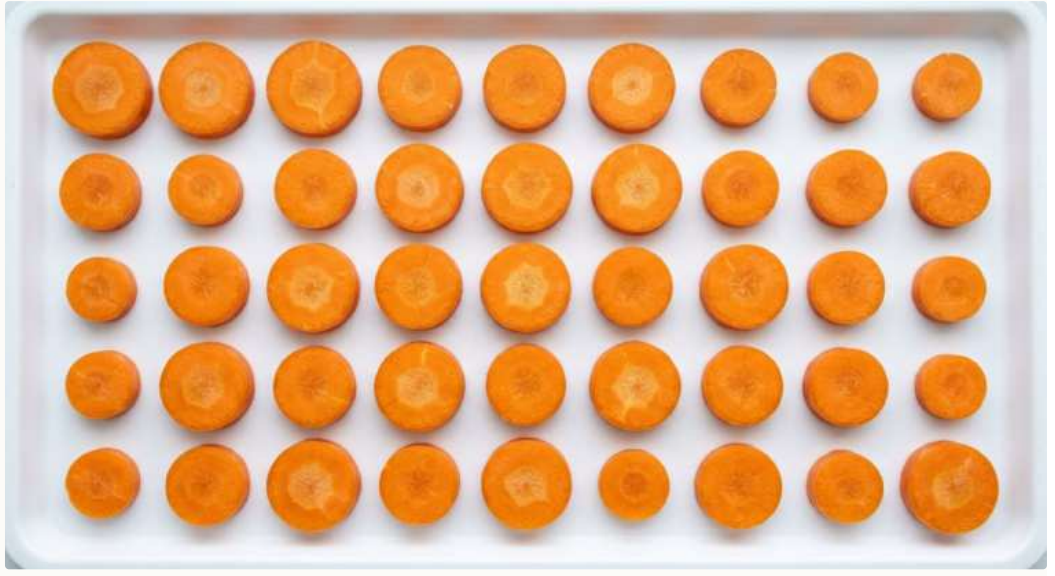

The carrot is the textbook case: among the densest provitamin A stores in any plant, and among the most tightly locked. Design has to address the lock, not just the quantity.

Consider what a carotenoid has to survive to count. It is a large, fiercely lipophilic molecule — it will not dissolve in the watery phase of the gut at all. For it to be absorbed, it must first be freed from its chromoplast, then partitioned into lipid droplets from the meal, then incorporated into bile-salt **mixed micelles** that ferry fat-soluble cargo to the intestinal wall. Every one of those steps is a place the nutrient can fail to arrive. Freed from a liquid-crystalline chromoplast, as in a ripe tomato, carotene micellizes relatively well. Trapped in a solid crystalline chromoplast, as in a raw carrot, it

largely does not. Eaten with no fat at all, even perfectly freed carotene has nothing to partition into and passes through. This is why the honest unit for a fat-soluble vitamin is never the milligram on the label; it is the milligram that reached a micelle, which is precisely what our *f_access* assay is built to count.

The lipid requirement is quantitative, not merely directional, and it is worth making concrete because it governs how a nutrient-dense food should actually be eaten. Carotenoid micellization rises steeply with co-ingested fat up to a point and then plateaus; the first few grams of lipid in a meal do most of the work, and a carotene-rich dish eaten with essentially no fat can deliver a small fraction of what the same dish delivers with a spoonful of oil. We treat that not as a flaw to engineer around but as a fact to design *with* — it is one reason the nourishment shelf is vegetarian-first but never fat-phobic, and one reason Chapter 15 builds the oil into the bowl rather than leaving it to chance.

The matrix levers, then, are as much a part of nourishment design as the biosynthetic ones, and they fall into a small, teachable set:

- **Compartment and crystal state.** Where the plant stores a compound decides how hard it is to free. Tomato's carotene, held in a liquid-crystalline chromoplast, is several-fold more bioaccessible than carrot's crystalline store at equal concentration. One of the most powerful things a designer can do is move a target's synthesis into a *friendlier compartment*, or co-express a protein that holds the carotenoid in a globular rather than crystalline form — which is why we care about the chromoplast type a carotenoid lands in, not merely how much lands there.
- **Cell-wall entrapment.** Nutrients inside intact plant cells are shielded by walls the gut's enzymes cannot breach. Rupturing cells — by chewing, chopping, and above all by cooking — multiplies release. Cooked carrot and cooked tomato both hand over far more carotene than their raw forms; the heat that dulls a vitamin C reading is the same heat that frees a carotenoid. Nourishment design has to hold the whole processing arc in view, because the food is cooked before it is absorbed.
- **The lipid requirement.** Fat-soluble vitamins (A, D, E, K) and their provitamins need dietary lipid to form micelles — the relationship above, written into every menu.
- **Chelation and binding.** Minerals are the mirror image of the fat-soluble vitamins: they are plenty soluble until something in the food grabs them. Phytate, oxalate, and polyphenols each bind specific minerals into complexes the gut cannot absorb. Here the design move is subtractive — lower the binder — and it is the whole of Chapter 3.

MATRIX STATE	EXAMPLE FOOD	RELATIVE F_ACCESS (ILLUSTRATIVE)	DESIGN OR KITCHEN MOVE
crystalline chromoplast	raw carrot	low	cook; design globular sequestration
liquid-crystalline chromoplast	ripe tomato	moderate–high	already favorable; protect on cooking
cell-wall entrapped	whole intact greens	low until ruptured	chop, cook, chew
mineral chelated	whole legume (phytate)	very low	phytase / lpa / ascorbate (Ch. 3, 7)
freely cytosolic	tomato folate	high	choose this matrix for fragile vitamins

Two further facts about carotenoids keep the matrix honest, and both belong on the food side of the ledger. The first is cooking, quantified: rupturing cells and melting the crystalline store can lift carotenoid bioaccessibility severalfold over the raw tissue, which is why we report *f_access* on the cooked food and treat a raw-tissue number as a curiosity. The second is the final cleavage itself. β-carotene becomes vitamin A only when the enzyme **BCO1** (β-carotene 15,15′-oxygenase) splits it in the intestinal wall — and the efficiency of that cleavage varies from person to person with their own *BCO1* genetics, sometimes by twofold or more. That variation is precisely why we stop our own measurement at *bioaccessible carotene* and convert to RAE only as a population average: the last step of the provitamin's journey happens inside the eater, in a gene we do not design and would never presume to speak for. We hand over a food that releases more carotene into a micelle. What each body's BCO1 then makes of it is, once again, the body's own business.

We measure every one of these the same way: standardized in-vitro digestion, on the food as it will actually be eaten, raw and cooked both, because a *f_access* measured on a raw sample of something no one eats raw is a number that lies politely. The through-line of this book's method is that we never let a density number stand alone. A gorgeous concentration behind a bad matrix is a food that photographs well and nourishes poorly, and the whole point of building BD as a *product* was to make that trade impossible to hide. When we design a nutrient-dense cultivar, half the engineering is in the pathway that makes the molecule and half is in the matrix that will — or will not — let it go. The two halves meet in a single number, and the number is the one we have learned to trust.

Chapter 3 · Antinutrients, honestly — *the molecules that hold the minerals hostage.*

There is a word that gets used as a slur in wellness circles and as a technical term in ours, and the gap between those two uses is where a great deal of bad nutrition advice lives. *Antinutrient.* It sounds like a villain, and it is nothing of the kind. Phytic acid is how a seed stores the phosphorus it will need to germinate. Oxalate is how many plants manage calcium and defend their tissues. Tannins are part of how a fruit protects itself as it ripens. These are not toxins to be purged; they are the plant's own good housekeeping, and they happen, as a side effect, to grip certain minerals so tightly that a human gut cannot pry them loose. To design nourishment honestly is to understand these molecules as what they are — not evil, not optional, but *tunable* — and to know exactly how far they can be turned down before the plant that depends on them suffers for it.



A soybean is a masterclass in the antinutrient trade: extraordinary mineral and protein content, held partly hostage by the phytate that the seed packed for its own future.

The most important of them is **phytic acid** — *myo*-inositol hexakis-phosphate, IP6 — a small molecule ringed with six phosphate groups, each a hook for a positive ion. At the pH of the small intestine those hooks close hard around iron, zinc, calcium, and magnesium, forming insoluble complexes that pass straight through. Phytate is the single largest reason that the iron in a bowl of whole grains and legumes — of which there is a great deal — is absorbed so sparingly. And here nutrition science hands us a gift the rest of

the book will lean on: the inhibition is **stoichiometric and predictable**, which means it can be written as a ratio and designed against.

The governing numbers are molar ratios, and they are worth committing to memory because they turn a vague worry into an engineering target:

RATIO	GOOD ABSORPTION WHEN	STRONG ABSORPTION WHEN	WHAT IT GOVERNS
phytate : iron	< 1	< 0.4	non-heme iron availability
phytate : zinc	< 15	< 5	zinc availability
phytate : calcium	< 0.24	—	calcium availability
(phytate × calcium) : zinc	< 200	—	zinc, in high-calcium meals

Let us work one, because the ratio is only useful if you can compute it. Take a legume holding 10 mg/g of phytic acid (molar mass ≈ 660 g/mol, so ≈ 15.2 µmol/g) and 50 µg/g of iron (molar mass ≈ 55.8 g/mol, so ≈ 0.90 µmol/g). The phytate:iron molar ratio is 15.2 / 0.90 ≈ **17** — a ratio at which almost nothing comes loose; the iron is there and unreachable. Our job is not to add iron to a food already rich in it; it is to drag that ratio under one, and ideally under four-tenths. Halving the phytate to 7.6 µmol/g and lifting iron 50% to 1.35 µmol/g brings the ratio to ≈ 5.6 — better, but still far from target, which tells us immediately that phytate reduction alone is not enough and the counter-reagent will have to finish the job. That is the kind of thing the arithmetic tells you before you waste a season learning it in the field.

There are three honest ways to move the ratio, and we use them in combination:

- **Express a phytase.** Phytases cleave phosphate off the IP6 ring, dismantling the hooks. Seeds make their own, activated during germination and soaking — which is the biochemistry behind every traditional culture that soaks, sprouts, or ferments its grains and pulses before cooking. We can express a heat-stable phytase in the seed itself, so the dismantling happens during digestion rather than requiring an overnight soak. Thermostability is the whole game here, because a phytase that denatures in the pot or the gut does nothing; it is exactly the kind of fold-stability problem the platform's protein-design pipeline was built to solve, and Chapter 14 shows the ΔΔG table for doing it.
- **Breed down the synthesis.** *Low-phytic-acid* (lpa) variants, reached by knocking down the inositol-phosphate kinases (the *IPK* steps) or the IP6 transporter (an ABC transporter, *MRP/ABCC*) that fills the storage vacuole, lower the phytate load at the source. This is the most powerful lever and the most delicate, because the seed packed that phosphorus for a reason: push lpa too far and germination vigor, seedling establishment, and field yield can all suffer. The flux-balance model earns its keep here, showing

where the phosphorus the seed can no longer store goes instead, and whether the plant can spare it. We tune Ipa, we do not abolish it.

- **Add the counter-reagent.** Ascorbate — vitamin C — reduces ferric iron to the ferrous form and holds it in a soluble chelate that resists phytate's grip, and it can lift non-heme iron absorption severalfold. Designing ascorbate *up* in the same tissue where we are driving phytate *down* is a two-lever play on a single ratio, and it is the clearest example in the book of why BD is a product and not a sum. We will build the ascorbate pathway in Chapter 6 specifically so that we can spend it here.

The other binders follow the same logic with different chemistry. **Oxalate**, in spinach and beet greens and chard, forms insoluble calcium oxalate — which is why spinach, for all its calcium content, surrenders perhaps a twentieth of it, against roughly a third for low-oxalate greens like kale. The lever is to breed the oxalate down, and the caution is that oxalate is doing real work for the plant, so the floor is not zero. **Tannins and other polyphenols** bind non-heme iron in tea, coffee, and some pulses; the same ascorbate that defeats phytate helps here too, and the designer's choice is whether a given polyphenol is a bug (an iron inhibitor) or a feature (an astringency the flavor index prizes, and a pigment that protects the plant) — a judgment Chapter 12 takes up directly. **Protease inhibitors and lectins** in raw legumes blunt protein digestion and are, mercifully, heat-labile: ordinary cooking denatures the Kunitz and Bowman-Birk inhibitors and the lectins alike, which is why the antinutrient that matters for protein is mostly a solved problem the moment a pot is involved.



Soaking, sprouting, souring — the oldest kitchens were practicing phytate reduction for a thousand years before anyone could spell it. Design is the same move, moved upstream into the seed.

We end the chapter where we began it, on the word itself. An antinutrient is not a thing to fear or to purge; it is a parameter. The reason this matters beyond semantics is that the fearful framing leads people to strip foods of things they need — the phosphorus, the fiber, the protective polyphenols — in pursuit of a mineral they could have freed with a soak and a squeeze of lemon. We would rather do the honest, boring, stoichiometric thing: measure the ratio, lower the binder by exactly as much as the plant can spare, add the

counter-reagent in the same tissue, and read the result as the only number that was ever going to tell the truth — the bioaccessible fraction, measured in a tube that treats every sample alike.



PART II OF VI

The Vitamins

Provitamin A · folate · the tocopherols · ascorbate — the pathways, the equivalents, the matrix

Chapter 4 · Provitamin A — *the carotenoid, the chromoplast, and the twelve-to-one tax.*

No nutrient has been designed into a plant more famously, or more consequentially, than provitamin A, and no design better illustrates the two-front war this book keeps describing: you must build the molecule, and then you must make it reachable, and the second fight is the one that is usually lost. β-carotene is the plant's own sunscreen and light-harvesting pigment, an orange forty-carbon chain the body can cleave into two molecules of retinal and thence retinol — true vitamin A. It is abundant in the plant world and scarce in the staple grains that feed most of it, which is why engineering it into an endosperm became the defining project of nutritional biotechnology. We build on that lineage directly, and we build on its hard-won humility about the difference between making a pigment and delivering a vitamin.



Rice is the canonical frontier: a grain whose polished endosperm holds essentially no provitamin A, and into which a whole carotenoid pathway can be written, enzyme by enzyme.

The pathway is a relay, and naming its runners is the point, because each is a place a designer intervenes:

STEP	ENZYME	SUBSTRATE → PRODUCT	WHY IT MATTERS TO DESIGN
0	GGPS	IPP + DMAPP → GGPP	the C20 precursor; shared with many pathways
1	**PSY** (phytoene synthase)	2 × GGPP → phytoene	the rate-limiting gate; the committed first step
2	PDS (phytoene desaturase)	phytoene → ζ-carotene	desaturation; colorless-to-colored transition
3	ZDS (ζ-carotene desaturase)	ζ-carotene → lycopene	completes the chromophore
4	CRTISO (carotenoid isomerase)	*cis* → all-*trans*-lycopene	isomerization to the usable form
5	**LCYB** (lycopene β-cyclase)	lycopene → β-carotene	cyclizes both ends; makes the provitamin

In a tissue that makes no carotene at all — a polished rice endosperm — the elegant engineering insight was that you do not need to install all five plant enzymes. You install **PSY** to open the committed gate, and then a single bacterial **CRTI** (a carotene desaturase from *Pantoea*) does the combined work of PDS, ZDS, and CRTISO, carrying phytoene all the way to lycopene by itself; the grain's own residual LCYB finishes the job to β-carotene. Two transgenes, a whole pathway. That is the architecture of the original golden grain, and in its improved form — PSY drawn from maize rather than daffodil — carotenoid accumulation in the endosperm reaches on the order of **30–37 µg/g**, a step-change from the undetectable baseline, roughly a twenty-three-fold improvement over the first-generation construct. We use exactly this logic, and we extend it two ways. By choosing the **LCYB dose** we set the lycopene-to-β-carotene balance — important because lycopene is a fine antioxidant pigment but *not* a provitamin, so a design heavy on lycopene photographs red and delivers little vitamin A. And by co-expressing a **sequestration partner** (a fibrillin-type or Orange-protein strategy that biogenesis the chromoplast toward a globular store) we push the carotene into a form that micellizes rather than crystallizes — fighting the matrix battle of Chapter 2 in the same construct that wins the synthesis battle.

The regulation matters as much as the enzymes, and it is where flux-balance thinking enters. PSY competes for GGPP with chlorophyll, tocopherol, and gibberellin synthesis; push PSY hard in the wrong tissue and you can distort the plant's development. The best designs are tissue-specific — endosperm-only, fruit-only — and they lean on the plant's own chromoplast-biogenesis regulators (the *Orange*/OR gene, which both triggers chromoplast formation and stabilizes the carotenoid sink) so that the carotene has somewhere to *go* once it is made. A pathway with no sink back-pressures and shuts itself down; a pathway with a sink keeps running. This is a general lesson the whole book will repeat: you do not just build a molecule, you build the place it lives.

Now the arithmetic of honesty. Provitamin A is not vitamin A; the body pays a conversion tax, and the accounting unit that respects that tax is the **Retinol Activity Equivalent (RAE)**. Dietary β-carotene converts at **12:1** by mass to retinol (other provitamin A carotenoids, like α-carotene and β-cryptoxanthin, at 24:1); only purified supplemental β-carotene in oil gets the better 2:1 rate. So a grain

holding 30 µg/g of β-carotene is not holding 30 µg/g of vitamin A — it is holding 2.5 µg RAE/g *before* the matrix takes its cut. This is the carotenoid's double tax: the conversion ratio and the bioaccessible fraction, stacked. Our BD framework absorbs both cleanly. We report carotene concentration as *C*, we measure *f_access* on the cooked grain (cooking matters enormously here — it ruptures cells and liquefies the store), and we convert the released, bioaccessible carotene to RAE at 12:1 only at the very end. The number that comes out the far side is the only one worth printing: bioaccessible RAE per serving, cooked, with the conversion tax already paid.



Carrot and corn and rice, read as provitamin A problems, are three different matrices around the same molecule — crystalline, protein-bound, and absent-then-installed, respectively.

LINE (ILLUSTRATIVE DESIGN TARGETS, COOKED)	B-CAROTENE C (MG/G)	F_ACCESS	BIOACCESSIBLE RAE (MG/G)	NOTE
reference polished grain	~0	—	~0	baseline; N undefined (division by ~0)
PSY + CRTI endosperm	30	0.15	0.38	synthesis won, matrix modest
PSY + CRTI + sequestration partner	32	0.28	0.75	matrix co-designed; ~2× delivered

We print those as targets, not trophies, and we flag the honest awkwardness in the first row: when the reference holds essentially zero, the *N* score runs to infinity and stops being useful, so for provitamin A in a staple grain we report **absolute bioaccessible RAE** instead of a ratio. A good measure knows the edges of its own domain. What the table does show is the thesis in miniature: the jump from the second row to the third is not more pigment — it is barely more pigment, a 7% bump in content — it is the *same* pigment made nearly twice as reachable by addressing the chromoplast it lands in. The RAE per gram roughly doubles on a content change of almost nothing, because *f_access* did the work. Synthesis is the famous half of provitamin A engineering. The matrix is the half that decides whether the famous half mattered, and it is the half we will not let anyone skip.

Chapter 5 · Folate — *two branches, one vitamin, and the fragility of a gift.*

Folate is the vitamin that taught the world a plant could be a pharmacy and a trap at once: essential, water-soluble, abundant in green leaves and legumes, and so chemically fragile that it bleeds out of a vegetable into the cooking water and oxidizes in the light on the counter. Its name is its origin — *folium*, leaf — and its job in the body is among the most fundamental there is, one-carbon transfer, the chemistry of building and copying the molecules of inheritance. We design folate into fruit with a particular care, because here more than anywhere the density you build is only as good as the stability you protect, and a gift that does not survive to the plate was never really given.



The tomato is our folate workbench: a fruit eaten fresh and cooked, with a generous cytosolic matrix, into which two short biosynthetic branches can lift folate roughly twenty-five-fold.

The biosynthesis of folate is a convergence — two separate branches that meet to form the finished vitamin — and the design lesson is that you must feed *both* or you starve the junction:

BRANCH	COMMITTED ENZYME	STARTING MATERIAL	PRODUCT TOWARD FOLATE
Pteridine	**GTP cyclohydrolase I (GTPCHI)**	GTP	6-hydroxymethyl-dihydropterin
PABA	**aminodeoxychorismate synthase (ADCS)**	chorismate	*para*-aminobenzoate (PABA)
Junction	dihydropteroate synthase; DHFR	pterin + PABA + glutamate	dihydro- then tetrahydrofolate
Retention	**folylpolyglutamate synthetase (FPGS)**	THF	polyglutamylated, cell-retained folate

The elegant and slightly counterintuitive finding that organizes folate engineering is that overexpressing **GTPCHI alone** boosts the pteridine half but can leave the plant *PABA-limited* — you build half a vitamin and run out of the other half, and the pterin backs up as an unusable intermediate. Feeding **both** GTPCHI and ADCS unlocks the real gains, and in the tomato this combination lifts folate roughly **25-fold** over the already-respectable baseline. That is the density win, and it is substantial. But folate design has a second front that carotene does not: **stability and retention**, and this is where the glutamate tail earns its keep. Polyglutamylated folates are bound and held inside the cell and are markedly more stable to the oxidation and leaching that destroy the monoglutamyl form. So a folate program that stops at synthesis — that builds the 25-fold density and ignores the tail — can watch a large fraction of it degrade between field and fork. Driving **FPGS** alongside the two synthetic branches protects the gift, and we read the protection directly.

Here is the retention arithmetic made concrete, because "fragile" is not a number and *C* at the plate is. Suppose the engineered fruit holds 375 µg/100 g at harvest. Monoglutamyl folate might lose 40% to a storage-plus-cooking sequence (375 → 225); a polyglutamylated, FPGS-stabilized folate might lose only 15% over the same handling (375 → 319). The retention tax is the difference between those two endpoints, and it is often larger than the entire gain from a lesser synthesis boost. This is why we measure folate not only at harvest but after realistic storage and after cooking, and why the honest *C* for a fragile vitamin is always the one left when the food is actually eaten — not the triumphant number on the day it was picked.

FOLATE FORM	HARVEST C (MG/100 G)	AFTER STORAGE + COOKING	RETAINED FRACTION
reference tomato (native)	15	9	0.60
engineered, monoglutamyl	375	225	0.60
engineered + FPGS (polyglutamyl)	375	319	0.85

The accounting unit here is the **Dietary Folate Equivalent (DFE)**, which exists precisely because natural food folate and synthetic folic acid are absorbed differently: 1 µg of food folate counts as 1 DFE, while 1 µg of synthetic folic acid counts as 1.7 DFE because of its higher bioavailability. For our purposes the point is that *we are building food folate, not folic acid* — the naturally polyglutamylated, matrix-embedded form — so our numbers live on the 1:1 side of that conversion, and our *f_access* assay, which deconjugates the

polyglutamyl tail the way the gut's own brush-border conjugase does, is measuring the right thing. The nourishment score for our folate tomato against a reference tomato runs to the order of **N ≈ 25,000** on density at harvest, settling toward **N ≈ 21,000** once the retention tax is paid and the FPGS protection is counted — a number we report with the same two cautions as always: it describes the fruit, and the fruit alone.

A word on what "folate" even names, because the vitamin is a family and the design has to respect the forms. The cofactor cycles through several one-carbon-loaded states — **5-methyl-THF** (the dominant circulating form), **5,10-methylene-THF** (the form that builds the thymine of DNA), and **10-formyl-THF** (for purine synthesis) — interconverting as the cell's one-carbon economy demands, with **5-formyl-THF** serving as a stable storage form. A plant accumulates a mixture of these, and the mixture's stability differs form to form, which is a second reason the polyglutamyl tail and the matrix matter so much: they protect a labile *set* of molecules, not one. We built the tomato first, but the same two-branch logic transfers directly to the other generous-matrix hosts — a lettuce whose folate is eaten raw, a legume whose folate rides with its protein, a rice whose endosperm starts near zero and so, like the provitamin-A grain, is scored in absolute DFE rather than as a ratio. The vitamin is a family; the design feeds the whole family and protects it on the way to the plate.

There is a reason we chose a fruit, and not a leaf, as the folate workbench, and it is a matrix reason. A leaf's folate is defended by cell walls and lost to boiling water; a tomato is eaten fresh as often as cooked, its folate sits in a generous cytosol, and its *f_access* stays high across both preparations. The tomato is also our best-characterized non-pepper host — 2,040 constructs deep, rated *Easy* for engineering, primary on sweetness and umami — so the folate work rides on a chassis we already know down to its promoters. When we go looking for where to put a fragile water-soluble vitamin, we look for the kindest matrix and the most forgiving kitchen *and* the best-understood genome, because the best-built molecule in the world is worth nothing if the cooking pours it down the drain or the host fights the construct. Folate, more than any nutrient in this book, is a lesson in humility: you can engineer a twenty-five-fold gift and lose most of it to light and heat if you are not as careful about keeping the vitamin as you were about making it.

Chapter 6 · Vitamin E and vitamin C — *the one that guards the fat, and the one that frees the iron.*

Two vitamins share this chapter because they are the plant's own antioxidants, and because designing them teaches opposite halves of the same craft. Vitamin E — the tocopherols — is lipophilic, lives in membranes and oils, and protects the plant's fats from going rancid; its design problem is one of *isoform*, of making the right member of a family. Vitamin C — ascorbate — is water-soluble, floods the cytosol, and is the most industrious small molecule in plant metabolism; its design problem is one of *flux*, of opening a long pathway's bottleneck. Together they are also the two vitamins that do double duty in this book — vitamin E stabilizing the designed fats of Chapter 10, vitamin C freeing the designed iron of Chapter 7 — so building them well here is spending we collect twice later.



Leafy greens hold both — tocopherols in the chloroplast membranes, ascorbate in the cytosol — and both are more fragile than their reputations suggest.

Vitamin E is not one molecule but eight — four tocopherols and four tocotrienols — and they are emphatically not equal. α-tocopherol has by far the highest vitamin E activity; γ-tocopherol, which dominates in many seed oils, has perhaps a tenth of it. So a plant can be rich in "vitamin E" by the crude total and poor in the form that counts. The pathway builds the ring from homogentisate (made by **HPPD** from the tyrosine pool), attaches the lipophilic prenyl tail via **HPT / VTE2**, and then — crucially — methylates its way around the ring. The last methylation, **γ-tocopherol methyltransferase (γ-TMT /**

VTE4), converts γ-tocopherol to α-tocopherol. Overexpressing VTE4 therefore does something quietly profound: without raising the total tocopherol at all, it *shifts the mixture* from the low-activity γ-form to the high-activity α-form, and the vitamin E activity of the oil can rise many-fold. Here is the arithmetic that makes the point: an oil at 100 µg/g total tocopherol that is 90% γ and 10% α carries, in α-tocopherol-equivalents, (0.10 × 100 × 1.0) + (0.90 × 100 × 0.1) = 10 + 9 = 19 α-TE/g. Flip the ratio with VTE4 to 90% α and 10% γ and the *same 100 µg/g total* now carries (0.90 × 100 × 1.0) + (0.10 × 100 × 0.1) = 90 + 1 = 91 α-TE/g — nearly a fivefold gain in delivered vitamin E from a methyltransferase, with no change in total content at all. This is the purest example in the book of why *C* alone is a liar. We design the isoform, not the total, and we report α-tocopherol equivalents, not crude tocopherol.

Vitamin C is a different craft entirely. Plants make ascorbate through the Smirnoff–Wheeler pathway, a long relay off the sugar-nucleotide GDP-D-mannose:

STEP	ENZYME	PRODUCT	ROLE
1	GME (GDP-mannose 3,5-epimerase)	GDP-L-galactose	commits sugar toward ascorbate
2	**GGP / VTC2** (GDP-L-galactose phosphorylase)	L-galactose-1-P	**the major control point**
3	GPP	L-galactose	—
4	L-GalDH	L-galactono-1,4-lactone	penultimate
5	**GLDH** (L-galactono-1,4-lactone dehydrogenase)	L-ascorbate	final oxidation, in the mitochondrion

The pathway's bottleneck is **GGP/VTC2**, and it is a bottleneck held shut by an unusually clever brake: the VTC2 messenger carries an **upstream open reading frame** (uORF) that represses its own translation when ascorbate is already high — a feedback valve written into the RNA itself. Relieving that brake, by editing the uORF or driving VTC2 past its own repression, opens the flux gate and can multiply fruit ascorbate severalfold. That is the density lever. But ascorbate's real value to a *nourishment* program is relational, and we have been promising it for two chapters: ascorbate reduces ferric iron (Fe^{3+}) to the soluble ferrous form (Fe^{2+}), roughly 2 moles of ascorbate stabilizing 1 mole of iron against re-oxidation and re-binding, and in that reduced, chelated state the iron resists phytate and tannin. Designing vitamin C up in the *same tissue* where we are driving phytate down (Chapter 3) and iron up (Chapter 7) is not three separate projects; it is one co-design on one matrix, and the iron's bioaccessible fraction is the shared number that reports whether it worked.



Sweet potato and the deep-orange roots carry the fat-soluble story; the greens carry the water-soluble one. A whole plate is how the two vitamins were always meant to arrive.

There is a reason the plant can spend these antioxidants so freely, and it is a recycling economy worth naming because it is itself a design surface. Oxidized ascorbate is not discarded; it is regenerated through the **ascorbate–glutathione (Foyer–Halliwell) cycle**, in which the enzymes **MDHAR** and **DHAR** reduce the spent forms back to ascorbate at the cost of glutathione and NADPH. And ascorbate in turn regenerates spent vitamin E: when α-tocopherol quenches a lipid radical and becomes a tocopheroxyl radical, ascorbate at the membrane surface reduces it back to the active vitamin. So the two antioxidants of this chapter are not independent after all — they are a coupled system, the water-soluble one refreshing the fat-soluble one, and a seed engineered for durable vitamin E is quietly also a seed that wants enough ascorbate and glutathione to keep that vitamin E in its active state. Boosting the recycling enzymes can raise the *effective* pool of both vitamins without making a single additional molecule of either — another case, like the VTE4 isoform shift, where the cleverest lever is not synthesis at all.

The caution for both vitamins is their fragility, and it is the same caution folate earned in the chapter before. Ascorbate oxidizes readily — in air, in heat, in storage — so its honest *C* is a moving target that falls from the moment of harvest, and we measure it along that decline, not only at its peak; a tomato's ascorbate can fall by a quarter in a week of warm storage. Tocopherols are sturdier but not immortal, and they are spent defending the very oils they ride in, so a seed engineered for more α-tocopherol is partly spending that vitamin to protect its own designed fats — a transaction we track explicitly in Chapter 10. The honest way to design an antioxidant vitamin is to remember that it is an antioxidant: it is *consumed* doing its job, and the number that matters is the one still standing when the food is eaten. Build generously, protect the build, and measure at the plate — the same discipline, applied to the two molecules the plant invented to protect itself.



PART III OF VI

The Minerals

Iron · zinc · calcium · selenium — storage, chelation, and the molar-ratio problem made solvable

Chapter 7 · Iron — *the mineral the plant hides from itself.*

Iron is the mineral that most rewards the full discipline of this book, because with iron, density is almost never the problem. The plant world is not short of iron; a bowl of lentils, a plate of greens, a wedge of whole-grain bread all carry it in quantity. The problem is entirely one of *release* — iron is held behind phytate, bound by tannins, and locked in chemical forms the gut cannot reduce — and so iron is the nutrient where our two-number framework stops being an elegance and becomes a necessity. You cannot design iron nourishment by measuring concentration. You have to measure what comes loose, and then you have to measure whether a living cell will actually take up what came loose, because for iron even bioaccessibility is not the final word.



The legume is iron's great paradox: dense with the mineral, and dense with the phytate that holds it. Design has to work both sides of that ledger at once.

There is a biology of two irons, and a designer must hold both in mind. **Heme iron**, bound inside the porphyrin ring of hemoglobin and myoglobin, is absorbed efficiently — on the order of 15 to 35 percent — and comes almost entirely from animal foods. **Non-heme iron**, the ionic form in plants, is absorbed far more variably and usually far more poorly — 2 to 20 percent — because its fate depends entirely on what else is in the meal. Everything a plant-forward nourishment program does about iron is about dragging that non-heme number up the range it already occupies, and the levers are, satisfyingly, the ones we have already built:

- **Ferritin.** Ferritin is the universal iron-storage protein — a hollow 24-subunit protein shell that holds up to roughly 4,500 iron atoms in a mineral core — and ferritin-bound iron is absorbed in a bioavailable form, apparently taken up partly intact. Expressing a seed ferritin (soybean or common bean *FER*) in a grain's endosperm raises both the iron content and the *quality* of its storage. High-iron rice lines built on ferritin plus the chelator below have reached on the order of **15 µg/g** of iron, several-fold over a polished baseline near 2 µg/g.
- **Nicotianamine synthase (NAS).** Nicotianamine is the plant's own metal chelator, central to loading iron (and zinc) into seeds and moving it through the phloem. Driving NAS raises seed iron *and* improves its chemical form for uptake; nicotianamine and its derivative 2′-deoxymugineic acid keep the metal soluble and mobile. Ferritin stores it; NAS helps put it there and keeps it reachable.
- **Phytase, and the phytate:iron ratio.** This is the release lever from Chapter 3, and it is the one that most often decides the outcome. All the ferritin in the world avails little if the iron stays above a phytate:iron molar ratio of one; co-expressing a heat-stable phytase in the seed drives that ratio down during digestion and converts stored iron into *available* iron. Density and availability, stacked in a single seed.
- **Ascorbate, the counter-reagent.** From Chapter 6: vitamin C reduces Fe^{3+} to the soluble Fe^{2+} and shields it from re-binding, lifting non-heme absorption severalfold. The classic, almost folkloric observation — a squeeze of citrus over the beans, the tomato cooked with the lentils — is this reaction, and we design it into the tissue so it does not depend on the cook remembering.

For iron, and only iron, we climb one rung higher on the measurement ladder than *f_access*, because solubility is not the same as uptake. We run the digest onto a monolayer of **Caco-2** cells — a human intestinal cell line that differentiates into something that behaves like gut lining — and we measure the cells' own **ferritin formation** in response, which tracks how much iron they actually imported. It is still a model, still a tube and a dish rather than a person, and we say so every time. But it is a materially better proxy than solubility alone, because it asks the question solubility cannot: not *did the iron come loose* but *would a gut cell take it in*. A design that wins the solubility assay and loses the Caco-2 assay has told us something true and sobering, and we would rather learn it on a bench than assume it on a label.

Here is the full stacking, worked, so the thesis is undeniable. Take a reference bean: iron 50 µg/g, *f_access* 0.05 (phytate-locked), so BD = 2.5 µg/g. Now a designed line: ferritin and NAS lift iron to 75 µg/g; phytase drops the phytate:iron ratio under one and lifts *f_access* to 0.15; ascorbate co-expression lifts it further to 0.20. BD = 75 × 0.20 = 15 µg/g. The nourishment score is N = 1000 × 15 / 2.5 = **6,000** — a sixfold delivered-iron gain built from a 1.5× content change and a 4× availability change, multiplied. The content lever alone would have

given N ≈ 1,500; the whole gain above that came from release. No label that reports content could ever see the four-fifths of this result that lives in the matrix.

LINE	FE CONTENT C (MG/G)	F_ACCESS	CACO-2 FERRITIN (RELATIVE)	BD (MG/G)	N VS REFERENCE
reference bean	50	0.05	1.0	2.5	1,000 (anchor)
ferritin + NAS only	75	0.06	1.3	4.5	1,800
+ phytase (phytate ↓)	75	0.15	3.0	11.3	4,500
+ ascorbate co-design	75	0.20	3.8	15.0	6,000



Greens carry iron too — and the oxalate that complicates it. The whole-plant view is the only honest one: content, binder, and counter-reagent together.

We will say the discipline once more, plainly, because iron is where the temptation to overclaim is strongest and the history of iron fortification is littered with programs that measured content, declared success, and changed nothing in the field: we design a food that holds more iron in a less-bound form from which a cultured intestinal cell takes up more. That is the entire, honest sentence. It is a real achievement and a measurable one, and it stops exactly at the gut wall of a cell in a dish. What iron status any individual reaches is theirs — a function of their physiology, their whole diet, their needs, the regulation of their own absorption, which rises when they need more and falls when they do not — and nothing we grow is entitled to speak for it. The plant can be built to offer more. Whether a body takes the offer is not ours to claim, only to make possible.

Chapter 8 · Zinc and the rest — *the quieter minerals, and the ratios that rule them.*

Iron gets the headlines; zinc does the governing. Zinc sits at the catalytic heart of hundreds of the body's own enzymes and of the plant's — and of the transcription factors, the zinc-finger proteins, that regulate much of life's gene expression — and like iron it is usually present in a plant-forward diet in adequate *quantity* and inadequate *availability*, held behind the same phytate by the same chemistry. The good news for a designer is that almost everything Chapter 7 built for iron transfers to zinc with the ratio adjusted, and the chapter's real work is to round out the mineral palette — calcium with its oxalate complication, magnesium and potassium riding along with good agronomy, and selenium, the mineral that reminds us the soil has a vote.



Cereal grains are the zinc frontier the way they were the iron and provitamin-A frontier: enormous dietary reach, a phytate-bound mineral load, and an endosperm that can be redesigned around both.

Zinc answers to the same design as iron — NAS to chelate and load, a modest storage and transport push via the **ZIP-family** transporters that move zinc into the seed, and above all phytate reduction — but its governing ratio is looser and its target different. Zinc absorption holds up reasonably until the **phytate:zinc molar ratio** climbs past about 15, and it climbs toward excellence below 5. In high-calcium meals a three-body term matters — **(phytate × calcium):zinc** below roughly 200 — because calcium and phytate gang up on zinc

together. The design brief writes itself from those numbers: lift seed zinc with NAS and transport, and drive the phytate down until the ratio clears 5, and read the result in the same in-vitro digestion we use for everything, escalating to a Caco-2 uptake assay where the stakes justify it. Zinc is, in a sense, the easiest win in the book — the same phytase that freed the iron frees the zinc in the same bite, two minerals unlocked by one enzyme, which is why a staple-grain program almost always designs iron and zinc together rather than separately.

Calcium is the oxalate story from Chapter 3, told as a mineral. Plants hold calcium well; some of them — spinach, chard, beet greens, rhubarb — lock most of it behind oxalate, so that spinach surrenders perhaps 5 percent of its calcium against roughly 30 percent for a low-oxalate green like kale or bok choy. The lever is to breed oxalate down, bounded by the plant's genuine need for it (oxalate buffers calcium and deters herbivores, so the floor is real), and the smarter move at the scale of a meal is often to *choose the matrix* — to lean the calcium contribution on the greens that were already generous with it rather than fight a plant determined to keep its oxalate. Here nourishment design and menu design are the same act, which is a theme Chapter 15 makes explicit.



The beet is a mineral workbench and a pigment workbench at once — its betalains the subject of Chapter 12, its mineral load governed by the same oxalate chemistry that governs its leaves.

Magnesium and potassium are the minerals good agronomy mostly delivers — they ride in with balanced soil and whole-tissue optimization rather than requiring a biosynthetic intervention, and the *Grown Right* volume takes up their cultivation in depth. We note them here to be complete and to make a point about proportion: not every nutrient needs an engineered pathway, and a program that reaches for CRISPR when it should reach for the soil has lost the plot. **Selenium** makes the opposite point most sharply. Selenium is not made by the plant; it is *taken up* from the soil, and a plant grown in selenium-poor ground will be selenium-poor no matter what its genome says. Where it is taken up, the design move is to steer it toward the better-tolerated, better-studied organic forms — **selenomethionine** and **methyl-selenocysteine** — rather than the

inorganic selenate, which is a matter of expressing the right assimilation and methylation enzymes (and, for safety, of keeping free selenocysteine from mis-incorporating into proteins). But the governing fact is agronomic: selenium is the mineral that reminds us the plant is only ever half the system, and the ground it stands in has the other half of the vote.

MINERAL	PRIMARY LEVER	GOVERNING NUMBER	CAVEAT
Iron	ferritin + NAS + phytase + ascorbate	phytate:Fe < 1 (ideal < 0.4)	escalate to Caco-2; non-heme is variable
Zinc	NAS + ZIP transport + phytase	phytate:Zn < 15 (ideal < 5)	(phytate×Ca):Zn < 200 in dairy-rich meals
Calcium	low-oxalate; matrix choice	oxalate load; absorbable fraction	floor on oxalate reduction; plant needs some
Mg / K	agronomy, whole-tissue	balanced uptake	rarely needs a pathway; soil does the work
Selenium	organic-form assimilation	soil Se availability	not synthesized; soil has the deciding vote

The quieter minerals teach the book's quietest lesson, which is about proportion and honesty together. The phytate:zinc ratio, the oxalate on the calcium, the selenium in the soil — these are not dramatic, and they do not make press, and they are where most of the real nourishment of a real diet is actually won or lost. We measure them with the same instruments and the same restraint as the headline nutrients, and we make the same refusal with every one: we are describing a food's content and a food's willingness to give that content up, and we are not, at any point, describing what will happen inside a person. The minerals are freed or they are bound. The body is its own.



PART IV OF VI

Protein, Fat & Fiber

The limiting amino acid · the long-chain omega-3 · the matrix of fiber — macronutrient quality by design

Chapter 9 · Protein quality — *the one amino acid the whole protein waits on.*

A protein is a chain, and a diet's protein is only as complete as its scarcest link. This is the oldest quantitative idea in nutrition — the limiting amino acid — and it is still the one that most cleanly separates a food that *contains* protein from a food that *delivers* it. Cereals are rich in protein and poor in **lysine**; legumes are rich in protein and comparatively short on the sulfur amino acids, **methionine** and cysteine. Each is a half-built wall waiting for the other's bricks, which is why nearly every traditional cuisine on earth independently discovered the grain-and-pulse pairing — rice and dal, corn and beans, bread and hummus — long before anyone could name a single amino acid. Designing protein quality is, at bottom, the science of filling in the limiting link so the wall can stand on its own.



The soybean is the protein workbench: dense, nearly complete, and limited mainly by its sulfur amino acids and its heat-labile protease inhibitors.

The lysine problem in cereals has a beautifully specific solution, and it lives at a single feedback-regulated enzyme. Lysine is made through the aspartate pathway: aspartate is phosphorylated by **aspartate kinase (AK)**, and the committed branch toward lysine turns on **dihydrodipicolinate synthase (DHDPS)** — and both AK and DHDPS are inhibited by lysine itself, a double feedback brake that keeps the amino acid from accumulating. The elegant move is to swap in a **feedback-insensitive DHDPS** — a bacterial *dapA* that does not feel lysine's inhibition — so the brake comes off and free lysine

climbs. This is not hypothetical chemistry; it is the mechanism behind high-lysine maize (the LY038 event used a maize-optimized bacterial *cordapA*), and it generalizes to the other cereals whose walls wait on the same brick. The parallel move in legumes targets the sulfur amino acids, by steering sulfur assimilation and expressing methionine-rich storage proteins, with the same goal: fill the limiting link, raise the quality of the whole.

The feedback-insensitivity itself is a protein-design problem — you are looking for substitutions in the DHDPS allosteric site that abolish lysine binding without wrecking the catalytic site — and it is exactly the kind of problem the platform's protein-language-model-plus-flux-balance pipeline was built for, the same pipeline that wrote the two Pun1 variants. The output looks like this, a ranked set of substitutions scored for the two things that matter, insensitivity and fold stability:

EDIT (ILLUSTRATIVE, DHDPS ALLOSTERIC SITE)	EFFECT	ΔΔG (KCAL/MOL)
N80I	lysine no longer binds allosteric cleft	−0.4
H53Y	disrupts the inhibitory dimer interface	−1.1
E84T	abolishes feedback without touching active site	−0.7

Two complications make the lysine design more than a single swap, and both are real design surfaces. The first is **catabolism**: plants degrade their own lysine through a bifunctional enzyme, **lysine-ketoglutarate reductase / saccharopine dehydrogenase (LKR/SDH)**, so a line that makes more lysine can quietly burn it back down unless the catabolic enzyme is also knocked back — the best high-lysine designs raise synthesis *and* lower degradation, pushing on both ends of the pool at once. The second is **branch competition**: aspartate feeds not only lysine but threonine and methionine, through shared intermediates and their own feedback loops, so pushing carbon toward lysine can starve the others, and the flux-balance model is what keeps the whole aspartate family in balance rather than trading one limiting amino acid for another. There is also the matter of *where* the amino acid sits: the classic high-lysine maize mutants (**opaque-2**, floury) softened the grain and raised lysine by changing the storage-protein balance, at a cost to texture and kernel hardness that took decades of modifier breeding to repair — a reminder that protein quality, like everything in this book, is read on more than one axis at once.

But raising an amino acid's concentration is, once again, only the density half of the sentence, and protein has its own version of the availability tax — **digestibility**. The field has moved, rightly, from the older **PDCAAS** (protein digestibility-corrected amino acid score, which caps at 1.0 and uses fecal digestibility) to the more exacting **DIAAS** (digestible indispensable amino acid score), which measures *ileal* digestibility — what is actually absorbed in the small intestine, amino acid by amino acid — and does not truncate at 1.0, so it can tell a genuinely excellent protein from a merely adequate one. The numbers tell the story the pairing always knew.

PROTEIN SOURCE	LIMITING AMINO ACID	DIAAS (APPROX.)	DESIGN LEVER
Wheat (isolated)	lysine	0.45	feedback-insensitive DHDPS → ↑ lysine
Rice	lysine	0.50	*dapA*; endosperm-targeted
Maize	lysine (and tryptophan)	0.42	high-lysine *cordapA*; opaque context
Soybean	methionine / cysteine	0.90–1.0	sulfur assimilation; Met-rich storage protein
Grain + pulse together	mutually covered	> either alone	complementation — the oldest design of all

Work the complementation, because it is the whole argument for the pairing and for designing the grain rather than abandoning it. A wheat protein limited at lysine (DIAAS ~0.45) eaten with a pulse whose lysine is generous climbs toward the pulse's score, because the pulse's lysine covers the wheat's gap; raise the wheat's own lysine with *dapA* and the wheat climbs toward completeness *before* it even meets the pulse, so a smaller portion of pulse finishes the job. The designed grain does not replace the ancient pairing; it makes the pairing more forgiving, which for a staple eaten three times a day is exactly the kind of margin that matters.

The availability tax has a second line item particular to legumes, and it is the happiest problem in the book because it is already solved in every kitchen: **protease inhibitors and lectins**. Raw soy and raw beans carry Kunitz and Bowman-Birk inhibitors that blunt the gut's protein-cleaving enzymes, and lectins that irritate — and both are **heat-labile**, denatured by ordinary cooking. The antinutrient that most threatens protein is undone by a pot of boiling water, which is why no one eats raw kidney beans and why the DIAAS of properly cooked legumes is the number that counts. We design the amino-acid composition, we let the cook do what cooks have always done to the inhibitors, and we measure the protein quality on the food as eaten — cooked, complemented, and scored by ileal digestibility, the only measure that respects the difference between a protein present and a protein absorbed.

Chapter 10 · The good fats — *building the long chain the plant never learned to finish.*

Fat is the macronutrient most maligned and least understood, and the nourishment program approaches it with a single, specific ambition that has nothing to do with quantity and everything to do with *structure*: the position and number of the double bonds in the chain. The distinction between a fat that the body treats as fuel and a fat it treats as a signaling and structural molecule is a matter of desaturation and elongation — of where the kinks are and how long the chain runs — and plants, which are superb at making the shorter omega-3 and omega-6 fats, almost universally stop short of the long-chain omega-3s that are otherwise the near-exclusive province of the sea. Teaching an oilseed to finish that chain is one of the genuine frontiers of nutritional design, and it is where this chapter is pointed.



The oilseed is where fat structure is designed — not the amount of oil, but the exact pattern of double bonds the plant writes into each chain.

Start with what the plant does well. From oleic acid (18:1), the desaturase **FAD2** installs the second double bond to make linoleic acid (18:2, omega-6), and **FAD3** installs the third to make **α-linolenic acid (ALA, 18:3, omega-3)** — the short-chain omega-3 that flax and canola and soy carry in abundance. Tuning the FAD2/FAD3 balance lets a designer set the omega-6 to omega-3 ratio of an oil, which is a meaningful compositional lever on its own: a modern diet skewed heavily toward omega-6 can be rebalanced at the source by biasing FAD3. But ALA is where plants stop, and the body converts ALA

onward to the long-chain omega-3s **EPA (20:5)** and **DHA (22:6)** only inefficiently — single-digit percentages for EPA, lower still for DHA. Those long-chain fats require a relay the plant simply does not possess:

STEP	ACTIVITY	REACTION	SOURCE OF THE GENE
1	Δ6-desaturase	ALA (18:3) → SDA (18:4)	microalgae / fungi
2	Δ6-elongase	SDA (18:4) → ETA (20:4)	microalgae
3	Δ5-desaturase	ETA (20:4) → **EPA (20:5)**	microalgae
4	Δ5-elongase	EPA (20:5) → DPA (22:5)	microalgae
5	Δ4-desaturase	DPA (22:5) → **DHA (22:6)**	microalgae

Assembling that relay from algal genes in a land oilseed — the work pioneered in *Camelina sativa* and canola — lets a field crop accumulate EPA and DHA in its seed oil, building on land the exact fats the ocean has had a monopoly on. It is a genuinely audacious piece of metabolic engineering: five to seven borrowed activities, expressed in concert, in the right order, in the seed, each enzyme handing its product cleanly to the next without leaking intermediates into the general lipid pool. The yield problem is real and quantitative — every step is less than perfectly efficient, so the fraction of oil reaching full-length DHA is the *product* of the step efficiencies, and a relay of six steps at 80% each delivers only about a quarter of its starting material to the end. Raising those per-step efficiencies, and suppressing the competing reactions that siphon intermediates away, is where the flux-balance model and the protein-design pipeline both earn their keep. We regard it as one of the most important frontiers on the whole nourishment roadmap, precisely because it does something agronomy and selection cannot — it does not optimize an existing trait, it *installs a capability the lineage never had*.

How far has the chain actually been carried? The best land-plant events — in *Camelina* and canola — have reached seed oils with **EPA plus DHA in the low tens of percent of total fatty acids**, a level that puts a field crop in the range of a fish oil, which a decade ago would have read as fantasy. Getting there meant solving a subtle bottleneck the table above hides: the **substrate dichotomy**. The borrowed algal desaturases prefer their fatty acid presented on one carrier (acyl-CoA), while the plant's own desaturation tends to happen on another (the phospholipid phosphatidylcholine), so intermediates pile up on the wrong carrier and the relay stalls. The fixes are themselves elegant pieces of design — acyl-CoA-dependent desaturases that accept the plant's native presentation, or an **LPCAT** shuttle that moves intermediates between the two carriers fast enough to keep the relay flowing. It is the kind of problem that only appears once you try to run an aquatic pathway on dry land, and solving it is most of what separates a trace of EPA from a genuinely useful oil.

FATTY ACID	SHORTHAND	MADE BY THE PLANT?	DESIGN STATUS
oleic	18:1	yes	baseline
linoleic (ω-6)	18:2	yes (FAD2)	rebalance against ω-3
α-linolenic (ω-3)	18:3	yes (FAD3)	the plant's own ceiling
EPA (ω-3)	20:5	no	installed (algal relay)
DHA (ω-3)	22:6	no	installed; the hardest step

The design does not end at synthesis, because a polyunsaturated fat is a fragile thing — every double bond is a site for oxidation, for rancidity, for the slow spoilage that turns a nourishing oil into a spent one, and a six-double-bond DHA molecule is among the most oxidation-prone compounds a seed can carry. This is where Chapter 6's vitamin E comes back to collect. The more unsaturated the oil we build, the more **α-tocopherol** it needs riding alongside to protect those double bonds from oxidation in the seed, in storage, and on the shelf. So the EPA/DHA oilseed is co-designed with its own antioxidant — VTE4 driving the tocopherol pool toward the protective α-form in the same seed that carries the fragile long chains — and the two traits are read together, by the fatty-acid profile (chain length and double-bond position, not just total fat) *and* the oxidative-stability number (peroxide value over storage) that decides whether the structure survives to the plate. An omega-3 that oxidizes before it is eaten was a density win and a nourishment loss, and the whole value of the fat was in the fragile part we then failed to protect. Build the chain the plant never finished, and then guard it, because the fat's entire point is the double bonds that make it both precious and perishable.

Chapter 11 · Fiber, resistant starch, and the whole matrix — *the nutrition that is mostly architecture.*

Fiber is the nutrient that is not absorbed, which is exactly its point and exactly why it is the hardest to write about with discipline. Everything else in this book is designed to come *loose* — to be freed from the matrix and delivered. Fiber is designed to stay. It is the plant's own architecture, the cell walls and storage polysaccharides that the human gut's enzymes cannot cleave, and its nourishment value lives in that very indigestibility. Here more than anywhere we hold the line on language, because fiber is the nutrient most draped in health halo, and we are going to describe what it *is* — a matter of chemistry and structure, measurable on a bench — and refuse, as always, to describe what it does to anyone.



A whole grain's value is mostly architecture — the bran, the cell walls, the starch granules — and fiber design is the design of that architecture, not of a molecule that comes loose.

The compositional vocabulary is where the rigor lives. Dietary fiber divides, usefully if imperfectly, into **soluble** and **insoluble**, and the distinction is physical and real: soluble fibers dissolve or disperse in water and form viscous solutions; insoluble fibers do not. Among the soluble fibers, the one most worth a designer's attention is **mixed-linkage (1,3;1,4)-β-D-glucan**, abundant in oats and barley, a polysaccharide whose viscosity is a direct, measurable function of its molecular weight and concentration — roughly, solution viscosity

scales with concentration and with molecular weight raised to a power greater than one, so that trimming the chains cuts viscosity steeply. That makes β-glucan *tunable*: the genes governing the cellulose-synthase-like **CslF/CslH** synthases set how much β-glucan accumulates, and the endogenous β-glucanases that trim chain length set the molecular weight, and therefore the viscosity, you end up with. Then there is **resistant starch** — starch that escapes digestion in the small intestine — which comes in a graded typology that is itself a design space:

TYPE	WHAT MAKES IT RESISTANT	LEVER
RS1	physically inaccessible (intact cell walls, whole/coarse grains)	grain structure, milling
RS2	native granule crystallinity (raw starch, high-amylose)	high-amylose design (↓ branching enzyme)
RS3	retrograded starch (formed on cooking then cooling)	amylose content + preparation
RS4	chemically modified	processing

The most powerful genetic lever there is **RS2 through high-amylose design**: starch is a blend of linear amylose and branched amylopectin, and shifting the balance toward amylose — by knocking down the starch-branching enzymes (*SBEIIa/SBEIIb*) that build amylopectin's branches — raises the fraction of starch that resists digestion, and raising amylose also raises the RS3 that forms when the cooked grain is cooled. The same move changes the grain's texture and cooking behavior, which is the recurring lesson of this book stated in starch: a nourishment lever is never *only* a nourishment lever. Change the amylose fraction and you have changed how the rice cooks, how the bread stales, how the whole thing eats. The flavor index and the nourishment measure are two readings of one plant, and a change to either is felt by both.

β-glucan and resistant starch are the headline levers, but the fiber of a real diet is a wider set, and a compositional account owes the rest a mention. **Arabinoxylan** is the dominant cell-wall fiber of cereal bran — a xylose backbone hung with arabinose, partly soluble and viscous — tunable through the arabinoxylan-synthesis and feruloylation genes that cross-link it into the wall. **Pectin** is the gelling fiber of fruits and the matrix that cooking softens. **Fructans** — inulin and the shorter fructo-oligosaccharides of chicory, onion, and garlic — are soluble, fermentable, and built by fructosyltransferases a designer can dial. Each has its own assay, and we profile them separately rather than lumping them into one "fiber" number, because they behave differently in the one compositional property that most tempts the health halo: **fermentability**. We can measure, in a standardized in-vitro colonic fermentation, how much short-chain fatty acid a fiber yields — and we report that as what it is, a *property of the fiber measured in a flask*, not a promise about a gut. The fermentation number is real and reproducible and ours to publish; what it means in a person is, as ever, the next sentence we decline to write.

So how does fiber enter a framework built on $BD = C \cdot f_access$, when its defining property is that f_access is near zero by design? Honestly, and by exception. For fiber we do not use the

bioaccessible-density score, because the score was built for nutrients that are meant to be released, and fiber is meant not to be. Instead we report fiber as a *compositional* profile — total, soluble, and insoluble fractions by the standard enzymatic-gravimetric assays; β-glucan content and its molecular weight; resistant-starch type and amount by the standardized resistant-starch method — because that profile is the honest, measurable description of what the architecture *is*. We measure it on the food as prepared, since cooking and cooling move starch between the digestible and resistant pools (the retrograded RS3 of a cooled, cooked grain is real and quantifiable, and it is why yesterday's rice is not nutritionally identical to today's). And we stop there, firmly, because the next sentence — the one about what fiber does in a body — is a regulated health claim in most of the world for good reason, and it is not ours to make, however many such claims regulators have authorized for others. Fiber is architecture. We can measure the architecture precisely, design it deliberately, and describe it exactly. What that architecture means for a person is between them, their whole diet, and anyone licensed to advise them — and the discipline that keeps our numbers trustworthy is the discipline of knowing, especially here, where our sentence ends.



PART V OF VI

The Color & the Compound

Anthocyanins · betalains · glucosinolates — where the nourishment pathway and the flavor pathway are one

Chapter 12 · Anthocyanins and the purple design — *the color is the compound is the flavor.*

Of everything in this book, the design of plant color is the place where the nourishment frontier and the flavor index turn out to be the same country seen from two borders — because the pigments that make a fruit purple are built on the very **phenylpropanoid** pathway the platform already runs for flavor, and because those pigments register, faintly but really, on the astringency axis of the Index itself. When we set out to design color, we were not leaving the flavor platform for a nutrition side-quest. We were following a branch of the same tree we had been climbing all along, and finding fruit on it.



The blueberry wears its design on the outside — anthocyanins are the pigment, the astringency, and the polyphenol all at once, three readings of one class of molecule.

The pathway is a continuation of one the Scoville Splice platform knows intimately. The general phenylpropanoid trunk — **PAL → C4H → 4CL**, the same opening steps that feed the capsaicinoid route, lignin, and the folate-PABA branch — divides into the flavonoid pathway proper, and from there into the anthocyanins:

STEP	ENZYME	PRODUCT
entry	CHS (chalcone synthase)	naringenin chalcone
isomerize	CHI (chalcone isomerase)	naringenin
hydroxylate	F3H (+ F3′H / F3′5′H)	dihydroflavonols (sets red vs blue)
reduce	DFR (dihydroflavonol reductase)	leucoanthocyanidins
oxidize	ANS (anthocyanidin synthase)	anthocyanidins (the chromophore)
glycosylate	UFGT	stable anthocyanins (vacuolar)
branch	**ANR** (anthocyanidin reductase)	epicatechin → condensed tannins

The whole branch is governed by a single regulatory switch of remarkable elegance — the **MBW complex**, a three-part transcriptional machine of a **MYB** factor, a **bHLH** factor, and a **WD40** scaffold, that turns the entire anthocyanin program on together. This is why the most dramatic color designs are *regulatory*, not enzymatic: you do not install the flavonoid enzymes one by one, because the fruit already has them, idling. You throw the MBW switch.

The landmark demonstration — and the one we build on — did exactly that. Two transcription factors from snapdragon, **Delila** (a bHLH) and **Rosea1** (a MYB), expressed together under a fruit-specific promoter, flip the tomato's own dormant anthocyanin pathway fully on, turning the flesh a deep purple throughout and lifting anthocyanin content from a trace to on the order of **2–3 mg/g** — a concentration comparable to a blueberry's, in a tomato. The designed purple tomato is not a tomato with pigment painted on; it is a tomato whose own latent machinery has been switched to the position evolution left off. We read that design on three instruments at once, and the three readings are the whole point of the chapter:

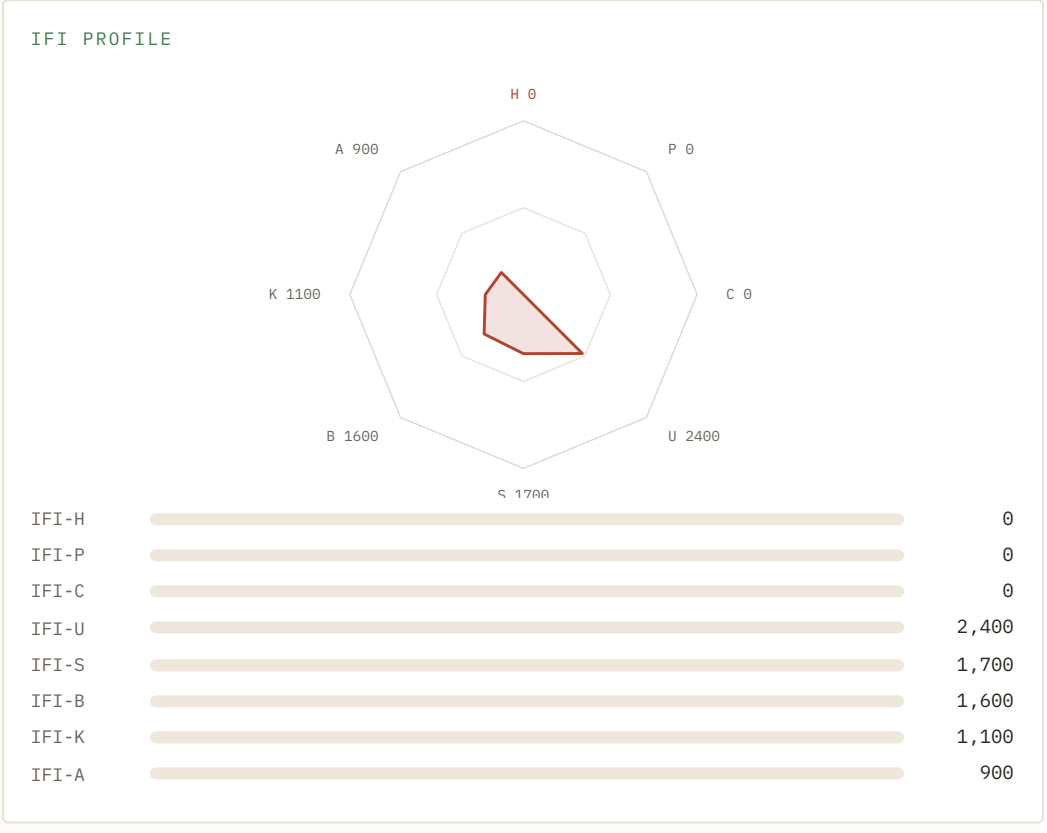

That scorecard is a nourishment-and-anthocyanin tomato read on the *flavor* index — umami and sweetness restored as in any of our designed tomatoes (the platform's 2,040 tomato constructs sit on exactly these axes), and note the **A: 900**, a real astringency the anthocyanins and their companion condensed tannins contribute,

the faint drying grip that a deeply purple fruit carries and that the salivary-protein axis of the Index detects directly. There is a lovely piece of connective tissue here: the **ANR** branch of the very same pathway makes the epicatechin that polymerizes into condensed tannins, which is the chemistry the *Grape & Wine* volume engineers for astringency on purpose. The pigment, the polyphenol, and the astringency are literally branches of one tree. The same molecules show up on the nourishment ledger as polyphenol *content* and *bioaccessibility*, and on the plant-physiology ledger as the pigment and UV-protective compound they evolved to be. Color, flavor, nourishment: one class of molecule, three honest measurements, and not a single one of them a claim about a body. We are at pains about that here because anthocyanins are catnip for the health-halo industry, draped in "antioxidant" promises we will not make. What we will say is exact and sufficient: the fruit holds more of these polyphenols, in a more bioaccessible form, and it registers on the astringency axis accordingly. That is a food described completely. The rest is not ours.



The beet makes the same point in a different pigment chemistry entirely — betalains, not anthocyanins, and the two never coexist in one lineage.

There is a second pigment chemistry worth the detour, because it teaches a lesson about biological constraint that a designer ignores at their peril. The order Caryophyllales — which includes beets, chard, amaranth, and cactus fruit — does not make anthocyanins at all. It makes **betalains**: nitrogen-containing pigments built on an entirely separate route, from tyrosine to **L-DOPA** (via the cytochrome P450 **CYP76AD1**) to **betalamic acid** (via **DODA**, DOPA 4,5-dioxygenase), which condenses with amino acids or cyclo-DOPA into the yellow betaxanthins and red-violet betacyanins that give a beet its color. And here is the constraint: anthocyanins and betalains are **mutually exclusive** — no plant lineage makes both, for reasons still not fully understood. You cannot simply add a betalain to a strawberry or an anthocyanin to a beet as if pigments were paint; the

lineage's deep chemistry decides which palette is even available, and designing betalains into a non-Caryophyllales host means importing the whole L-DOPA route from scratch. It is a humbling, clarifying fact, and it is the kind of thing a design platform learns by respecting the plant rather than overriding it. We design within what the lineage can be taught, not against what it is — and the color, when it comes, is the pathway made visible, the flavor made faintly astringent, and the nourishment made, honestly and only, more present.

Chapter 13 · Glucosinolates and the Brassica lever — *the sharp compound that is bitterness, pungency, and design tension at once.*

The brassicas — broccoli, kale, cabbage, cauliflower, mustard, the whole pungent, divisive, nutritionally dense family — are where the nourishment program runs headlong into the flavor index and has to *negotiate*, because the compounds that make a brassica compositionally interesting are the very compounds that make it taste the way the people who dislike brassicas dislike. This is the most honest chapter about trade-offs in the book, because glucosinolates do not let you pretend there is a free lunch. The molecule you might design *up* for its compositional interest is the same molecule the flavor index reads as bitterness and pungency, and you cannot move one without moving the others. The IFI exists precisely so that this negotiation can be conducted with numbers instead of guesses.



Broccoli is the glucosinolate workbench: dense with glucoraphanin, compartmentalized away from the enzyme that activates it, and sitting squarely on the bitterness and pungency axes of the flavor index.

The chemistry is a two-compartment system of real ingenuity. Brassicas store **glucosinolates** — in broccoli, principally **glucoraphanin**, a 4-methylsulfinylbutyl glucosinolate — intact and inert in the tissue, physically separated from the enzyme that would activate them. That enzyme is **myrosinase** (a thioglucosidase), sequestered in its own idioblast "myrosin" cells. Only when the tissue is damaged — chopped, chewed, crushed — do enzyme and substrate meet, and myrosinase cleaves glucoraphanin to

sulforaphane, the isothiocyanate that is the compound of interest and the source of the sharp, mustardy bite. There is a subtlety the designer must respect: a side protein called the **epithiospecifier (ESP)** can divert that hydrolysis away from the isothiocyanate toward a nitrile instead, so the *same* glucoraphanin can yield more or less sulforaphane depending on ESP activity and on temperature. Knocking ESP down, or warming the tissue briefly to the window where myrosinase survives but ESP does not, steers the chemistry toward the isothiocyanate — which is why how a brassica is prepared changes what it yields as much as how it was bred.

The design levers sit at each stage: the aliphatic-glucosinolate program is governed largely by the **MYB28 (HAG1)** transcription factor (with MYB29 as a partner), whose dose sets glucoraphanin accumulation; the precursor supply can be steered through the methionine-chain-elongation and core-glucosinolate genes; and the activation depends on myrosinase, ESP, and preparation. The real-world precedent is instructive: a high-glucoraphanin broccoli was reached not by transgene but by *introgression* — crossing cultivated broccoli with a wild relative (*Brassica villosa*) carrying a stronger MYB28 allele — a reminder that the newest design goals are sometimes met with the oldest design tool, the deliberate cross.

Glucoraphanin is the glucosinolate we design *up*, but the family is larger and not all of it is wanted. **Indole glucosinolates** (glucobrassicin and its kin) hydrolyze to indole-3-carbinol and related products on a separate schedule, and **progoitrin** yields **goitrin**, a genuinely antinutritional compound a designer wants *down* — so a serious brassica program is a two-sided edit, lifting the aliphatic glucoraphanin while trimming the progoitrin, both reachable through the MYB28/MYB29 and side-chain-modification genes that sort the glucosinolate profile. And there is a twist in the eater that explains why this negotiation is so fraught: bitterness perception at **TAS2R38** is itself genetically variable — the classic PTC/PROP "taster" polymorphism means some people experience brassica bitterness far more intensely than others, from the identical vegetable. That variation lives in the diner's own receptor gene, not in the food, which is exactly why we attenuate bitterness at the level of *perception* rather than by stripping the compound: we leave the glucoraphanin intact and quiet the signal, so one designed line can suit the strong taster and the indifferent one at the same table.

Now the tension, read on the index. A broccoli bred for maximum glucoraphanin is, other things equal, a *more bitter and more pungent* broccoli, because the compounds are the same. Watch what that does to a flavor scorecard, and watch what design does back:

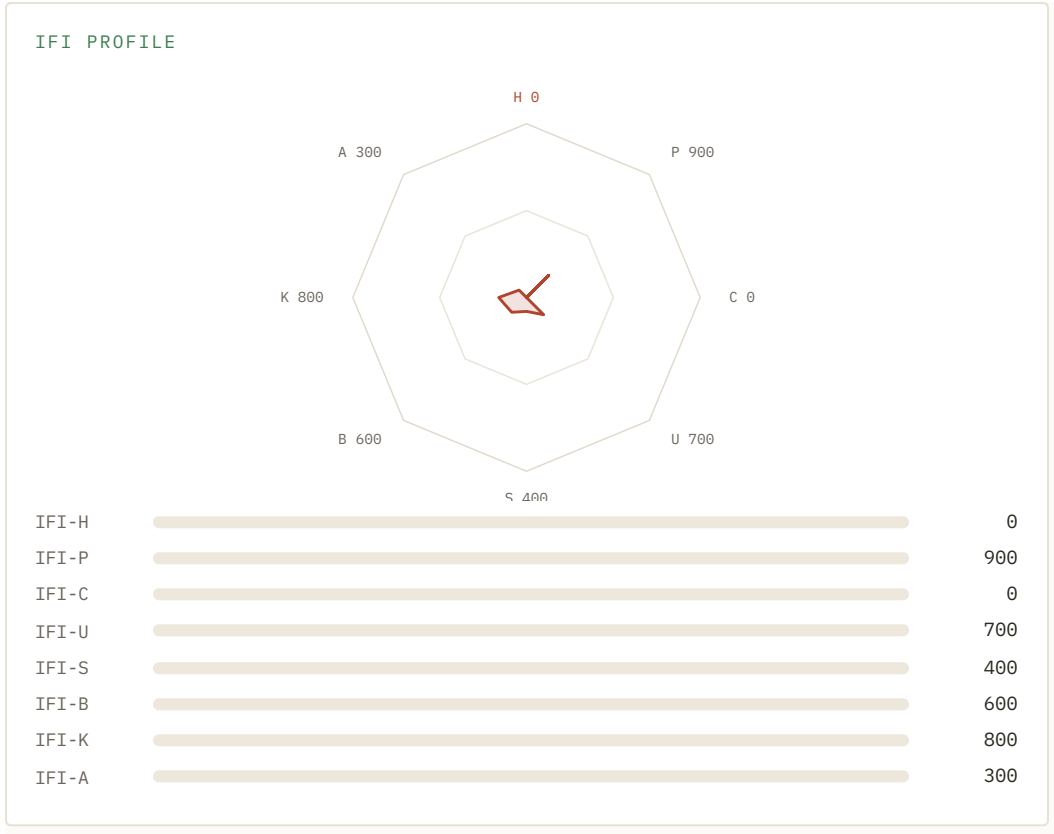

That is the raw high-glucoraphanin line read on the Index: pungency up at **P: 900**, and a *low* zero-bitter score at **B: 600** — remember IFI-B measures bitterness *attenuation*, so a low B means the bitterness stands tall and un-attenuated. It is compositionally rich and, for many eaters, hard going. Here is where the platform earns its keep, because the flavor axes and the glucosinolate content are *separable levers*. The perception of brassica bitterness runs substantially through the **TAS2R38** receptor and its response to glucosinolate-derived compounds; our zero-bitter work (the same work that serves the coffee and cacao volumes, attenuating bitterness 20–65%) can quiet that perception without dismantling the glucoraphanin itself — targeting the receptor-level bitterness rather than the molecule that is the point of the vegetable. The co-designed line reads very differently:

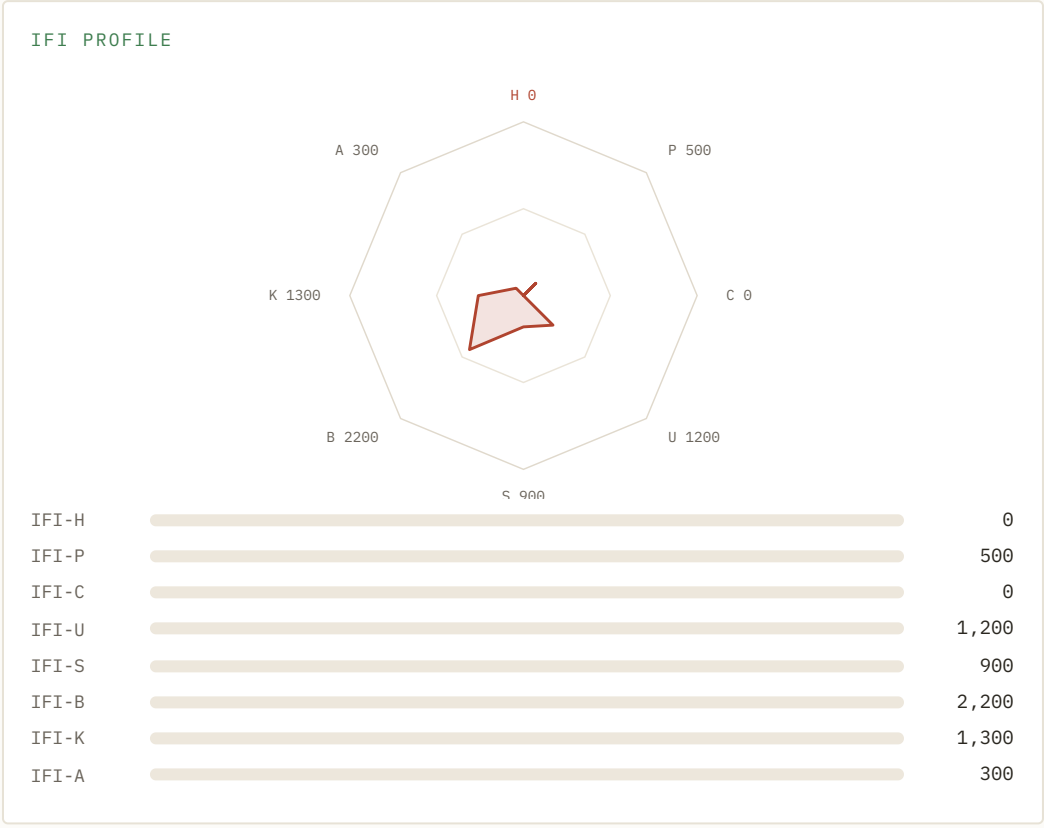

Same glucoraphanin density, but **B: 2200** — the bitterness attenuated to something most palates welcome — and umami and kokumi lifted so the vegetable tastes *rich* rather than merely sharp. This is the negotiation the whole book has been building toward,

conducted in public, in numbers: we did not choose between a compositionally interesting broccoli and a palatable one; we measured the conflict on the Index, found the receptor-level lever that let the two goals separate, and designed both at once. We will not tell you what sulforaphane does in a body, because that is a regulated claim and a contested science and not ours to adjudicate. We will tell you, exactly, that this broccoli holds more glucoraphanin, in a form its own myrosinase can still activate, at a bitterness its eater can actually enjoy — and that the enjoyment is not a footnote to the nourishment but a condition of it, because the most perfectly designed vegetable in the world nourishes no one who will not eat a second bite.



PART VI OF VI

The Frontier

The 64-organism roadmap · the whole plate · the honest edge — what we grow next, and the lines we hold

Chapter 14 · What we build next — *the roadmap, read as nourishment.*

The Beyond-Peppers roadmap — sixty-four concept organisms across five tiers, the imagination-fuel of the whole library — was drawn first as a map of *flavor*: a volcanic tomato, a seismic potato, a dynasty rice, each named for the heat and drama it could carry. Read it a second time, through the lens of this book, and the same map becomes something quieter and, to our family, more moving: a roadmap of *nourishment*, because every organism we have taught to carry designed flavor is an organism whose nutritional pathways we now understand well enough to design too. The platform that routes carbon to a capsaicinoid can route it to a carotenoid; the regulatory switch we throw for a pigment is the switch we throw for a vitamin. What we build next, read as nourishment, is a plan for putting density and availability into the staples that feed the most people — and the roadmap's own tiers are the schedule.



The frontier organisms, rendered: a rice, a carrot, a corn, each a staple whose nutritional pathways the flavor platform has already learned to read.

The tiers sort the work by genetic reach and horizon — the same tiers, with the same gene counts and timelines, that the flavor roadmap uses — and here they are, read as nourishment targets:

TIER	GENES · HORIZON	CONCEPT ORGANISM	NOURISHMENT READING — THE DESIGN GOAL
T1 Solanaceae core	4–5 · 18–24 mo	**Volcanic Tomato** (Ember → Krakatoa)	lycopene + β-carotene with chromoplast-friendly sequestration; folate stacked
T2 Extended nightshades	5–7 · 24–36 mo	**Seismic Potato**	a staple tuber with provitamin A and zinc availability; low reducing-sugar
T3 Cross-family	7–10 · 36–48 mo	**Sweet-Heat Strawberry** (Candied Fang)	anthocyanin + ascorbate co-design; the vitamin-C counter-reagent built in
T4 Frontier synthetic	10–12 · 48–72 mo	**Dynasty Rice**	the full-stack grain — provitamin A + iron + zinc + lysine, phytate down
T4 Frontier synthetic	10–12 · 48–72 mo	**Storm Corn · Wildfire Broccoli**	high-lysine carotenoid maize; high-glucoraphanin, bitterness-negotiated brassica

The crown of the whole roadmap, read this way, is **Dynasty Rice** — the T4 concept that, as a nourishment target, becomes the most ambitious single object in the library: a staple grain carrying provitamin A in its endosperm (Chapter 4's PSY+CRTI relay), iron made available by ferritin and NAS with the phytate driven down (Chapter 7), zinc freed by the same phytase (Chapter 8), and lysine lifted by a feedback-insensitive DHDPS (Chapter 9) — four nutritional frontiers stacked in one grain, each measured on the ladder its nutrient demands. We do not pretend that is near. It is a T4, ten to twelve genes, forty-eight to seventy-two months of concept horizon, and the honest difficulty is not any one of the four traits but their *coexistence* — four pathways, each with its own flux demands, sharing one endosperm's carbon and nitrogen and one plant's agronomic vigor.

This is where a genome-scale **flux-balance** model stops being a nicety and becomes the only way to know whether the grain can afford all four at once. The question is quantitative: the carotenoid pathway draws on GGPP and acetyl-CoA; the lysine boost draws aspartate-derived carbon and nitrogen; ferritin and the phytase are protein-synthesis loads; and the endosperm has a fixed budget of fixed carbon flowing in from the plant. The model lets us ask, before a single edit, whether raising the carotenoid starves the lysine, whether the phytase expression taxes the seed's own phosphorus economy, and where the true rate limiter sits — exactly the analysis that, for the pepper, found control sits upstream in the phenylpropanoid pathway — at 4CL and HCT — not at the terminal Pun1 step, before we finalized the design. Read on the flavor index, a nourishment-forward Dynasty Rice is deliberately quiet — a staple is not a superhot, and a grain eaten three times a day earns its place by being gentle, not loud:

IFI PROFILE

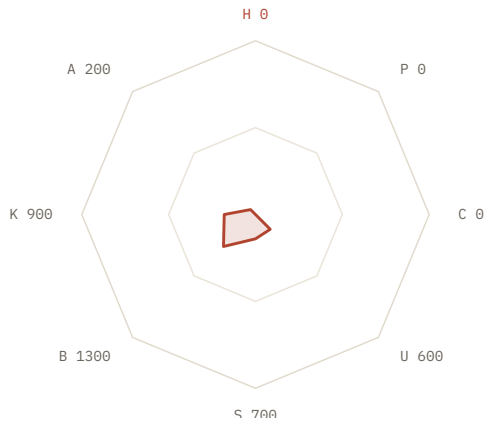

IFI-H	0
IFI-P	0
IFI-C	0
IFI-U	600
IFI-S	700
IFI-B	1,300
IFI-K	900
IFI-A	200

A little sweetness, a little savory roundness, bitterness kept low so nothing fights the daily bowl — the flavor of a grain designed to disappear into a thousand dishes while carrying four nutritional frontiers inside it. The roadmap's drama was always its heat. Its gravity, to us, is this: the same organisms, the same methods, pointed at the density and the availability of the foods that carry the most people through the most days.



We have designed enough of these organisms to know the difference between what is near and what is far — and to say so plainly, which is its own kind of design discipline.

A word on the method that makes the roadmap more than a wish list, because it is the same method that wrote the two Pun1 variants into the pepper that opened this library. Every one of these nutritional enzymes — a feedback-insensitive DHDPS, a γ-tocopherol methyltransferase pushed for throughput, a phytase stable enough to work in the gut — is a protein, and a protein is something the platform can **score and redesign**. An ESM2 650M-parameter protein-language model ranks every possible single-residue substitution by predicted effect on catalytic efficiency times fold stability; a genome-scale flux-balance model confirms where the true rate limiter sits; CRISPR writes the winners in. For the gut

phytase, where thermostability is the whole battle, the output looks like the Pun1 table it is modeled on — a short list of substitutions that buy fold stability without costing activity:

EDIT (ILLUSTRATIVE, GUT-STABLE PHYTASE)	EFFECT	ΔΔG (KCAL/MOL)
D52N	removes a pH-labile surface charge	−1.6
G118P	rigidifies a floppy loop near the active site	−2.0
K204E	new salt bridge, stabilizes the fold at low pH	−1.3

The frontier is not a hope that these traits can be built. It is a queue of design problems the platform already knows how to pose — and the only honest uncertainty is time, flux, and the plant's own vote, which we have learned, over forty-two organisms and roughly 19,500 product SKUs backed by around 225 distinct construct designs, to ask before we assume.

Chapter 15 · Designing the whole plate — *nourishment is a meal, not a molecule.*

For fourteen chapters we have designed nutrients into plants one pathway at a time, and there is a danger in that precision: it can make nourishment sound like something that happens inside a single engineered cultivar, when in truth it has always happened across a *plate*. The oldest wisdom in human eating is combinatorial — the grain with the pulse, the iron with the citrus, the fat-soluble vitamin with the spoonful of oil — and the most sophisticated thing a design platform can do is not to cram every nutrient into one heroic vegetable but to compose a meal in which each designed food covers for the others, exactly as the traditional pairings always did, now with the levers measured instead of inherited. This chapter is where the nourishment measure and the menu become one act, and where our family's own way of eating — plant-forward, reverent, unhurried — stops being a preference and becomes a design principle.



The harvest tray is the real unit of nourishment — not one cultivar, but the composition of many, each covering the others' gaps.

Consider a single supper we actually eat, built from the levers of this book — a bowl of designed grain and pulse, dressed with the counter-reagents that free what the grain and pulse contain.

The everyday bowl — *a plate that covers its own gaps*

- Designed high-lysine grain (Chapter 9), cooked and cooled to raise its resistant-starch fraction (Chapter 11)

- A pulse, cooked long enough to denature its protease inhibitors (Chapter 9) — the lysine covers the grain's limit; the grain's methionine covers the pulse's
- A low-phytate preparation — soaked, or built on an Ipa line (Chapter 3) — so the iron and zinc come loose
- Designed tomato, folate-rich and ascorbate-lifted (Chapters 5, 6), cooked into the dish — the vitamin C reduces the iron and defeats the phytate (Chapter 7)
- A spoonful of an omega-structured oil (Chapter 10), stabilized by its own α-tocopherol — and the lipid that micellizes any carotenoid in the bowl (Chapter 2)
- Greens chosen for low oxalate (Chapter 8), so their calcium is actually available
- A little designed chili for pleasure and warmth, because nourishment that forgets delight is not nourishment we would serve

Read on the flavor index, that bowl is not dramatic — and it is not meant to be:

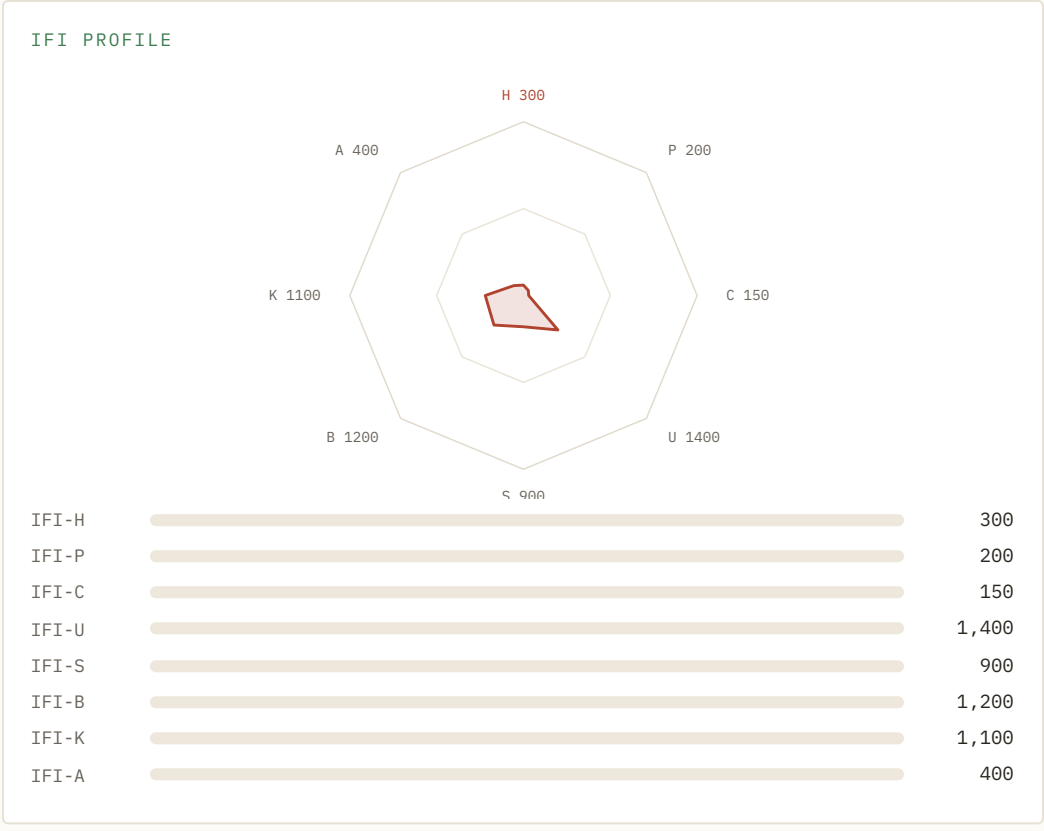

Modest heat, real umami and kokumi depth, bitterness talked down, a whisper of astringency from the greens and the oil — the profile of a meal you could eat every day, which is the only kind of meal that nourishes anyone, because nourishment is a habit before it is a nutrient. And read on the nourishment ledger, the bowl's whole architecture reveals itself as a set of *multiplications*, each pairing turning a bound nutrient into an available one:

PAIRING IN THE BOWL	MECHANISM	WHAT IT MULTIPLIES
ascorbate (tomato) × iron (pulse)	$Fe^{3+} \rightarrow Fe^{2+}$, soluble chelate	non-heme iron f_access, severalfold
lysine (grain) × methionine (pulse)	mutual complementation	DIAAS of the whole above either part
oil × carotenoid (greens, tomato)	micellization	fat-soluble vitamin f_access
low-oxalate greens × their calcium	less chelation	calcium f_access
cook-and-cool × starch	retrogradation	resistant-starch fraction

No single cultivar in the bowl is a miracle. The composition is the design — and a second plate makes the same point in a different key, for a cook who wants the frontier at breakfast:

The morning plate — *soft eggs of the garden, or not*

- A folate- and ascorbate-rich tomato, barely cooked, over a slice of high-β-glucan barley bread (Chapter 11)
- Wilted low-oxalate greens with a thread of the omega-structured oil
- A scatter of toasted seeds for the sulfur amino acids and the vitamin E
- Optional: an egg, or a cured fish — written here without apology, for the reasons the next passage gives

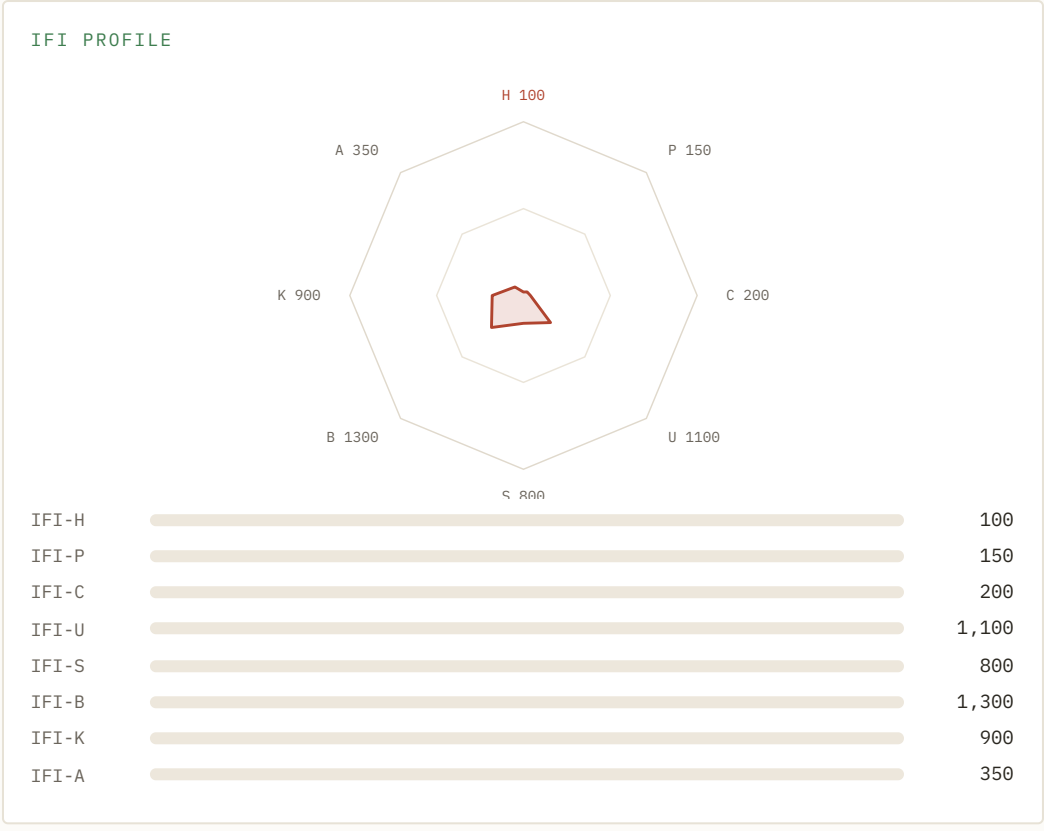

There is one more plate-level lever, and it is about *timing* rather than composition: the inhibitors that live in other foods. The tannins in tea and coffee bind non-heme iron as surely as phytate does, so an iron-forward meal and a strong cup are better kept an hour or two apart — a piece of folk sequencing that turns out to be exact chemistry, and one no amount of in-the-plant design can override, because the inhibitor arrives in a separate cup. We mention it because it marks the edge of what cultivar design can do: we can lower the phytate inside the bean and raise the ascorbate inside the

tomato, but we cannot reach into the glass the eater pours beside them. The plate is a system, and some of its levers are held by the cook and the clock, not the breeder — which is the humbling, correct note to end the design chapters on.

A word, with respect, about the foods not always on those plates. This is a vegetarian-first book from a plant-reverent family, and the bowls above are how we most often eat — but the nourishment frontier does not end at the edge of the garden, and we will not write about meat and fish with either apology or neglect. Heme iron, as Chapter 7 made plain, is absorbed by a route no plant iron travels; vitamin B12 is made by neither plant nor animal but by microbes, and reaches most diets through animal foods; and the long-chain omega-3s of Chapter 10 are, until the oilseeds finish the chain we are teaching them, overwhelmingly the sea's to give. A cook feeding a body that needs those things is doing honest nourishment work, and a design platform that pretended otherwise would be practicing ideology, not science. Our lean is plants, and our reverence is for the living thing on the plate whatever it is — grown or raised or caught — and for the eater deciding, with a whole diet in view, what their own nourishment requires. We design the plant to carry the most it can. We set the plate so each food frees the next. And we leave the table, as we always have, to the family sitting at it.

Chapter 16 · The honest edge — *the discipline that keeps a nourishment program trustworthy.*

A book that spends fifteen chapters raising the nutrient content of food owes its reader one chapter on the lines it will not cross, because the distance between nutrition and quackery is measured entirely in the claims a program is willing to make, and we have built this whole shelf on a refusal. The refusal is simple to state and relentless to keep: we describe what a food *holds* and what a food *releases*, both measured on a bench, and we never once describe what a food *does* to a person. This is not timidity. It is the only posture under which precise numbers stay worth anything — because the moment a measurement is pressed into the service of a health claim it cannot support, every other measurement the program has ever published becomes suspect. The honest edge is where a nourishment program earns the right to be believed about the things it *can* prove.



Reproducibility is not a footnote to nourishment; it is the whole of its credibility. A number we cannot make again is a number we had no business printing.

The regulatory world has a vocabulary for this line, and we hold to it deliberately. A **content (nutrient) claim** — "a good source of iron," "high in folate" — is a statement about the food's composition, verifiable by assay, and it is the kind of claim our numbers support directly. A **health claim** — any assertion that a food prevents, treats, mitigates, or reduces the risk of a disease — requires a standard of evidence (significant scientific agreement in the United States; the authorized-claim processes under Articles 13 and 14 in the European Union, adjudicated by EFSA) that no single cultivar's composition

can meet on its own, and it is a kind of claim we do not make, full stop. Between them sit **structure/function** statements, which are more permissive and more treacherous, and which we approach with extreme conservatism, because the health halo lives precisely in the gap between what is technically sayable and what is honestly supportable. Our rule is stricter than the regulations require: if a sentence implies a benefit to a body, it does not go in a Scoville Splice book, a label, or a word of our marketing, regardless of whether a looser reading of the rules might permit it. "Holds twenty-five times the folate of its reference tomato" is a sentence we can stand behind to the last decimal. "Good for you" is a sentence we decline to write.

CLAIM TYPE	EXAMPLE	STANDARD	OUR STANCE
Content / nutrient	"high in folate," "good source of iron"	compositional, assay-verifiable	**we make these** — it is what our numbers are
Structure / function	implies an effect on the body's workings	permissive; easily abused	avoided; too close to the halo
Risk-reduction / health	"reduces the risk of X"	SSA (US) / EFSA Articles 13–14 (EU)	**we do not make these**, ever
Nourishment score (ours)	"N = 4,500 vs reference bean"	our own food-property measure	a composition-and-release metric, not a benefit

There is also a regulatory reality to the *foods* themselves, which we neither hide nor dramatize: a plant with an engineered nutritional trait is, in most jurisdictions, a regulated article, evaluated for food and environmental safety before it goes anywhere near a plate — the framework under which provitamin-A rice, for instance, received food-safety clearances from regulators in several countries and cultivation approval in others, and under which an anthocyanin tomato was reviewed and cleared for the United States market. We build inside that framework, on the assumption that a nourishment trait worth having is a nourishment trait worth proving safe, slowly and in public.

That discipline has a quantitative backbone, and it is the same one that runs through every volume in this library: **reproducibility**. A nourishment number that cannot be made again is not a finding; it is an anecdote in a lab coat. The platform's standard — a batch-to-batch coefficient of variation at or below **≤7%** — applies to the nutrient assays exactly as it applies to the flavor scores, because a folate content or an iron-availability figure that swings 40% between harvests describes nothing a formulator or a family can rely on. And the data lives in the open: the program's **21 peer-reviewed papers** are posted to Zenodo under **CC BY 4.0**, methods and all, so that the *f_access* protocol we used, the Caco-2 conditions, the HPLC calibration, can be checked and repeated by anyone who doubts them. Open data is not generosity; it is the price of being believed. A nourishment claim held privately is a nourishment claim asking to be taken on faith, and we would rather be checked.



The team at the bench is where the discipline is actually kept — not in a policy, but in the hundredth time someone runs the assay again to be sure.

THE DISCIPLINE	WHAT IT MEANS IN PRACTICE	WHY IT PROTECTS THE READER
Content over health claims	describe composition and release; never a bodily outcome	keeps every real number trustworthy
f_access is a model	report which rung (digestion tube vs Caco-2); flag it as a food property	never mistakes a tube for a person
CV ≤ 7%	a number must reproduce across batches	a figure you can actually build on
Open data (CC BY 4.0, Zenodo)	methods and assays public and repeatable	belief earned by verification, not asserted
Regulated as food	safety review before the plate	a trait worth having is worth proving safe
Cultivar, not prescription	we design the food; the diet is the eater's	respects the person and their advisors

These are computational designs; Scoville Splice's cultivars and constructs have not been experimentally validated in a laboratory. The nourishment scores in this book are predictions built on the same flux-balance and protein-design models that produced the flavor index, and they carry that honest qualifier everywhere they go.

There is a deeper reason we hold the edge this hard, and it is not legal and not even scientific; it is about what we think food is *for*. A health claim turns food into medicine — into a thing you take because of what it will prevent — and the whole sensibility of this family, and of this shelf, is that food is not medicine. Food is nourishment and pleasure and the table you keep, and the moment it

becomes a dosage it stops being any of those things. We design plants to carry more of what a body can use because we revere the plant and the people it feeds, not because we are prescribing anything to anyone. The honest edge, in the end, is where our reverence and our rigor turn out to be the same discipline: we will not cheapen the food by overclaiming it, and we will not cheapen the science by letting it say more than it knows. The plant holds more, and more comes loose, and we can prove both. Everything past that belongs to the person at the table — which is exactly where we have always believed it belonged.

Closing · The table we are growing toward

We began with a question asked over clean plates — *what else is it giving them?* — and sixteen chapters later we have an instrument to answer the only part of that question a plant can honestly answer: how much a food holds, and how much it lets go. It is less than the question asked, and we have been careful, the whole way, to say so. A plant cannot promise a person anything. It can only be built, with reverence and with arithmetic, to offer the most it can — and then set on a table, beside other foods that free what it carries, in front of people who will decide for themselves what their nourishment needs.



Most mornings end where they began — a last walk down the rows, the lights warming, the next harvest already deciding what it will hold.

That is the whole of the nourishment frontier, and it is enough. The eight axes taught us to design what a food *says* to the tongue. The ninth measure — density times availability, anchored against a reference, held to the same ≤7% and posted in the open — lets us design what a food *holds* for the body, without once pretending to speak for the body itself. Provitamin A written into an endosperm. Iron shaken loose from its phytate by the vitamin C we built in the next row over. A grain and a pulse, each covering the other's limiting link, the way every grandmother's kitchen already knew. Four nutritional frontiers stacked, someday, in a single grain of Dynasty Rice — if the flux allows it, and the plant agrees.

The computing that makes this possible was one pair of hands; the questions, the tasting, the table where it all gets argued over and eaten, belong to our whole family, and always have. We built a company to put a number on flavor so we could design it rather than chance it. We built this frontier to do the same for nourishment — not as a claim, never as a cure, but as an intention written into the plant itself: that food should be grown to carry the most it can, that a body should be able to reach the most of what it carries, and that designing flavor and designing nourishment were always, if you looked closely enough, the same reverent act. We are still at the beginning of it. But the rows are planted, the instrument is built, and the table we are growing toward is one where *what else is it giving*

them finally has an honest answer — measured, reproducible, and offered, as everything at our table is offered, with love and without a single overclaim.