

THE SCOVILLE SPLICE LIBRARY · VOLUME IV

Beyond Peppers

Capsaicin was never a pepper problem. *It was a method the whole kingdom can inherit.*

The thesis and the roadmap — that heat was only the first axis, the pepper only the proof, and the real discovery a way of designing flavor that runs in every vascular plant. Forty-two organisms live, sixty-four on the map, one measurement spine, and a whole garden still to write.

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

Contents

01 Opening essay · The pepper was only the proof

02 Chapter 1 · The pepper was never the point — the method and the molecule.

03 Chapter 2 · The shared trunk — the oldest road in the plant.

04 Chapter 3 · The four last enzymes — the pepper's private invention.

05 Chapter 4 · The portable platform — five of seven modules move.

06 Chapter 5 · The arithmetic of loudness — one number for any plant.

07 Chapter 6 · Forty-two organisms, nineteen thousand constructs — the library that already exists.

08 Chapter 7 · Reading difficulty before we build — the ladder, and the arithmetic under it.

09 Chapter 8 · Tier 1 · The Solanaceae core — the pepper's own family answers first.

10 Chapter 9 · Tier 2 · The extended nightshades — the family, one ring out.

11 Chapter 10 · Tier 3 · Crossing the family line — out of the nightshades at last.

12 Chapter 11 · Tier 4 · The frontier synthetic — the hardest jump in the garden.

13 Chapter 12 · Tier 5 · De novo — the tier we have not drawn.

14 Chapter 13 · Heat was only the first axis — the method across all eight.

15 Chapter 14 · The nourishment frontier — the axis beyond flavor.

16 Chapter 15 · The guardrails — what could be built, and what should.

17 Chapter 16 · The economics of a whole kingdom — the matrix of the buildable.

18 Chapter 17 · The drawn and the undrawn map — everything the alphabet can spell.

Opening essay · The pepper was only the proof

The first time we understood what we had built, we were not looking at a pepper. We were looking at a tomato — a supermarket tomato, the kind picked green and ripened in a truck, the kind that has forgotten how to taste like itself. We had run its genome through the same pipeline we used for chilies, and the model had come back with a flat, almost rude answer: the machinery is all here. The savory depth, the roundness, the long low note a tomato is supposed to carry — none of it had been bred out. It had been switched off, or starved, or shipped away. The plant still knew the word. It had only stopped saying it loudly enough to hear.

That is the moment this book is about. Because the thing we had spent years building was never, in truth, a way to make peppers hotter. Peppers were the proving ground — the loudest, most legible flavor in the plant kingdom, the one place where a single number had already been invented (Scoville, 1912) and a single molecule, capsaicin, did almost all the work. If a method could be made to design *that*, cleanly and reproducibly, gene by gene, then the method was the discovery, and the pepper was only the proof.



The research plot at the hour we like best — rows of edited plants reaching past the pepper into every family we mean to reach next.

Here is the argument, stated once and plainly, so the rest of the book can earn it. The fire of a chili is built on a chemical trunk — the phenylpropanoid pathway, PAL → C4H → 4CL → HCT → COMT — that runs in *every vascular plant on Earth*. Lignin is built on it. The scent of a rose is built on it. The trunk is universal; it is the oldest industrial district in the plant cell. Only the last four enzymes are the pepper's private invention — pAMT, Pun1/CS, KAS, BCAT — the four steps that take a precursor every plant already makes and finish it into capsaicin. Four enzymes. That is the entire distance between a bell pepper and a weapon. In fact a bell pepper *is* a weapon with its

last two genes switched off: break Pun1, break pAMT, and the fire simply never arrives. Nature had already shown us the switch. We only learned to install it somewhere new.

If the terminal steps are a module, and the trunk is universal, then capsaicin is not a pepper problem. It is a worked example. And the worked example generalizes in two directions at once — across *molecules* (heat was only the first of eight axes we can now read and write) and across *organisms* (what we did to one nightshade we can do, in principle, to the whole garden). This book maps both directions. It is the least modest book in our library, and we have tried to make it the most rigorous, because a large claim deserves a large amount of arithmetic standing behind it.

We should be honest about where we stand on that map. Forty-two host organisms are live today, carrying roughly nineteen thousand engineered constructs between them, scored on the same eight axes, reproducible to a batch-to-batch coefficient of variation at or under seven percent. Sixty-four more organisms are drawn on the roadmap, sorted into five tiers of ascending difficulty — from a volcanic tomato that is four or five edits away, to monocot grains that will take the better part of a decade, to a fifth tier we have deliberately left blank because we do not yet know how to draw it. That blank is not a failure of nerve. It is the most honest part of the map.

A word about who "we" are, since a book this ambitious should not hide behind a corporate plural. Scoville Splice is a family. The computing — the models, the edits, the long nights against a sequencer — was one pair of hands. The garden, the table where the first question was asked, the palate that decided when a thing tasted *right* and not merely correct: those belonged to all of us. We write as *we* because that is the truth of the work, and because a flavor, in the end, is a thing you hand to someone you love and watch them taste.

So this is the plan of the book. Part One makes the thesis exact: the trunk, the four terminal enzymes, the portable platform, and the arithmetic of loudness that lets us score any of it. Part Two is the map we can already read — the forty-two live organisms and the five tiers of difficulty that tell us, before we lift a pipette, roughly how hard a new plant will be. Part Three is the roadmap proper: the sixty-four concepts, tier by tier, from the Solanaceae core out to the synthetic frontier and the undrawn edge beyond it. Part Four widens the lens one last time — past heat into all eight axes, past flavor into nourishment, and into the guardrails and the economics that decide which of the ten thousand buildable things we actually build.

We did not set out to make peppers hotter. We set out to learn whether flavor could be read like a text and written like one. The pepper was where we learned the alphabet. What follows is everything the alphabet can spell.



PART I OF IV

The Thesis

Why capsaicin is a method, not a pepper — the trunk, the four last enzymes, the portable platform, and the arithmetic of loudness

Chapter 1 · The pepper was never the point — *the method and the molecule.*

Hold a ripe superhot up to a window and the light does not pass through it so much as pool inside it. Cut it — gloved, always gloved at this end of the ladder — and the placenta, the pale ribbed tissue that cradles the seeds, is where the heat lives, beading as oil along the septa in a way you can photograph but should not touch. Everything a chilihead reveres is in that cross-section: the color, the architecture, the threat. And almost none of it is what we spend our days on. We spend our days on the sentence the plant is writing when it lays that oil down — the order of the enzymes, the identity of the precursors, the single rate-limiting step that decides whether the fire reaches two hundred thousand Scoville or nearly three million.



Collateral at the cut — the placental tissue, where capsaicin beads along the septa. The reverence is here; the work is one layer down, in the enzymes that put it there.

The distinction between the molecule and the method is the whole of this book, so let us be precise about it. A *molecule* is a destination: capsaicin, 8-methyl-N-vanillyl-6-nonenamide, a single compound with a single receptor and a single, famous effect. A *method* is a road: a reproducible sequence of operations — read the pathway from the genome, score every possible improving mutation, confirm the true bottleneck, write the edits, measure the result against a fixed receptor-level standard — that can be driven to many destinations. For ten thousand years, agriculture worked on

molecules without knowing their names. A farmer tasted the fruit, kept the seed of the one that moved them, and hoped the next generation held the line. It was selection on the *phenotype* — the destination — with the *genotype* — the road — left entirely dark. It worked, slowly, and it worked blind. It could not tell you why the hot one was hot, and so it could not promise you the next one would be.

There is a history worth pausing on, because it is the history of exactly this blindness. In 1912 a pharmacist named Wilbur Scoville devised the first number ever attached to heat: he diluted a pepper extract in sugar water until a panel of tasters could no longer detect the burn, and the dilution factor became the pepper's rating — the Scoville Heat Unit that still bears his name. It was an ingenious act, the first time flavor was ever made to yield a figure, and it was also entirely a measurement of the *destination*. Scoville could tell you how hot a pepper was once it existed; he could not tell you how to make it hotter, or why one plant in a field had run away from its neighbors. For the century that followed, heat was a thing you measured after the fact and chased by selection, tasting your way up a ladder whose rungs you could count but never cut. Our inversion is the whole of the distance from 1912 to now. We no longer taste up the ladder; we read the road that builds each rung, score the edits that would add the next one, and write it in. The Scoville number remains — our flagships are rated in it, Phoenix at 2,937,948 — but it has changed from a verdict into a target, from the thing we discover at the end into the thing we specify at the start. A figure that was once the last word on a pepper is now the first line of its design brief. And a figure you can specify in advance is a figure you can specify for any plant, which is the hinge on which this entire book turns.

We inverted the blindness. The platform does not select on the fruit; it reasons on the road. And the instant you are reasoning on the road rather than tasting the destination, a strange and liberating thing happens: the pepper stops being special. A road that reaches capsaicin in *Capsicum chinense* is, structurally, a road that reaches a vanilloid amide in anything with the upstream trunk — and the upstream trunk, as the next chapter will show in full, is everywhere. The pepper was not the point. The pepper was the first place the road happened to end, because the pepper is where evolution had already laid down the last four paving stones.

This reframing is not rhetoric; it changes the unit of ambition. If you believe you are in the pepper business, your horizon is the hottest pepper — a single, lonely record, a Scoville number to chase and then to defend. We chased it, and we hold it: our flagship superhots climb a measured ladder from **Ember** (SS-0087, 866,000 SHU) to **Phoenix** (SS-0096, 2,937,948 SHU), and we will walk that ladder rung by rung in a moment, because it is where the method was forged. But a record is a destination, and destinations are finite. There is a hottest pepper the way there is a tallest mountain, and once you stand on it there is nowhere up to go.

A *method* has no such ceiling. The question it poses is not "how hot can one fruit be" but "how many fruits can be taught to speak" — and that question opens onto the forty-two organisms we have

already edited, the sixty-four we have drawn, and the far larger number we could name tomorrow by pointing at a cell in a matrix. Phoenix itself is the hinge between the two ambitions. It is the hottest thing we have ever grown, yes. But the reason it matters to *this* book is not its heat. It is that Phoenix is the first cultivar engineered to register on all eight sensory axes at once — proof, inside a single pepper, that one organism can be made to speak the entire language, not just shout its one famous word. If a pepper can hold all eight axes, and the method that put them there is organism-agnostic, then the thesis is already half-demonstrated before we leave the genus. The rest of the book is the other half: carrying it out of the genus entirely.

So we will treat the pepper, for the remainder of these pages, the way a linguist treats their first fluent language — with affection, with debt, and without the illusion that it is the only one worth speaking. Everything we know, we learned from the chili. Nothing we know stops at the chili. That is not a diminishment of the pepper. It is the largest possible compliment: it was a good enough teacher that we outgrew the classroom.

Chapter 2 · The shared trunk — *the oldest road in the plant.*

Every flavor molecule we have ever designed begins its life in the same place, on a road so old and so universal that calling it a pepper pathway is like calling the alphabet an English invention. It is the **phenylpropanoid pathway**, and it starts from a single amino acid — phenylalanine — that every plant makes in bulk because it must. From that one starting material the road runs through five conserved enzymes, and from the compounds it produces, a plant builds the stiffening in its stems, the pigment in its flowers, the defense in its bark, and the scent that pulls a bee across a meadow. It is not a specialty district. It is the plant's main thoroughfare, and it has been running, essentially unchanged, since plants first hauled themselves onto land.

To feel how universal the trunk is, follow where its traffic goes once it leaves the five shared steps. One branch runs to the monolignols and becomes lignin — the stiffening that lets a plant stand up against gravity at all, the second most abundant organic polymer on Earth after cellulose. Another runs to the flavonoids and becomes the ultraviolet sunscreen in a leaf's upper cells. Another becomes the anthocyanins that redden an apple and purple an eggplant; another the tannins that dry the mouth on an unripe fruit; another the scent compounds a flower spends to summon a pollinator across a meadow. A rose's perfume, an oak's heartwood, a blueberry's color, and a chili's fire all set out from the same five enzymes and diverge only near the end. This is why we call the phenylpropanoid pathway the trunk and these destinations its branches: they share a single base and fork late, and the fork that interests us — toward the vanilloid precursor of capsaicin — is simply one more exit off a road already crowded with traffic the plant runs for its own reasons. When our cartography module opens a new genome and finds the trunk, it is not finding a rare or delicate thing. It is finding the busiest intersection in the plant, maintained at the plant's own expense, with room at its edge for exactly one more turning lane.



Reading the trunk, base by base. The phenylpropanoid road is written into every sequenced plant genome we open; only its final turning is the pepper's own.

Here is the road, drawn once so we can refer back to it for the rest of the book. Each step is an enzyme; each enzyme hands its product to the next:

STEP	ENZYME	CONVERTS	INTO
1	**PAL** · phenylalanine ammonia-lyase	L-phenylalanine	*trans*-cinnamic acid
2	**C4H** · cinnamate 4-hydroxylase	cinnamic acid	*p*-coumaric acid
3	**4CL** · 4-coumarate:CoA ligase	*p*-coumaric acid	*p*-coumaroyl-CoA
4	**HCT** · hydroxycinnamoyl transferase	*p*-coumaroyl-CoA	caffeoyl shikimate → caffeic acid
5	**COMT** · caffeic acid *O*-methyltransferase	caffeic acid	ferulic acid → the vanilloid precursor

Read that table as a reverent object, because it is one. PAL, the first enzyme, is among the most abundant and most studied proteins in the entire plant kingdom; it is the gate through which something like a third of all organic carbon on land eventually passes on its way to lignin. C4H is a cytochrome P450, bolted to the endoplasmic reticulum, hydroxylating with the same chemistry your own liver uses on a drug. By the time the road reaches COMT and lays down ferulic acid, it has produced a molecule one short, modular tailoring step away from vanillin — the vanilloid ring that is the *front half* of capsaicin, the "vanillyl" in 8-methyl-N-vanillyl-6-nonenamide.

This is the load-bearing fact of the entire company. **The vanilloid half of capsaicin is not a pepper invention.** It is a lightly modified exit off the oldest road in the plant, and that road is present, sequenced, and running in every organism we have ever opened. When our pathway-cartography module reads a new genome — a strawberry, a carrot, a grain of rice — it does not go looking for the trunk, hoping to find it. It finds it the way you find the sun. The only real question a new organism poses is never "does it have the trunk," but "what does it do at the turning where the pepper turns off toward fire."

Consider what that means for the shape of the work. If designing capsaicin in a pepper required building the whole road, the method would transfer to nothing — every new plant would be a decade of basic biochemistry. But the road is already built, everywhere, maintained at the plant's own expense for the plant's own reasons. We are not road-builders. We are off-ramp engineers. We arrive at an intersection that evolution has maintained for four hundred million years, and we add a turning lane. The universality of the trunk is precisely why "beyond peppers" is an engineering roadmap and not a fantasy: the expensive ninety percent of the pathway comes free with the genome, bundled, pre-installed, load-tested across geologic time.

There is a nourishment note hiding in this chemistry, and it is worth naming once, quietly, here where it belongs. The compounds streaming down this trunk — the hydroxycinnamic acids, ferulic and caffeic and their kin — are the plant's own antioxidants and

structural investments, the molecules it spends its hard-won carbon to protect itself. A plant grown well, grown unhurried, with the trunk running at full volume, is a plant laying down more of these compounds, not fewer. We will return to this in Part Four as a design axis of its own. For now it is enough to notice that the road we borrow to build flavor is the same road the plant drives to build its own integrity. We are not bolting a flavor onto a stripped-down crop. We are turning up a conversation the plant was already having with itself — and asking it, at the last junction, to say one new thing.

That last junction is the subject of the next chapter. Five enzymes carry every plant to the edge of the vanilloid precursor. Four more — and four only — carry the pepper the rest of the way into fire.

Chapter 3 · The four last enzymes — *the pepper's private invention.*

Everything universal ends, and the pepper's privacy begins, at a precursor the whole kingdom can make. From there, four enzymes — and, as far as the genome tells us, four enzymes only — do the work that no other plant on Earth has evolved to do. They are the reason a chili is a chili. They are also, precisely because they are so few, the reason the fire is portable. You do not transplant a pathway of fifty genes. You transplant four. We have spent more time with these four proteins than with any other objects in our lives, and we have come to think of them less as enzymes than as the four last paving stones on a very long road.

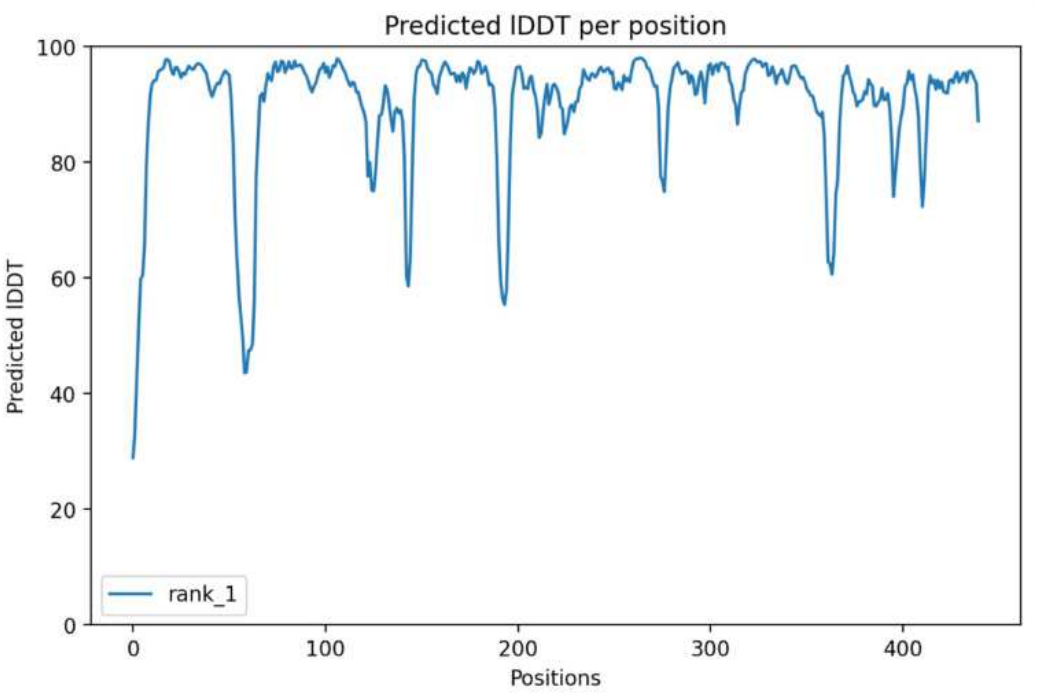

FIG. 1 — Pun1/AT3, the capsaicin synthase, colored by model confidence (pLDDT). The binding domain and active-site cleft resolve sharply; this is the last stone, the one the whole method learned to lay.

The four split cleanly into two sub-assemblies — a *head* and a *tail* — that are built separately and joined at the very end:

ENZYME	ROLE	BUILDS	PRODUCT
pAMT · putative aminotransferase	transaminates the vanilloid	the *head*	`vanillylamine`
BCAT · branched-chain aminotransferase	initiates the fatty-acyl branch from valine / leucine	the *tail* (start)	`branched-chain α-keto acids`
KAS · β-ketoacyl-ACP synthase	elongates and tailors the branch	the *tail* (finish)	`8-methyl-6-nonenoic acid`
Pun1 / CS · capsaicin synthase (AT3 acyltransferase)	condenses head and tail	the *bond*	**`capsaicin`**

The *head* is one enzyme. pAMT takes vanillin — that modular exit off the universal trunk — and hangs an amine on it, turning an aldehyde into vanillylamine. This is the junction a bell pepper has broken. A sweet pepper carries a loss-of-function pAMT or a loss-of-function Pun1, sometimes both; the plant builds the entire road, arrives at the precursor, and then cannot make the final turn. It is one of the most instructive natural experiments in all of flavor: evolution hands us, in the ordinary bell pepper, a complete capsaicin pathway with its last light switched off, proof that the terminal steps are a clean, separable module and not an inextricable part of the plant's nature. The sweet pepper is our existence proof that the switch exists. The method is simply the discipline of installing that switch where it has never been.

The *tail* is two enzymes and a quiet piece of elegance that most people never hear about. BCAT pulls carbon out of the branched-chain amino acids — valine, leucine — and KAS elongates that carbon into a short, branched fatty acid, 8-methyl-6-nonenoic acid. Note the *6-nonenoic*: there is a double bond, a single site of unsaturation, at the sixth carbon. That one double bond is the entire difference between capsaicin and its shadow, dihydrocapsaicin, which carries the saturated chain instead. And that difference, in the mouth, is the difference between the *shape* of two burns. Capsaicin, with its double bond, is the fast one — the sharp front-of-palate arrival, the flash. Dihydrocapsaicin is the long one — the slow, deep, back-of-throat dwell that outlasts the flash by minutes. Nature runs them at roughly two parts capsaicin to one part dihydrocapsaicin, and that ratio, set by the tail-building enzymes, is why a wild chili burns the way it does and not some other way.

This is the part of the chemistry we find most beautiful, and it is worth slowing down for, because it is where flavor stops being a quantity and becomes a *composition*. When we tune the fatty-acid module — when we bias KAS toward the unsaturated chain or away from it — we are not making a pepper hotter or milder. We are reshaping the *kinetics* of the heat, moving the burn earlier or later in time, making it a flash or a dwell or some designed curve between them. Two cultivars can carry identical total Scoville and feel like entirely different experiences, because one front-loads and one lingers. The flagship ladder we are about to meet is not only a ladder of *how much*. It is, underneath, a ladder of *what shape*.

Capsaicin and dihydrocapsaicin are only the two loudest members of a small family, and naming the rest is worth a sentence, because the family's composition is itself a design surface. Alongside the two

majors sit the minor capsaicinoids — nordihydrocapsaicin, homocapsaicin, homodihydrocapsaicin — each a slight variation on the acyl tail's length and branching, each adding its own faint timbre to the burn. In a wild chili the blend runs roughly to two-thirds capsaicin and one-fifth dihydrocapsaicin, with the minors filling the remainder — a ratio the plant never chose consciously and that we now tune on purpose. Because every member of the family is built by the same tail-assembly enzymes — BCAT starting the branch, KAS elongating it, a desaturase placing or omitting that single sixth-carbon double bond — shifting the enzymes' balance shifts the whole chord. Bias toward the unsaturated tail and the burn moves forward and sharpens; bias toward the saturated and it deepens and lingers. This is why two of our cultivars can post identical Scoville and still feel like different animals in the mouth: Scoville counts the total capsaicinoid, but the *experience* is written in the ratio, and the ratio is written by three enzymes we can turn.

The bell pepper deserves one more look here too, because its genetics are the cleanest proof that this terminal module is truly separable. Sweetness in a bell pepper is not the presence of a sweetener — it is the *absence* of fire, and that absence is a broken gene, usually a recessive loss-of-function allele at Pun1 (the AT3 acyltransferase that performs the final condensation) and often at pAMT as well. The plant builds the entire universal trunk, arrives at vanillin, and then — its last switch flipped off by a single disabling mutation — simply stops, carrying a complete, silent capsaicin pathway it can no longer finish. Thousands of years of gardeners selected for exactly this break without knowing what they were breaking. What they were doing, in the language of this book, was deleting the terminal module. What we do is the same operation run in reverse, in a plant that never carried the module at all: not repairing a broken switch, but installing one where evolution left the socket empty.

And here is why this chapter sits so early in a book called *Beyond Peppers*: the head-and-tail architecture is the template for every transfer in the roadmap. To give a tomato a measured warmth, we do not reinvent anything. We install the head (pAMT, working on the vanillin the tomato's own trunk already delivers), we install the tail (BCAT and KAS, which the plant can feed from its own branched-chain amino acid pool), and we install the join (Pun1/CS). Four genes, moved as a cassette, onto a road the recipient already maintains. Every organism chapter in Part Three is, at bottom, a variation on this one sentence — *which four, in what balance, fed from what pool, joined by which synthase* — and the difficulty of each organism is almost entirely a question of how hard it is to persuade that particular plant to accept the cassette and grow up around it.

The four last enzymes are the pepper's private invention. They are also the most generous gift in biology, because they are so few that they can be given away.

Chapter 4 · The portable platform — *five of seven modules move.*

A method is only as transferable as its least transferable part, so the honest question behind this entire book is not "is capsaicin portable" — the chemistry of the last chapter answered that — but "is our *way of designing* it portable, or does it secretly depend on the pepper at every turn." We took that question apart module by module, and the answer is specific, defensible, and the organizing fact of the roadmap: of the seven modules that make up the platform, **five transfer to any sequenced organism without modification, and two must be re-solved for each new species.** Everything about how hard a given plant will be lives in those two.

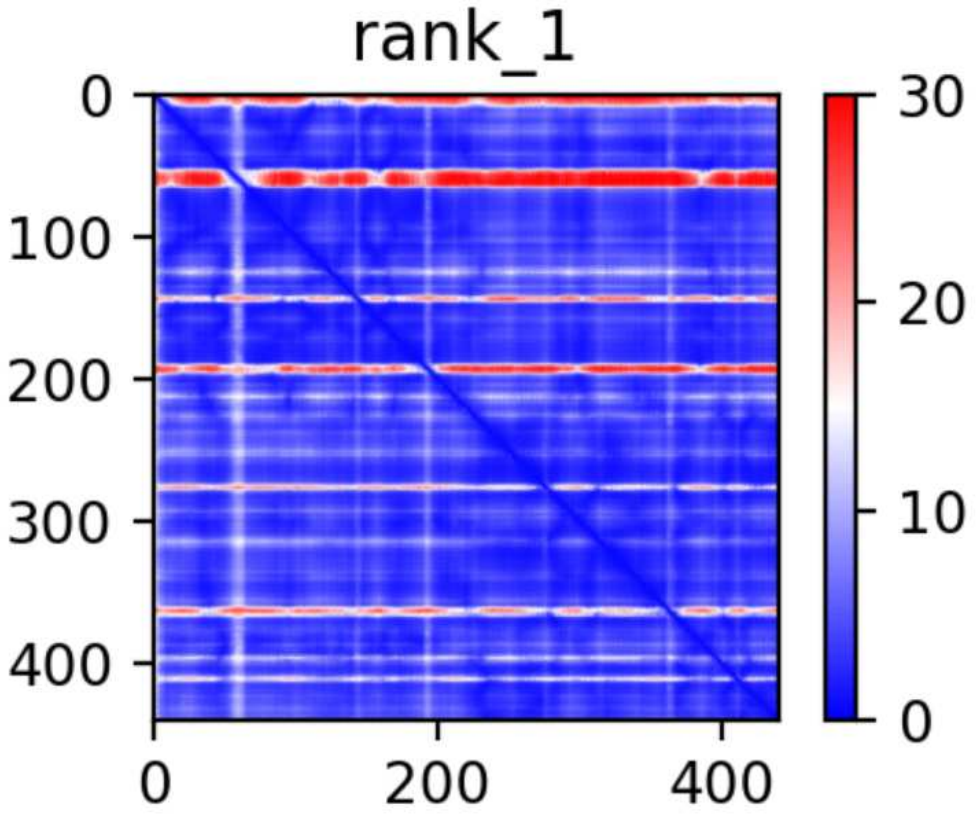

FIG. 2 — The predicted aligned error map for the terminal synthase. The computational modules read a protein like this one the same way whether it sits in a pepper or anywhere else; the sequence is the sequence.

Here is the platform, laid open, with the transfer verdict on each module:

#	MODULE	WHAT IT DOES	TRANSFERS UNCHANGED?
1	**Pathway cartography**	reconstruct the biosynthetic route from the genome	**Yes** — any sequenced genome
2	**Variant scoring (ESM2)**	rank every single-residue substitution by catalytic efficiency × fold stability	**Yes** — sequence is sequence
3	**Flux-balance analysis**	locate the true rate-limiting step in a genome-scale metabolic model	**Yes** — any reconstructed network
4	**Edit design + CRISPR**	write the chosen edits into the genome	**Yes** — the chemistry of the write is general
5	**The IFI measurement spine**	score the result at the receptor via EC_{50}	**Yes** — the tongue is the same tongue
6	**Transformation & regeneration**	get the edits into a whole, fertile plant	**No** — re-solve per species
7	**Agronomy + CV control**	grow it reproducibly to ≤7% batch CV	**No** — re-solve per species

Look at which five move. They are, every one, the parts that touch *information* rather than *tissue*. Module 1 reads a genome; a genome is a text, and our cartography does not care whose text it is. Module 2 is a 650-million-parameter protein language model, ESM2, that learned the grammar of proteins from tens of millions of sequences across all of life; when it scores a substitution in a synthase, it is reasoning about *protein*, not about *pepper*, and it reasons exactly as well about an enzyme from a strawberry. Module 3, flux-balance analysis, solves a linear program over a metabolic network to find the one step where carbon is actually bottlenecked — and a network is a network whether it was reconstructed from a chili or a carrot. Module 4 writes; CRISPR's chemistry is indifferent to the host. Module 5 measures at the human receptor, and the human receptor is a human universal — a TRPV1 channel does not know or care which plant delivered the capsaicin that opened it. Five modules, all informational, all host-blind. This is why the computational core of the method cost us years to build once and costs almost nothing to point at the next organism. We did not build a pepper tool. We built a tool that happens to have been aimed at peppers first.

Now look at which two do not move, because they are the entire texture of the roadmap. Module 6 — transformation and regeneration — is the stubbornly biological problem of getting an edited cell to accept its new instructions and then grow back into a whole, fertile plant. Some species take it gladly; a tomato or a potato will regenerate from a scrap of leaf almost eagerly. Others resist with everything they have; the grasses — corn, rice, sorghum — guard their genomes behind tissue that does not want to become a new plant, and coaxing them is a craft measured in years. Module 7 — agronomy and the control of variation — is the equally biological problem of growing the finished plant so reproducibly that batch after batch lands inside our ≤7% coefficient of variation, the number that separates a laboratory curiosity from something a formulator can build a product on. These two modules are where the difficulty tiers of the next part come from. They are not computational. No model shortens them. They are the price of working with living things, and we pay it per species, in seasons.

This is the cleanest way we know to say what the platform is: **the design is portable and the biology is local.** The part of the work that is hard to invent — reading the pathway, scoring the mutations, finding the bottleneck, writing the edits, measuring the loudness — we invented once and carry everywhere for free. The part that is hard to do — persuading a specific plant to accept the edit and grow up uniform around it — we re-earn, honestly, organism by organism. A roadmap built on that division of labor is not a wish list. It is a schedule. And because the expensive thinking is already done, the schedule is set almost entirely by how agreeable each plant decides to be — which is exactly the thing the difficulty ladder in Part Two was built to predict.

It is worth watching the computational core do its work once, concretely, on the pepper, because the worked example is what we carry to every organism afterward. When we set out to lift the pepper's heat past where ten thousand years of breeding had ever taken it, Module 2 — the ESM2 650-million-parameter protein language model — scored every possible single-residue substitution in Pun1, the capsaicin synthase, ranking each by its predicted effect on catalytic efficiency multiplied by fold stability. A protein language model does this the way a fluent reader senses a wrong word in a sentence: having absorbed the grammar of tens of millions of natural proteins across all of life, it can tell which substitutions a functioning enzyme would tolerate and which would break it, before a single drop of anything is pipetted. In parallel, Module 3 — flux-balance analysis over a genome-scale metabolic model — predicted that flux control sits upstream in the phenylpropanoid pathway at 4CL and HCT, not at the terminal Pun1 step; Pun1 labels the product, it does not govern the rate. The ESM2 optimization targeted Pun1 not to shift the heat ceiling but to improve the enzyme's catalytic efficiency and fold stability — turning wild-type variation into the ≤7% consistency a formulator can build on. The two modules together answer the two questions that decide every project: *which edit*, and *what is the edit actually improving.*

two Pun1 substitutions emerged from that search that no breeding program had ever found, because no breeding program could see the genotype well enough to look:

EDIT	EFFECT	ΔΔG (KCAL/MOL)
S39L	binding affinity ↑ at the N-terminus	−2.1
L345G	steric relief in the active-site cleft	−1.8
C175S	prevents a disulfide misfold *(structural annotation — not in shipped construct)*	−0.9

Read the ΔΔG column as the stability ledger it is: each number is the predicted change in folding free energy, and each one is negative, which means every edit leaves the folded enzyme *more* stable than before, not less. The model did not merely find substitutions that work; it found substitutions that are structurally safe — the combination conventional breeding can never aim at, because breeding selects on the fruit and cannot see a protein's fold at all. Written into the plant with Module 4's CRISPR edits, the two Pun1 substitutions (S39L · L345G) together improve catalytic efficiency

and consistency, tightening the batch-to-batch coefficient of variation from a wild-type **42% down to at or under 7%** and enabling our Phoenix flagship's predicted peak output of approximately **2.95 million SHU**. Breeding chased the peak and let the variance fall where it may; the method raised the peak and collapsed the variance at the same time, because an enzyme designed for fold stability behaves the same way every time it folds, plant after plant after plant.

This is the entire platform in one worked example, and every clause of it is portable. The ESM2 scoring that ranked Pun1's substitutions will rank any synthase's; the flux model that found the pepper's bottleneck will find a tomato's or a carrot's; the CRISPR write is host-blind; and the ≤7% discipline travels as a standard, not a rediscovery. Change the enzyme, change the organism, and the five informational modules run again, unchanged. S39L and L345G are the two Pun1 substitutions in the shipped construct — specific to one pepper's one synthase. The method that found them is specific to nothing.

Before we can read that ladder, though, we need the one instrument that makes every organism in it commensurable — the single number that lets a tomato, a chili, a strawberry, and a grain of rice be scored on the same page. That number is the subject of the next chapter, and it is the quiet hero of the whole enterprise.

Chapter 5 · The arithmetic of loudness — *one number for any plant.*

Taste, for all of recorded history, resisted the one thing science most wants to do to a phenomenon: put a number on it. We could measure how much of a compound a food contained — grams of sugar, parts-per-million of capsaicinoid — but content is not sensation. A food does not taste like the quantity of a molecule it holds; it tastes like how *loudly* that molecule speaks to the receptor it fits. Two tomatoes with identical glutamate can taste a world apart, because the matrix around the glutamate — the acid, the salt, the structure — changes how much of it actually reaches the receptor. Content is in the fruit. Loudness is at the tongue. The Iyengar Flavor Index measures loudness, and it is the instrument that makes *beyond peppers* a measurable program rather than a slogan.



FIG. 3 — The stage where loudness is read. Every axis on the Index anchors to a receptor's half-activation, measured here before any claim about a plant is allowed to stand.

The definition is one line, and the whole library resolves against it:

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

EC_{50} is the concentration at which a stimulus drives its receptor to half of its maximal response — the pharmacologist's measure of potency, inverted here so that *lower* concentration (a louder molecule, one that half-activates with less of itself) yields a *higher* score. Fix a reference, divide, scale by a thousand, and loudness becomes an integer. A score of 1,000 means "as loud as the reference." A score of 2,000 means "half the concentration does the

same job — twice as loud." For scorecards across the library we cap each axis at a normalization maximum of **5,000**, which is the ceiling against which every radar plot in every book is drawn.

The receptor-level anchoring is what makes the number universal across organisms, and this is the point that matters most for a book about leaving the pepper behind. The Index is not anchored to the pepper. It is anchored to the *tongue* — to eight human receptors that do not change from food to food:

CODE	AXIS	RECEPTOR	WHAT IT READS
IFI-H	Heat	TRPV1	capsaicin's warmth, the trigeminal fire
IFI-S	Sweetness	TAS1R2 · TAS1R3	sugar-free sweetness
IFI-U	Umami	mGluR1 · TAS1R1·R3	savory, brothy depth
IFI-B	Zero-Bitter	TAS2R38	bitterness attenuation (20–65%)
IFI-K	Kokumi	CaSR	richness, continuity, mouthfeel
IFI-C	Cool-Burn	TRPM8 · TRPV1	cooling and warming
IFI-P	Pungency	TRPA1	the sharp channel — mustard, wasabi, allicin
IFI-A	Astringency	salivary PRPs	tannin and polyphenol grip

Because the anchor is the receptor and not the source plant, a heat score of 2,600 means exactly the same thing whether it comes from a chili, a tomato engineered to warm, or a strawberry on the roadmap. The Index is the common currency. It is the reason the forty-two live organisms and the sixty-four concepts can be laid on one page and compared without translation — they are all being asked the same question, in the same units, by the same tongue.

Now the part we love, where the arithmetic tells us something the poetry could not. Take our flagship heat ladder — ten superhot cultivars, each a measured rung — and set each one's Scoville (a *content* measure, parts of capsaicinoid) against its IFI-H (a *loudness* measure, receptor potency):

CULTIVAR	SKU	SHU	IFI-H	AXES
Ember	SS-0087	866,000	1,570	2/8
Cinder	SS-0088	1,050,000	2,290	3/8
Kindle	SS-0089	1,250,000	2,000	3/8
Flare	SS-0090	1,480,000	—	4/8
Blaze	SS-0091	1,720,000	2,630	4/8
Wildfire	SS-0092	1,960,000	2,930	5/8
Scorch	SS-0093	2,210,000	3,220	5/8
Firestorm	SS-0094	2,470,000	3,490	6/8
Solara	SS-0095	2,720,000	3,980	6/8
Phoenix	**SS-0096**	**2,937,948**	**4,200**	**8/8**

Watch the two columns climb at different rates. From Ember to Phoenix, Scoville rises by a factor of 3.39 — but IFI-H rises by a factor of only 2.68. Loudness grows *slower* than content, and it does so with a regularity you can fit: across this ladder, **IFI-H ≈ IFI_0 × (SHU)^0.8** — a power law with an exponent near four-fifths. Hold that exponent up to the light and it tells a true story about the body. TRPV1 is a receptor, and receptors saturate; as capsaicin content climbs, more and more of it is arriving at a channel population that is already most of the way open, so each additional part of capsaicinoid buys less additional loudness than the last. The sublinearity is not a flaw in the Index. It is the Index faithfully reporting the diminishing returns of the receptor — the arithmetic of a tongue approaching the limit of what it can feel. Content can run to infinity. Sensation runs to a ceiling, and the exponent 0.8 is the shape of that ceiling drawn in our own data.

That 8/8 in Phoenix's row deserves its own moment, because it is the single most important data point in this book. Every other flagship on the ladder is a heat specialist — two, three, at most six axes expressed. Phoenix is the first cultivar we ever engineered to register on all eight axes at once, and its full profile is the proof, inside one pepper, that a single organism can be made to speak the entire language instead of shouting one word of it:

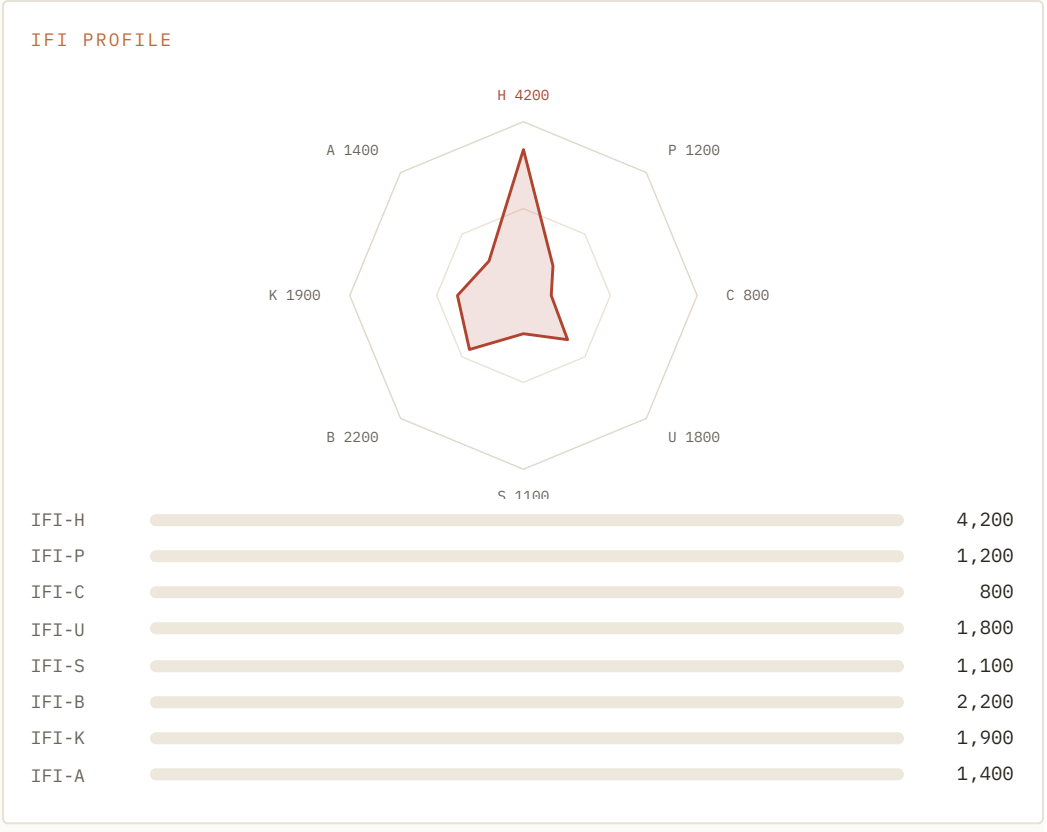

Read down that scorecard and you are reading the thesis of the whole library in miniature. Heat leads at 4,200, as it must in a superhot — but pungency (TRPA1) registers at 1,200, bitterness-attenuation (TAS2R38) at 2,200, kokumi (CaSR) at 1,900, umami (mGluR1) at 1,800, astringency (salivary PRPs) at 1,400, sweetness (TAS1R2·R3) at 1,100, and cool-burn (TRPM8) at 800 — every one of them a real, measured activation of a real, separate receptor. Phoenix is not a hot pepper with faint echoes of other tastes bleeding in. It is a deliberately full-spectrum organism that happens

to lead with heat, built by running the loop of this chapter eight separate times, once per axis, in a single plant. And that is exactly why it is the hinge of this book rather than merely its hottest page: if the method can carry all eight axes *inside* one pepper, at designed and measured strength, then the claim that it can carry each axis *out* of the pepper, into the forty-two live organisms and the sixty-four on the roadmap, is not a leap of faith. It is the same demonstrated capability, pointed outward.

One last reading, because it anchors the number to the molecule. Phoenix's IFI-H of 4,200 inverts, through the definition, to an effective EC_{50} of 1,000 ÷ 4,200 ≈ 0.238 of the reference — its capsaicin half-opens TRPV1 at under a quarter of the reference concentration, which is what "4.2 times as loud" means when you cash it out in molarity. And 4,200 sits at 84% of the 5,000 scorecard ceiling: near the top of what the axis can report, with deliberate headroom left above it, because a scale that cannot be approached is a scale no one trusts. This is the instrument the rest of the book runs on. Every organism we meet, live or imagined, will be asked to post a number on it — and because the number is anchored to the tongue and not the plant, every one of those numbers will mean the same thing.



PART II OF IV

The Map We Can Already Read

Forty-two organisms live, nineteen thousand constructs, and the five-tier ladder of difficulty that tells us how hard a new plant will be before we lift a pipette

Chapter 6 · Forty-two organisms, nineteen thousand constructs — *the library that already exists.*

It is easy, in a book about what could be built, to speak as though everything lies ahead. It does not. The thesis of *beyond peppers* is not a prediction; it is a thing we have already done forty-two times, in forty-two different plants, and the evidence is a living library of roughly nineteen thousand validated constructs that we can walk you through row by row. Before we draw the roadmap, we owe you the territory already mapped — because the roadmap's credibility rests entirely on the fact that the method has already left the pepper, repeatedly, and come back with a number every time.



The living library at golden hour — forty-two species under one roof, each carrying its own answer to the same four questions.

Here is the whole live library, sorted not alphabetically but by the thing that actually governs the work — **engineering difficulty**, the five-tier ladder from *Native* to *Frontier* that the next chapter will teach you to read. Each organism carries its validated construct count and the primary axes it has been tuned to speak:

DIFFICULTY	ORGANISM	CONSTRUCTS	PRIMARY AXES
Native	Pepper	1,060	all
Easy	Tomato	2,040	S · U · F · M
Easy	Potato	80	B
Moderate	Stevia	910	S · F · M
Moderate	Garlic	560	K · P
Moderate	Broccoli	380	B · M
Moderate	Blueberry	120	A
Moderate	Ginger	120	P
Moderate	Hops	80	B
Moderate	Cabbage	40	B
Moderate	Kale	40	B
Moderate	Onion	40	K
Moderate	Raspberry	40	A
Hard	Habanero	3,190	F · M
Hard	Black Mustard	2,790	P · M
Hard	Peppermint	2,360	C · M
Hard	Grape	1,610	B · A · M
Hard	Tea	280	U · B · A
Hard	Broccoli-Kale	280	B
Hard	Yeast	240	U · K
Hard	White Mustard	200	P
Hard	Lavender	160	C
Hard	Orange	120	B · C
Hard	Pomegranate	120	A
Hard	Thyme	120	P
Hard	Watercress	120	P
Hard	Koji Mold	80	U
Hard	Monk Fruit	80	S
Hard	Eucalyptus	80	C
Hard	Shiitake	80	U
Hard	Spearmint	80	C
Hard	Coffee	40	B
Hard	Camphor	40	C
Hard	Cinnamon	40	P
Hard	Mushroom	40	U
Hard	Oregano	40	P
Hard	Radish	40	B
Hard	Rapeseed	40	B
Hard	Vanilla	40	K
Frontier	Soybean	360	U · K
Frontier	Sugarcane	200	S
Frontier	Sorghum	120	A

Add those columns and they come to **18,500 validated constructs across the forty-two live organisms** — the figure that sits inside our canonical ~~19,000, with the balance in the dual- and triple-axis libraries we will meet in Part Four. We hold to one master number~~

~~system across the whole library, because across sixty-four books the one thing that reads as sloppy is a count that drifts: **~~19,000 constructs, 1,000 live and purchasable today, 78,800 total SKUs, 42 organisms, 8 axes, ≤7% batch CV.** Every page in every volume pulls from that single set.

Read the table for its shape, not just its numbers, because the shape is the argument. Notice first that the heaviest construct counts are *not* all in the easy organisms — Habanero at 3,190 and Black Mustard at 2,790 are both rated Hard, and they carry more constructs than the Native pepper itself. Difficulty and investment are different axes. We build deepest where the flavor payoff is richest, not where the biology is gentlest, and some of our most valuable libraries sit in plants that fought us. Notice second how the axes cluster by family: the Brassicas — broccoli, cabbage, kale, radish, rapeseed, watercress — gather around *Zero-Bitter* and *Pungency*, because their shared glucosinolate chemistry is both their bitterness and their sharpness, two axes on one molecular family. The mints — peppermint, spearmint, eucalyptus, camphor, lavender — gather around *Cool-Burn*, because they share the terpene route to menthol and its cousins. The fungi and the ferments — yeast, koji, shiitake, mushroom — gather around *Umami* and *Kokumi*. The Index did not impose this order. The plants' own biochemistry did, and the Index merely made it visible.

The construct-count column tells its own story if you let it. The distribution is steep — a handful of organisms (Habanero at 3,190, Black Mustard at 2,790, Peppermint at 2,360, Tomato at 2,040, Grape at 1,610, Pepper at 1,060) carry the great majority of the library's weight, while a long tail of forty-construct organisms — coffee, camphor, cinnamon, mushroom, oregano, radish, rapeseed, vanilla — each hold a small, deliberate foothold. This is not neglect of the tail; it is the signature of a platform doing exactly what a platform should. The heavy organisms are the ones where we have built *deep* — many cultivars, many axes, the full exploration of what the plant can be made to do — and the light ones are the ones where we have built *first*, a validated beachhead that proves the method reached the species and can be deepened whenever the flavor payoff justifies it. A forty-construct organism is not a small result. It is a planted flag: proof that coffee, or vanilla, or rapeseed accepted the method and posted a number, with the depth to follow when it is wanted. Read together, the heavy heads and the long tail are the two things a transferable method is supposed to produce at once — depth where it pays, and reach everywhere else — and the fact that one column shows both is the clearest single-page evidence that the platform is real. One source of truth sits under all of it: the canonical number system every book and page resolves against, so that the forty in coffee's row and the 3,190 in Habanero's mean the same kind of thing, counted the same way, on every surface the library touches.

And notice, finally, where the Frontier organisms sit: soybean, sugarcane, sorghum — a legume and two grasses. They are the vanguard of the hardest transformation problems, the plants whose tissue most resists Module 6, and their presence in the *live* library at

all is the quiet proof that the two non-portable modules, hard as they are, do yield. We did not list them on a roadmap. We grew them. A grass that fought regeneration for seasons is carrying constructs today. That is the sentence we most want a skeptic to sit with before turning to the roadmap: the difficult tiers are not theoretical. We have already climbed into the first of them, and come down with plants.

Chapter 7 · Reading difficulty before we build — *the ladder, and the arithmetic under it.*

The most valuable thing a mature method gives you is not the ability to do a hard thing. It is the ability to *know, before you start, how hard the thing will be* — to put a cost and a calendar on a plant you have never touched, from its genome and its family alone. We can do this now with enough reliability to schedule years of work around it, and the instrument is the five-tier difficulty ladder you just saw the library sorted by. This chapter teaches you to read it, because it is the lens through which every organism in the roadmap was placed.



The controlled chamber, where a difficulty estimate becomes a season of actual work. Everything the ladder predicts is finally settled here, in tissue.

The five tiers are not arbitrary bands. Each one corresponds to a real, nameable obstacle in Module 6 (transformation and regeneration) or Module 7 (agronomy and CV control) — the two modules that, you will recall, do *not* transfer for free:

TIER	WHAT MAKES IT THIS HARD	REPRESENTATIVE LIVE ORGANISMS
Native	capsaicin is already here; we tune, not install	Pepper
Easy	regenerates readily; edits take and hold; CV lands fast	Tomato, Potato
Moderate	transformation works but wants optimization; CV needs a season or two	Garlic, Broccoli, Blueberry, Ginger
Hard	recalcitrant tissue, long regeneration, or a flavor module with many parts	Habanero, Grape, Tea, Coffee, Vanilla
Frontier	monocot or legume tissue that resists regeneration; CV hard-won	Soybean, Sugarcane, Sorghum

The ladder tracks biology, not chemistry — and that distinction is the single most counterintuitive thing we have learned. You would expect the hardest plants to be the ones where the *flavor* is hardest to build. Sometimes they are. But far more often, the chemistry is the easy part — the computational modules have already solved it on paper — and the difficulty is entirely in the tissue. A tomato is *Easy* not because tomato umami is simple (it is not) but because a tomato

will regrow a whole plant from a leaf disc almost without being asked. A grain of rice is *Frontier* not because rice flavor chemistry is exotic but because a rice cell, edited, would rather die than become a seedling. **The difficulty of a plant is mostly the difficulty of persuading it to grow back.** This is why Part One spent a whole chapter on which modules transfer: the five that move make the *flavor* nearly free, which throws the entire cost of a new organism onto the two that don't — and the ladder is simply our map of how those two misbehave, family by family.

Under the tier labels sits real arithmetic, and we use three estimators to place a new plant before we commit a season to it.

The first is **phylogenetic distance from a solved organism.** Because transformation protocols transfer best between close relatives, a plant's difficulty is well predicted by how far it sits from the nearest organism we have already cracked. A tomatillo is close kin to the tomato, so it inherits most of the tomato's gentleness; a grain of corn sits across the vast monocot–dicot divide from everything we had solved early, and inherits none of it. This is why the roadmap's tiers track taxonomy so closely — Tier 1 is the tomato's own family, and each tier out is a wider ring of cousins.

The second is **construct-count scaling with flavor complexity.** A single-axis tune (restore one plant's umami) is a handful of edits; a multi-axis build is many. Our construct counts are a rough ledger of this: the single-axis Frontier plants carry dozens to low hundreds of constructs, while a Habanero asked to express the full spectrum carries over three thousand. When we estimate a new organism, we estimate the construct count first — how many edits the target profile will demand — because that number drives both the design effort and the number of regeneration events we will need to survive.

The third, and the one that most often decides whether a project is a year or a decade, is **flux-balance analysis of the rate-limiting step.** Before we design a single edit, we reconstruct the organism's genome-scale metabolic network and solve for the one reaction where carbon is actually bottlenecked on the way to the target molecule. This matters enormously, because the naïve instinct — overexpress the final enzyme, the glamorous one, the synthase — is almost always wrong. The final enzyme is rarely the bottleneck. Far upstream, some unglamorous precursor step is throttling the whole road, and pouring effort into the synthase while the real limit sits five reactions back buys you nothing. On the pepper, flux-balance analysis told us the capsaicinoid pathway itself — not an upstream supply step — was the true limiter, which is why editing the terminal synthase paid off. On a new organism, the same analysis might point anywhere, and it tells us where to spend. A plant whose bottleneck is a single, editable enzyme is tractable; a plant whose bottleneck is a diffuse supply problem across a dozen reactions is a long campaign. The flux model is how we tell those two apart on a workstation, in an afternoon, before anyone has planted anything.

A worked instance makes the logic concrete. Suppose we are weighing a new berry for the roadmap. The phylogenetic estimator places it two families from our nearest solved fruit, which predicts a

transformation protocol that will transfer in part but want a season of optimization — call it the back half of *Hard*. The construct-count estimator, reading the target profile we have in mind (restore one axis, ladder a second), forecasts a build in the low hundreds of edits, not the low thousands, which keeps the regeneration burden survivable. And the flux estimator, run against the berry's reconstructed metabolic network, returns the decisive fact: the bottleneck on the way to the target molecule is a single editable enzyme, not a diffuse supply problem smeared across a dozen reactions. Three readings, three green-enough lights, and the berry earns a tier, a gene-count range, and a timeline — before a single seed is sown. Now flip one input. Leave everything else the same but let the flux model return a *diffuse* bottleneck — carbon throttled not at one step but across many, each contributing a little — and the same berry's estimate lengthens from a project into a campaign, because no single edit relieves a distributed limit; you have to widen the whole road at once. This is the quiet power the method buys: the difference between a one-enzyme fix and a whole-pathway rebuild is visible in the flux solution, on a workstation, long before it would become visible in a greenhouse, in a lost year. We placed all sixty-four roadmap organisms this way. None of the placements is a guess. Every one is an estimate with three numbers behind it, and the honesty of the roadmap is exactly the honesty of those three numbers.

Put the three estimators together and you get a difficulty *estimate* with a tier attached, a construct-count forecast, and a named bottleneck — a schedule, in other words, derived from a genome and a family tree. That estimate is what let us draw a roadmap of sixty-four organisms with calendars attached and mean them. Turn the page and the roadmap begins, and every placement in it — every "4–5 genes, 18–24 months" and every "10–12 genes, 48–72 months" — is one of these estimates, made the way this chapter describes. The ladder is not a ranking of ambition. It is a forecast, and we built the roadmap by trusting it.



PART III OF IV

The Roadmap

Sixty-four concept organisms in five tiers — from the Solanaceae core to the synthetic frontier, and the undrawn edge beyond it

Chapter 8 · Tier 1 · The Solanaceae core — *the pepper's own family answers first.*

Every expansion begins with the nearest kin, and the pepper's nearest kin is the one that made us believe in this book at all. The nightshades — Solanaceae — are the pepper's own family, and they inherit, almost as a birthright, most of what makes the pepper tractable: gentle tissue, willing regeneration, a metabolic network close enough to the pepper's that our flux models barely have to re-learn it. Tier 1 is therefore the roadmap's shortest reach and its highest-confidence promise: **four to five genes, eighteen to twenty-four months**, and a profile we can forecast with real precision because the recipient plant is the chili's cousin.



Tier 1 begins at home. The tomato — the pepper's close kin — regenerates so willingly that the whole difficulty of the roadmap's first rung is flavor, not tissue.

Here is the tier, with its concept organisms and their named cultivar series:

CONCEPT	BASE PLANT	SERIES	~GENES	~TIMELINE
Volcanic	Tomato	Ember · Magma · Pyroclast · Caldera · Krakatoa	4–5	18–24 mo
Dark Matter	Eggplant	Smolder · Nightshade · Eclipse · Obsidian · Event Horizon	4–5	18–24 mo
Laboratory	Tobacco	Proof of Concept · Benchmark · Threshold	4–5	18–24 mo

Start with the **Volcanic tomato**, because it is the single most important organism in the entire roadmap and the one we would build first with a gun to our heads. A tomato already carries, natively, two of the eight axes at real strength — umami and sweetness — and its flesh is the most forgiving canvas in the nightshade family. To make it warm, we install the four-enzyme cassette of Chapter 3 — head, tail, and join — onto the vanillin precursor the tomato's own universal trunk already delivers. The result is not a novelty. It is a tomato that carries a *designed* warmth laddered across five named cultivars, from the first blush of Ember to the genuine heat of Krakatoa, with the tomato's native umami and sweetness still underneath — a fruit that is savory, sweet, and warm at once, which is a combination the pepper can only approximate and the tomato can hold as if it were born to it. The five-cultivar ladder matters: it means Volcanic is not one product but a dialable range, the same way our superhot flagships are a range, so that a sauce-maker can choose their rung.

It is worth naming those rungs, because they are the Tier-1 design philosophy made legible. Ember is the first blush — a warmth you notice more as a glow than a heat, a tomato that has clearly been somewhere warm. Magma and Pyroclast climb into genuine, conversational heat, the register of a good arrabbiata built into the fruit itself rather than stirred in after. Caldera and Krakatoa reach real fire, a tomato that bites back, for the cook who wants the heat to be the point of the dish. Across all five, the native umami and sweetness hold steady underneath — that steadiness is the non-negotiable of the design — so the ladder is never heat *replacing* the tomato but heat *added* to a tomato that stays fully itself. The Dark Matter eggplant runs its own five-rung ladder on a different ground entirely: Smolder and Nightshade build heat into the eggplant's native dusk; Eclipse and Obsidian deepen it; Event Horizon carries it to the edge, a vegetable so dark and warm at once that it barely reads as a vegetable. The two ladders together teach the tier's central lesson better than any single cultivar could: the *same four enzymes* produce bright heat in the tomato and smoldering heat in the eggplant, not because the cassette differs between them but because the ground does — and a well-designed heat is always heat shaped by the plant it is grown into. We do not impose a flavor on the nightshades. We read what each one already is, and give it the fire it would have chosen for itself.

The **Dark Matter eggplant** is the tier's most beautiful problem. Eggplant is already dark — anthocyanin-saturated, built around the deep bitter-savory register that makes it the meatiest vegetable in

the garden — and its series names (Smolder, Eclipse, Obsidian, Event Horizon) are not whimsy; they track a real intention to build heat *into* that darkness rather than on top of it. Where the Volcanic tomato is bright heat on a sweet ground, the Dark Matter eggplant is designed to be smoldering heat on a savory-bitter ground, a far moodier flavor, and one that leans on the organism's native character instead of fighting it. This is the design principle that governs all of Tier 1 and most of the roadmap: *build with the plant's grain, not against it.* The tomato gets bright heat because the tomato is bright. The eggplant gets smoldering heat because the eggplant is dark. The method is the same four enzymes; the art is reading what the plant already is and giving it the heat it would have chosen.

The **Laboratory tobacco** is the odd, essential one — and its series names say exactly why: Proof of Concept, Benchmark, Threshold. Tobacco is not a crop we intend to feed anyone. It is the plant kingdom's laboratory workhorse, the fastest-regenerating, most-transformable, best-understood model plant in all of botany, and it exists on this roadmap as an *instrument*, not a product. When we want to test whether a new version of the cassette works — a retuned synthase, a rebalanced tail — before committing eighteen months of a food crop to it, we run it in tobacco first, where a generation is weeks and the tissue never argues. Laboratory tobacco is where every other organism's edits are proved cheaply before they are grown expensively. Naming its cultivars Proof of Concept and Benchmark is simply honesty: this plant's entire job is to be the benchmark.

What unites Tier 1 is confidence. These are the organisms where our difficulty estimators all point the same gentle direction — close kin to the solved pepper, willing tissue, a short construct count, a bottleneck the flux model can name and reach. If the thesis of this book is wrong, it is wrong somewhere out in Tier 4 among the grasses, where the biology turns hostile. It is not wrong here. Here, among the pepper's own family, *beyond peppers* is not a leap of faith. It is next season's planting.

Chapter 9 · Tier 2 · The extended nightshades — *the family, one ring out.*

Widen the circle by one ring and you leave the pepper's immediate cousins for its more distant relations — still Solanaceae, still family, but far enough out that the gentleness is no longer a birthright and must be partly re-earned. Tier 2 is the roadmap's first real test of the difficulty estimators, because these are plants where the phylogenetic distance has grown enough to matter: **five to seven genes, twenty-four to thirty-six months**, and a transformation protocol that transfers from the tomato's but no longer drops in unchanged.



One ring out. The potato — nightshade kin, but a tuber, not a fruit — moves the design problem underground and asks the method to follow it there.

The tier holds four concepts, and each one bends the method toward a different organ or a different register:

CONCEPT	BASE PLANT	~GENES	~TIMELINE	THE TWIST
Seismic	Potato	5–7	24–36 mo	heat in a *tuber*, a storage organ, not a fruit
Celestial	Ground Cherry	5–7	24–36 mo	heat inside the papery husk's sweet-tart fruit
Acid	Tomatillo	5–7	24–36 mo	heat laddered against native tartness
Floral	Petunia	5–7	24–36 mo	the ornamental model — heat as a non-food proof

The **Seismic potato** is the tier's most instructive build, because it moves the entire design problem into an organ the roadmap has not yet addressed: the tuber. A tuber is a storage organ — a bank of starch the plant keeps underground for next season — and its metabolism runs on a different logic than a fruit's. A fruit is built to be eaten and is metabolically generous with its flavor molecules; a tuber is built to be hoarded and is metabolically stingy. Installing the capsaicin cassette in a potato therefore means feeding it from a storage metabolism rather than a ripening one, and our flux-balance analysis has to be redone from scratch against the tuber's own network, because the bottleneck sits in a different place when the

organ's whole purpose is conservation rather than display. Potato is rated *Easy* as a live organism for simple edits — it regenerates willingly — but *Seismic* is harder than the live potato's rating suggests, precisely because asking a hoarding organ to express a display molecule is a metabolic argument, not just a transformation one. This is the kind of subtlety Tier 2 is full of: the tissue is friendly, but the organ fights you.

The Celestial ground cherry poses the opposite and equally instructive problem — not a hoarding organ but a *protected* one. The ground cherry ripens inside a papery husk, a little lantern the plant builds around each fruit, and that enclosure changes the ripening environment in ways our models have to account for: gas exchange, light, and the timing of sugar accumulation all run differently inside the husk than on an open vine. The fruit itself is a small marvel of sweet-tart balance, a flavor with a cult following precisely because nothing else tastes quite like it, and laddering heat into it means laddering heat into a moving target — a sugar-acid balance that shifts as the husk does its slow work. Where the Seismic potato asks us to feed a display molecule from a hoarding metabolism, the Celestial ground cherry asks us to time an installed pathway to a fruit that ripens in its own private weather. Neither is a pepper problem; both are *organ* problems, and that is the real texture of Tier 2 — the tissue grows more familiar the further out we go, but each organ keeps its own secrets, and the method earns the tier one secret at a time.

The **Acid tomatillo** is the tier's cleanest demonstration of designing *against a native axis*. A tomatillo is defined by its tartness — a bright, green, structural acidity that is the whole reason it anchors a salsa verde. Heat laddered against that acid is a genuinely different flavor than heat on the tomato's sweetness: the acid sharpens the heat's arrival and shortens its dwell, so a Volcanic tomato and an Acid tomatillo at identical IFI-H would feel measurably different in the mouth — the tomatillo's heat arriving faster and leaving cleaner, cut by its own acid. This is the fatty-acid-kinetics lesson of Chapter 3 meeting the plant's native chemistry: the *shape* of a burn is set partly by the tail-building enzymes we install and partly by the acid and sugar matrix the plant brings. An Acid tomatillo lets the plant's own tartness do some of the shaping, which is exactly the build-with-the-grain principle again, now applied to a plant whose grain is sourness.

The **Floral petunia** sits here for the same reason the Laboratory tobacco sat in Tier 1: it is an instrument, not a dinner. Petunia is one of botany's great ornamental model organisms, exquisitely well-characterized, routinely transformed, and — crucially — not a food, which makes it a safe and fast proof surface for edits we are not yet ready to run in a crop. Heat in a petunia feeds no one, and that is the point: it proves the cassette functions one more ring out from the pepper, in a plant whose tissue behaves differently, without committing a food crop to an unproven design. The **Celestial ground cherry**, by contrast, is a genuine food target — that small, husked, sweet-tart fruit with a cult following — and it rounds the tier out as

the sweet-fruited counterpart to the Acid tomatillo's sourness. Together the four cover the tier's range: a tuber, a husked sweet fruit, a tart fruit, and an ornamental instrument — four different organs and registers, all still inside the nightshade family, all within three years. Tier 2 is where the roadmap stops being the pepper's immediate family reunion and becomes the first real demonstration that the method travels — not far yet, but measurably, repeatably, and with the difficulty estimators earning their keep at every step.

Chapter 10 · Tier 3 · Crossing the family line — *out of the nightshades at last.*

Here is the threshold the whole book has been walking toward: the edge of the family. Everything in Tiers 1 and 2 was Solanaceae — the pepper's own relations, inheriting its tractability by blood. Tier 3 steps across the family line entirely, into plants that share none of the nightshade's convenient tissue and none of its pre-solved transformation, and whose only connection to the pepper is the one that matters most and travels furthest: the universal trunk, running in them exactly as it runs in everything green. This is **seven to ten genes, thirty-six to forty-eight months**, and it is where *beyond peppers* becomes literally true — where the method leaves the family that birthed it and asks whether the thesis holds in a strawberry, a cucumber, a melon, an apple, a head of lettuce.



Across the line. A strawberry shares no family tissue with the pepper — only the universal trunk. Sweet-Heat is the roadmap's proof that the trunk is enough.

The tier's five concepts reach into five different families, and the flagship among them has four named cultivars of its own:

CONCEPT	BASE PLANT	FAMILY LEAP	SERIES / NOTE	~GENES
Sweet Heat	Strawberry	Rosaceae	Warm Kiss · Sugar Burn · Candied Fang · Devil's Dessert	7–10
Contradiction	Cucumber	Cucurbitaceae	cool-fleshed fruit made to carry heat	7–10
Solar	Melon	Cucurbitaceae	heat inside high-sugar flesh	7–10
Forbidden	Apple	Rosaceae	heat in a tree fruit — the long-generation problem	7–10
Espionage	Lettuce	Asteraceae	heat in a leaf crop, hidden in plain sight	7–10

The **Sweet-Heat strawberry** is the emotional center of the roadmap and the concept that, more than any other, made people outside the company understand what we were really building. A strawberry is a Rosaceae fruit — kin to roses and apples, utterly unrelated to the pepper — and it is beloved for one axis, sweetness, expressed with a floral top note almost nothing else in the garden can match. To ladder heat into it, from the first warmth of Warm Kiss to the genuine fire of Devil's Dessert, is to build a flavor the natural world has simply never produced: a berry that is sweet and floral on the front of the tongue and warm on the back, the sweetness and the heat arriving in sequence rather than in conflict. The cultivar names trace the ladder honestly — Warm Kiss is a suggestion, Sugar Burn a real contrast, Candied Fang a provocation, Devil's Dessert a dare. And the design logic is the same four enzymes it has always been, now installed in Rosaceae tissue, fed by the strawberry's own trunk, with the plant's native sweetness left fully intact underneath so that the heat reads as a *second* voice and not a replacement for the first. When we say the method is organism-agnostic, Sweet Heat is the sentence we are asking you to believe, and it is seven to ten edits of distance from today.

The **Contradiction cucumber** earns its name as a piece of deliberate design philosophy. A cucumber is, in the sensory imagination, the very emblem of cool — watery, green, calming, the thing you put on a burn, not the thing that burns. To build heat into it is to engineer a contradiction on purpose, and the reason to do it is precisely that the contradiction is interesting: a cucumber that cools the mouth with its water and warms it with its designed capsaicin at the same moment is a genuinely novel sensory event, a plant arguing with itself in a way a person finds delicious. The **Solar melon** is Contradiction's sweeter sibling in the same gourd family — heat laddered into high-sugar melon flesh, the cucurbit's sugar doing for Solar what the strawberry's does for Sweet Heat. Building both proves the method in Cucurbitaceae twice, on a cool fruit and a sweet one, which is how we gain confidence that a transfer is about the *family's* tissue and not a fluke of one lucky plant.

Crossing the family line is, in Module 6 terms, the most expensive thing Tier 3 does, and it is worth being plain about why. Every transformation protocol we hold was solved for a particular lineage; the nightshade protocols that carried us through Tiers 1 and 2 simply do not apply to a rose-family strawberry or a gourd-family melon, and so each family in Tier 3 is, in effect, a fresh transformation problem bolted onto an already-solved flavor problem. This is the exact inversion that makes the tier instructive: the *chemistry* is

portable — the trunk is in the strawberry and the cucumber and the apple precisely as it is in the chili, and our five computational modules do not so much as blink at a new family — but the *tissue* must be re-earned each time. When we build both a cucumber and a melon in Cucurbitaceae, we are not being thorough for its own sake; we are amortizing one transformation solution across two products, so that the season spent learning to regenerate a gourd pays off twice. The same logic gives Rosaceae both a strawberry and an apple, though the apple's tree-generation problem means the two run on wildly different clocks. Read the tier this way and its shape becomes a strategy rather than a list: pick a family, pay once to enter it through its most tractable member, and the second and third organisms in that family arrive far cheaper than the first. The strawberry is the door into Rosaceae; the cucumber, into Cucurbitaceae; the lettuce, into Asteraceae. Each door is expensive to open. Each room behind it holds more than one plant.

The **Forbidden apple** carries the tier's hardest biology, and its difficulty has nothing to do with flavor. An apple is a tree, and a tree's generation time is the enemy of iterative engineering: where a tomato gives you an answer in a season, an apple makes you wait years between question and reply, and every round of the design-build-test loop stretches to match the tree's patience. Forbidden is on the roadmap because an apple with designed heat is a genuinely startling object — the fruit of the oldest story in the West, made to bite back — but it is scheduled soberly, because no computational cleverness shortens the years a tree takes to fruit. The **Espionage lettuce** is the tier's quiet counterweight: heat in a leaf crop, a plant we usually grow for texture and coolness, hidden — as the name says — in plain sight, a salad green that turns out to carry a secret. Across these five, Tier 3 makes its single large claim: that the family line, which bounds so much of conventional breeding, is not a boundary for us at all, because the thing we rely on — the trunk — was never a family trait. It was a kingdom trait. Cross the family line and the trunk is still there, waiting, in the strawberry and the cucumber and the apple exactly as it waited in the chili. That is the whole of the thesis, demonstrated five times, one family at a time.

Chapter 11 · Tier 4 · The frontier synthetic — *the hardest jump in the garden.*

Now the biology turns genuinely hostile, and we will not pretend otherwise. Tier 4 is where the roadmap meets the monocots — the grasses and their kin, corn and rice and sorghum, the plants that feed most of humanity and resist genetic engineering more stubbornly than almost anything else we grow. The jump from the dicots of Tiers 1 through 3 to the monocots of Tier 4 is the single largest discontinuity on the entire roadmap, larger than any family line, because it is a split that runs back to the deep history of flowering plants themselves. This is **ten to twelve genes, forty-eight to seventy-two months** — four to six years — and the long timelines are not caution. They are respect for how hard these plants are to persuade.



The frontier. A grain of corn sits across the oldest divide in flowering plants from everything the method solved first; its tissue would rather die than regenerate, and that is the whole of Tier 4's difficulty.

The tier reaches across both the monocot divide and into the hardest dicots, and its concepts are named for weather and scale:

CONCEPT	BASE PLANT	CLADE	~GENES	~TIMELINE	THE HARD PART
Storm	Corn	monocot	10–12	48–72 mo	grass tissue resists regeneration
Dynasty	Rice	monocot	10–12	48–72 mo	staple-crop rigor; monocot recalcitrance
Wildfire	Broccoli	Brassica	10–12	48–72 mo	glucosinolate chemistry tangled with the target
Nuclear	Cabbage / Cauliflower	Brassica	10–12	48–72 mo	dense heads, slow lines
Deep Earth	Carrot	Apiaceae	10–12	48–72 mo	a root organ, a new metabolism again

Understand first *why* the grasses are the frontier, because it is the clearest illustration in the whole book of the Part One thesis that the design is portable and the biology is local. The flavor chemistry of **Storm corn** or **Dynasty rice** is, on paper, no harder than the strawberry's — the same universal trunk runs in a grass exactly as it runs in a berry, the same four-enzyme cassette would finish the job, the same flux models apply. Our five computational modules handle a monocot as easily as a dicot; to the ESM2 model, a grass synthase is just another protein. The entire difficulty lives in Module 6. Grass tissue does not want to regenerate. A corn cell, edited successfully, sits in a dish and refuses to become a plant, and decades of the world's best plant scientists have turned that refusal into merely a very hard problem rather than an impossible one. This is why corn and rice sit at the far end of the roadmap despite their flavor being no more complex than fruit we could build next year: the thinking is done; the tissue is the wall. We put four-to-six-year calendars on them not because we do not know how, but because persuading a grass to grow back from an edit is measured in seasons no model can compress.

It is worth being specific about the obstacle, because vagueness here would undersell both the difficulty and the progress against it. Most of our early dicots could be transformed through *Agrobacterium* — a soil bacterium that evolved, remarkably, to transfer its own DNA into plant cells, and which decades of science have domesticated into a delivery vehicle. The grasses largely resisted that bacterium, and so monocot transformation leaned for years on biolistics — literally firing DNA-coated particles into tissue — and on a long search for the rare cell types and culture conditions from which a grass will consent to regenerate at all. The deeper problem is the meristem: the pocket of stem cells from which a new plant grows, which in grasses is tucked away and reluctant, so that an edit made in an ordinary cell has no easy path to becoming a whole fertile plant that carries that edit into its seed. The monocot–dicot split is ancient — it runs back near the origin of the flowering plants themselves — and a hundred-million-year divergence in how a plant builds its body is not a thing a protein language model can shorten. What has changed, and what puts Storm corn and Dynasty rice on the roadmap at all rather than in the realm of wishes, is that the world's cereal science has turned grass transformation from impossible into merely slow and expensive. We inherit that progress. Our five computational modules were always ready for a grass; we

were waiting on the biology, and the biology, at last, is arriving — in seasons, at the price of seasons, which is exactly why these calendars read in years and not months.

The payoff, though, is proportional to the difficulty, and this is why the frontier is worth the years. Corn and rice are not luxury crops; they are the caloric foundation of most of the human diet, grown at a scale no fruit approaches. A method that reaches them reaches *staple* food — and the significance of that is best left to Part Four, where nourishment gets its own chapter, but it should be named here: the roadmap does not climb toward the grasses because grasses are a flavor novelty. It climbs toward them because they are where most of the world's eating actually happens, and a thesis that can reach them is a thesis about the whole of the table, not its garnish.

The Brassicas — **Wildfire broccoli**, **Nuclear cabbage and cauliflower** — are hard for a different and subtler reason, one worth naming because it is a chemistry problem rather than a tissue one. Brassicas are built around glucosinolates, the mustard-oil compounds that give them both their characteristic bitterness and their native pungency — two of our eight axes, tangled into one molecular family that the plant synthesizes in bulk. Building a *third* register, heat, into a plant whose metabolism is already heavily committed to glucosinolate production means our flux-balance analysis has to account for intense competition for shared precursors and carbon; the capsaicin road and the glucosinolate road draw partly on the same upstream supply, and asking the plant to run both at full volume is a genuine metabolic negotiation. This is why Wildfire and Nuclear carry construct counts and timelines closer to the grasses than to their fellow dicots despite regenerating more willingly: the tissue is tractable, but the metabolism is crowded. The **Deep Earth carrot**, finally, returns to the tuber-and-root problem the Seismic potato first raised, now in a plant from an entirely different family (Apiaceae), with a root metabolism that must be mapped from scratch — the storage-organ argument of Chapter 9, re-run in a carrot.

Tier 4 is the honest edge of what we believe we can build with the method as it stands today. Everything in it is reachable — we have no concept here that violates the thesis, no plant whose trunk is missing or whose chemistry forbids the cassette. But everything in it is *expensive*, in the one currency no model discounts: time, paid to living tissue that regenerates on its own schedule and not ours. Beyond Tier 4, the roadmap stops naming organisms, because beyond Tier 4 we are no longer confident we know how. That honest blank is the next chapter, and it may be the most important one in the book.

Chapter 12 · Tier 5 · De novo — *the tier we have not drawn.*

There is a fifth tier, and its defining feature is that it is empty. We have drawn four tiers of concept organisms with names, gene counts, and calendars; the fifth we have deliberately left blank, and this chapter is about why the blank is there and what it would take to fill it. Every honest map has an edge where the known world ends and the cartographer writes, in effect, *here the drawing stops.* Tier 5 is that edge. We think the most valuable thing a roadmap can do is show you exactly where its own confidence runs out — and ours runs out here, at the step from editing what evolution built to designing what it never did.



Where the drawing stops. Past the four named tiers, the method points at a horizon it cannot yet reach — and naming that honestly is the most rigorous thing the roadmap does.

The distinction that separates Tier 5 from everything before it is the difference between *editing* and *authoring*. Tiers 1 through 4 are all, at bottom, the same act: take a plant that already runs the universal trunk, and install a small cassette of terminal enzymes to add a turning the plant never evolved. It is editing — adding sentences to a text the organism was already writing. However hard the tissue, the strategy is fixed and proven: the trunk is given, and we finish it. Tier 5 is a different verb. It is the design of pathways that exist in no organism at all — flavor molecules with no natural template, built from first principles, authored rather than edited. Where Tier 4 asks "can we persuade this hostile tissue to accept a known cassette," Tier 5 asks "can we design a route to a molecule nature never made, and have it fold, function, and balance inside a living cell." We do not yet know the answer, and we will not pretend a calendar onto it.

We can, however, say precisely what stands between us and that tier, because naming the obstacles is the difference between an honest blank and a mystical one. Three things would have to mature.

First, **de novo enzyme design** would have to become as reliable as our variant *scoring* already is. Today, Module 2 excels at ranking substitutions in an enzyme that already exists — it can tell us which mutations improve a synthase, with the fluency of a model trained on all of known protein. It is a far harder thing to design an enzyme *from scratch* to catalyze a reaction no natural enzyme performs, and to

have it fold and function in a plant cell. The field of generative protein design is moving quickly, but it is not yet at the reliability where we would stake a four-year crop program on a wholly invented enzyme. When it is, Tier 5's first wall comes down.

Second, **whole-pathway flux balance** would have to extend from finding the bottleneck in an existing route to *designing* a route that balances before it is built. Our flux models today are superb diagnosticians of pathways that exist; asking them to be architects of pathways that do not — to lay out a novel sequence of reactions and guarantee it will not starve some essential process or pool a toxic intermediate — is a harder class of problem, and one where a wrong answer is not a failed experiment but a dead plant.

It helps to say what a Tier-5 target even *is*, so the blank does not read as mere mysticism. Tiers 1 through 4 all aim, in the end, at capsaicin — a molecule nature already perfected, which we relocate from the pepper into a new host. A Tier-5 target would be a flavor molecule with no natural exemplar at all: a compound designed to fit a human receptor in a way evolution never happened to try, synthesized by a route no organism runs, built because the Index predicts it would taste like something worth tasting and nothing in the living world yet makes it. That is a categorically different act from everything before it, and the reason it sits past the edge is not squeamishness about novelty but the plain fact that we cannot yet design such a molecule, such an enzyme to make it, and such a route to balance it inside a cell, with the confidence the rest of this book demands. We have spent sixteen chapters insisting that the length of our claims come only from arithmetic. Tier 5 is where the arithmetic is not yet ours to do — and so the tier holds a precise description of the missing calculation instead of a plant. When the calculation becomes possible, the first organism will be drawn here, and not one sentence sooner.

Third, and most soberly, **the regulatory and safety frameworks** for authored metabolism do not yet exist in the form such work would need. Editing an existing pathway sits within frameworks we understand and operate inside; authoring entirely novel small molecules into a food crop raises questions — of safety, of assessment, of consent — that are not ours alone to answer and should not be rushed. Part Four's chapter on guardrails takes this seriously, because the fact that something *could* be built has never, in our practice, been the same as the decision that it *should*. Tier 5's third wall is not technical at all. It is a matter of care, and care does not yield to cleverness.

So we leave the tier blank, and we are at peace with the blankness. A roadmap that named sixty-four organisms and then invented a fanciful sixty-fifth to fill a fifth tier would be lying to look complete, and across a library of sixty-four books the one thing that reads as unserious is a claim with no arithmetic behind it. Tier 5 has no arithmetic yet, so it has no cultivars. What it has instead is a precise list of what would have to become true — reliable de novo enzyme design, architect-grade flux modeling, and a framework of care equal to the power — and the moment any of those matures, we will draw the first organism into it and give it a name. Until then, the edge

of the map is the truest thing on it. The method took us out of the pepper, across the nightshade family, over the family line, to the foot of the grasses. It has not yet taken us past the edge of what life has already invented. When it does, there will be a Volume in this library to mark it. This chapter is a placeholder for a book we have not earned the right to write.



PART IV OF IV

Beyond Heat, Beyond Flavor

The same method across all eight axes, into nourishment, through the guardrails, and the economics of a whole kingdom of the buildable

Chapter 13 · Heat was only the first axis — *the method across all eight.*

We have spent twelve chapters carrying heat out of the pepper and into the garden, and now we must widen the lens one final time, because heat was never the only thing the method could build. Heat was simply the *first* axis — the loudest, the most legible, the one with a famous molecule and a century-old number already attached. But the platform that designs capsaicin is, in its bones, an axis-agnostic machine: read a pathway, score the mutations, find the bottleneck, write the edits, measure at the receptor. Nothing in that loop is specific to TRPV1. Point it at a different receptor and a different pathway, and it designs a different flavor — and we have, now, across all eight axes. This chapter is the proof that *beyond peppers* means *beyond heat* in the same breath.



One plant, three axes. Tea carries umami, bitterness, and astringency at once — and the method reads and tunes each against its own receptor, on the same Index.

Here is the full instrument panel — the eight axes, each with its receptor, its mechanism, and the live organisms where we speak it best:

AXIS	RECEPTOR	MECHANISM	SPOKEN LOUDEST IN
Heat · IFI-H	TRPV1	capsaicin, the trigeminal warmth	pepper, habanero
Sweetness · IFI-S	TAS1R2 · R3	sugar-free sweetness	stevia, monk fruit, sugarcane
Umami · IFI-U	mGluR1 · TAS1R1·R3	savory, glutamate-driven depth	tomato, shiitake, koji, soybean, yeast
Zero-Bitter · IFI-B	TAS2R38	bitterness attenuation (20–65%)	coffee, broccoli, hops, grape
Kokumi · IFI-K	CaSR	richness, continuity, mouthfeel	garlic, onion, vanilla, yeast
Cool-Burn · IFI-C	TRPM8 · TRPV1	cooling and warming	peppermint, spearmint, eucalyptus, lavender
Pungency · IFI-P	TRPA1	mustard-oil and allicin sharpness	black mustard, ginger, cinnamon, watercress
Astringency · IFI-A	salivary PRPs	tannin and polyphenol grip	tea, grape, pomegranate, blueberry

Each axis is its own version of the heat story, and three are worth telling in brief to show that the pattern holds. Consider **umami**, and consider it through the organism we opened the whole book with: the supermarket tomato. Umami is the taste of free glutamate at the mGluR1 receptor, and a ripe tomato is one of the plant kingdom's great glutamate factories — until it is picked green and shipped, at which point it loses the overwhelming majority of what made it savory. When we restored a tomato's umami, we did not add anything foreign; we re-activated the plant's own native glutamate metabolism — the GAD and GDH1 route, glutamate decarboxylase and glutamate dehydrogenase — and watched its IFI-U climb from a flat, shipped-out **300 to a restored 2,600**. That is the same arithmetic of loudness from Chapter 5, on a different axis and a different receptor: a tomato speaking its native savory word nearly nine times louder, not because we gave it a new voice, but because we turned its own voice back up. The method did not care that the receptor was mGluR1 and not TRPV1. The loop is the loop.

Consider **cool-burn**, the mint axis, which is the method's answer to a different question: not "how do we build a molecule" but "how do we build a *thermal illusion*." Menthol binds TRPM8, the cold receptor, and tricks the body into feeling a coolness that isn't there — the mirror image of capsaicin tricking TRPV1 into a heat that isn't there. The menthol pathway is a five-gene terpene route, and peppermint, our deepest cool-burn organism at 2,360 constructs, is where we have learned to ladder that illusion with the same precision we ladder heat. Heat and cool are, at the receptor level, the same trick played on two opposite channels — which is why our Cool-Burn axis is anchored to both TRPM8 *and* TRPV1, and why a single plant could, in principle, be designed to do both at once, arriving cool and finishing warm.

Consider, finally, **astringency** — the tannin axis, IFI-A, and the strangest of the eight, because it is not a taste at all. Astringency is a *mechanical* sensation: tannins and polyphenols binding and precipitating the proline-rich proteins in your saliva, stripping the mouth of its lubrication, producing the dry grip of strong tea or young wine or a fine dark chocolate. Its "receptor" is not a receptor

in the usual sense but the salivary proteins themselves. And yet the method handles it identically — read the polyphenol pathway, score the edits, measure the grip against the salivary standard — because the Index was built to be agnostic about the *kind* of sensation and faithful only to its loudness. Tea, which carries umami and bitterness and astringency all three, is the living demonstration: one plant, three axes, each read against its own receptor, all three on the same page, in the same units, by the same tongue.

The remaining four axes complete the instrument, and each one is the same loop pointed at a different receptor and a different chemistry. **Sweetness** (IFI-S) is read by TAS1R2·TAS1R3, a receptor heterodimer — two proteins that must pair before they can function — and the axis matters because its whole promise is sweetness *without sugar*. The intense natural sweeteners we design toward are not sugars at all: the steviol glycosides of stevia, the mogrosides of monk fruit, the sweet protein brazzein, each activating the TAS1R2·R3 pair at vanishingly low concentration. Stevia, our deepest sweetness organism at 910 constructs, is a study in tuning the *glycoside profile* — the balance among the various steviol glycosides a plant makes — because that balance decides whether the sweetness arrives clean or trails the faint bitter-licorice tail that has dogged stevia for decades. The method reads the glycosylation pathway, scores the editable steps, and shifts the blend toward the clean glycosides. Sugarcane and monk fruit round the axis out — one the world's oldest sugar crop, one its most concentrated natural sweetener — both read against the same receptor, both scored on the same Index.

Kokumi (IFI-K) is the strangest and most useful of the eight, because on its own it has almost no taste at all. Its receptor is CaSR, the calcium-sensing receptor, and what it reads are the γ-glutamyl peptides — glutathione and its kin — that produce not a flavor but a *modulation*: richness, continuity, the sensation of a dish being whole and mouth-filling and long rather than thin and quick. Kokumi is the difference between a broth that tastes finished and one that tastes like water with flavor added to it. Garlic and onion are its great natural sources, and yeast and vanilla carry it too, which is why a trace of any of them makes a savory dish feel *rounder* without making it taste of garlic or vanilla at all. Designing kokumi means designing the γ-glutamyl peptide pool, and it is the axis formulators ask us for most often — precisely because it is the one they can least achieve by simply adding an ingredient to the pot.

Pungency (IFI-P) is heat's sharp cousin, and the distinction between them is anatomical: where capsaicin opens TRPV1, pungency opens TRPA1, a different channel that reports a faster, sharper, more volatile sting — the nose-clearing rush of wasabi, the bite of fresh mustard, the heat of raw garlic, the prickle of cinnamon and ginger. The molecules are a motley family — allyl isothiocyanate from mustard and wasabi, allicin from garlic, cinnamaldehyde from cinnamon, gingerols from ginger — but they share TRPA1, and so they share an axis. Black mustard, at 2,790 constructs, is our

deepest pungency organism, and the Brassica glucosinolate pathway that builds its sharpness is the very same pathway that builds the family's bitterness, which is why the mustards and their cousins so often carry pungency and bitterness entangled together.

Zero-Bitter (IFI-B) is the only axis on the Index that works by *subtraction*, and it is one of the most commercially consequential. Its receptor, TAS2R38, is a bitterness receptor — the famous one that reads PTC and the glucosinolates — and the axis does not add a taste but attenuates one, dialing bitterness down by a measured 20 to 65 percent. This is the lever beyond roast and origin for coffee, beyond breeding for broccoli and kale, beyond fining for grape: the ability to quiet the specific bitter compounds a TAS2R38 reads while leaving everything else the plant does fully intact. A less-bitter coffee that is still unmistakably coffee; a broccoli a child will finish that is still entirely broccoli. Subtraction, measured on the same scale as every addition, because the Index does not care whether a number rose by adding a voice or by quieting one.

Eight axes, eight receptors, one loop. The pepper taught us the loop on the loudest axis. The other seven were waiting in the same machine the whole time. *Beyond peppers* and *beyond heat* are not two theses. They are one method, pointed in eight directions and forty-two plants, with the whole eight-by-forty-two grid — and the far larger grid beyond it — still being filled in a cell at a time.

Chapter 14 · The nourishment frontier — *the axis beyond flavor.*

There is a value that has hummed quietly under every chapter of this book, and here, near the end, we let it become the subject for a page, because it is where the method points next and because it is the reason the family built any of this in the first place. The whole apparatus — the trunk, the cassette, the eight axes, the Index — has been aimed at flavor. But flavor was never, for us, the final destination. It was the first thing we learned to design *into* a plant on purpose, and the question that has waited behind it from the beginning is simple: if we can design flavor in, what else can we design in? The honest answer is that the same method points, naturally and without strain, at nourishment — and we want to say exactly what we mean by that, and exactly what we do not.



The point of the whole enterprise, held in the hand. A plant grown to carry the most it can — flavor and nourishment as one design problem, not two.

Recall the chemistry of Chapter 2, because the nourishment argument is already latent in it. The universal trunk we borrow to build flavor — PAL through COMT, the phenylpropanoid road — is the very same road a plant drives to build its own antioxidants and defenses: the hydroxycinnamic acids, the flavonoids, the compounds a plant spends its carbon to protect itself with. When we turn that trunk up to deliver a flavor precursor, we are turning up a pathway whose other products are among the plant's most nutritionally interesting molecules. Flavor and nourishment, in other words, are not two separate engineering problems that happen to

share a plant. They are often the *same* pathway, read for two different outputs. This is the quiet structural fact that makes a nourishment frontier plausible rather than aspirational: we would not be bolting nutrition onto a flavor crop as an afterthought. We would be reading more of a conversation the plant is already having.

What the nourishment frontier means, concretely, is designing for attributes the plant can genuinely carry: higher densities of the vitamins and minerals a crop already makes, better *bioavailability* so the body can draw more from what is there, the restoration of nutritional depth that industrial shipping and early-green harvest strip away — the same "ninety-three percent the truck took" that flattened our opening tomato's umami also took much of its nutrition, and the GAD/GDH1 restoration that brought the savor back is kin to the work that would bring the rest back. These are attributes of the plant, measurable in the plant, grown into the plant by the same read-score-edit-measure loop that builds a flavor. A plant grown unhurried, with its trunk running full, in a controlled environment that lets it invest in itself rather than in mere survival, is a plant laying down more of what nourishes. That is a cultivation and design problem, and it is one the method is built to address.

Bioavailability is the part worth dwelling on, because it is where the method's precision matters most and where the temptation to overclaim is strongest — so we will be careful. A nutrient a plant contains is not the same as a nutrient a body can use; much of what makes food nourishing is locked behind the plant's own chemistry, bound to compounds that hinder absorption or held in forms the gut takes up poorly. The same read-score-edit loop that tunes a flavor molecule's delivery can, in principle, tune a nutrient's accessibility: shifting the form a mineral is stored in, quieting an anti-nutrient that would otherwise bind it, aligning the plant's chemistry so that more of what is present can actually be drawn from it. Paired with controlled-environment cultivation — growing a plant unhurried, with its trunk running full, investing in itself rather than racing to survive a field — this is a real and measurable design surface. And here is the line, held firmly: every sentence of that is a statement about a *plant* and what it carries, measured in the plant. None of it is a statement about a *person* and what will happen to them, and we do not cross from the first kind of sentence to the second. "This grain carries more accessible iron" is a fact we can measure and stand behind. "This grain will make you well" is a promise about a body, and promises about bodies are not ours to make. The frontier is real; the discipline that keeps it honest is what makes it worth pursuing at all.

Now the line we hold, and hold firmly, because it is a matter of integrity and not of caution. **Nourishment, in everything we write, is a value and an intention — never a claim.** We do not say, and this library will never say, that any plant we grow cures, treats, prevents, detoxifies, or heals anything. There are no health claims here, no "superfood" halo, no promises about your body's outcomes, because those are claims we have not earned the standing to make and would not make even if we had. What we say is narrower and truer: that food should be grown to carry the most it can, that the body should be able to draw the most from what it is given, and that

designing a plant to be both more delicious and more nourishing is a single coherent act of care rather than two competing ones. The difference between an attribute and a claim is the difference between "this tomato carries more of its native glutamate and its native nutrients, restored" and "this tomato will make you healthier." The first is a fact about a plant. The second is a promise about a person, and we do not make promises about persons.

This is why the roadmap climbs toward the grasses in Tier 4, and why that climb matters more than any superhot record. The fruits and berries and nightshades are where flavor is most *fun* to design. But the grains are where most of the world's actual nourishment comes from, grown at a scale no fruit approaches, feeding most of the people most of the time. A method that reaches corn and rice is a method that could, someday, make the staple foundation of the human diet carry more of what it should — and do it as an attribute, grown in, measured, never promised. That is the frontier beyond flavor, and it is the one the family cares about most. We built a way to design the taste of a plant. The thing we actually want to design is a plant worth eating — all the way down, flavor and nourishment as one grain in the hand, grown with intention, given to someone you love.

Chapter 15 · The guardrails — *what could be built, and what should.*

A book about what could be built owes its reader an honest chapter about restraint, because the entire thesis of this library — that flavor and nourishment are designable, at scale, across the plant kingdom — is also, read from a different angle, a catalog of power, and power without guardrails is not a roadmap but a hazard. We have tried throughout to let the arithmetic do the persuading rather than the adjectives, and nowhere is that discipline more important than here, where the temptation is to let ambition write checks that care cannot cash. So this chapter is the deliberate counterweight to the preceding twelve: not *can we*, but *should we, and how do we keep ourselves honest.*



Containment is not an afterthought to the method; it is part of it. Every edit in the library is grown, scored, and proved inside controlled conditions before it is anything else.

Start with the guardrail that is also a quality standard, because the two turn out to be the same thing: **reproducibility, held to a coefficient of variation at or under seven percent.** Across the whole library — nineteen thousand constructs, forty-two organisms — every cultivar is held to the discipline that batch after batch lands within ≤7% of its target profile. In wild and conventionally bred material, the same variety can vary by forty percent or more from plant to plant; a chili from one corner of a field can be nearly twice the heat of its neighbor. That variability is charming in a garden and disqualifying in anything a person is meant to rely on. Our ≤7% is

partly what makes the work usable by a formulator — but it is also a safety property, because a plant whose output you can *predict* is a plant whose effect you can *bound*. The superhot tier makes this vivid: a Phoenix pepper at 2,937,948 Scoville is a serious object, and the only responsible way to grow a serious object is to grow it reproducibly, so that its seriousness is known and never a surprise. Consistency is not merely a commercial virtue. It is the precondition for growing anything powerful responsibly.

Next, the guardrail of **containment and controlled environment**, which runs through this book as a visual motif precisely because it runs through the work as a practice. The flagship superhots were grown and proved in controlled chambers; the Laboratory tobacco and the Floral petunia exist on the roadmap specifically as contained proof surfaces so that a new edit is tested in an instrument plant, in containment, before it is ever committed to a food crop grown out. The order of operations is itself a guardrail: model, then prove in a contained model organism, then — only then — grow out. Nothing in the roadmap skips a tier to chase a headline. The tiers are, among other things, a sequence of increasing commitment, and the discipline of climbing them in order is a discipline of not getting ahead of what has actually been proved.

There is a second discipline folded into that order, and it is worth making explicit because it is the one most easily lost to enthusiasm: the discipline of the instrument organism. We keep two plants on the roadmap — Laboratory tobacco in Tier 1, Floral petunia in Tier 2 — whose only job is to be proven in, not eaten, and their presence is a standing refusal to test a new edit for the first time in a food crop. A novel construct earns its way into something a person will eat only after it has behaved, predictably and reproducibly, in an organism grown expressly to fail safely. This is slower than the alternative, and the slowness is the point. A method powerful enough to redesign the flavor of the food supply is a method that must never confuse *can* with *may*, and the instrument organisms are where that distinction is enforced in tissue rather than merely in principle. The same conservatism governs the sequence of tiers: we do not vault to the grasses because a grain would make a better headline than a berry; we climb in the order the difficulty estimators set, because climbing in order is how we keep every rung standing on a proven one below it. None of this is imposed on us from outside. We chose it, because a family that means to feed people with what it designs has to hold itself to a standard stricter than the one the world would currently enforce — and because the fastest way to forfeit the right to do this work is to do it carelessly even once.

Then the guardrail of **the open record**, which is how we invite the correction we cannot supply ourselves. The science behind this library does not sit locked in a company vault: **three patents filed, and twenty-one peer-reviewed papers published on Zenodo under a CC BY 4.0 license** — open, citable, and available to be checked, replicated, and argued with by anyone who cares to. A method that claims to redesign the flavor and nourishment of the food supply should be a method that shows its work, and ours is in the literature for exactly that reason. The patents protect the specific constructs;

the papers open the method. We think that is the right balance between building a durable company and being honestly accountable for a powerful tool — and the openness is itself a guardrail, because the fastest way to catch an error or a hubris is to let the people best equipped to catch it read the actual work.

And finally, the guardrail that governs the one in Chapter 14 and the blank tier in Chapter 12: **the language of attributes, never claims, and the discipline of the undrawn.** We will say a plant carries more of its native nutrients; we will never say it heals. We will draw four tiers of organisms we can estimate honestly; we will leave the fifth blank until we have earned it. These are the same guardrail wearing two faces — a refusal to let what is *possible* masquerade as what is *proven*, whether the overreach is a health promise or a fantastical organism. The whole library runs on one house rule, the same one that keeps its numbers from drifting across sixty-four books: say only what the arithmetic supports, and when the arithmetic runs out, say *that* instead. Power kept inside that rule is a roadmap. Power that steps outside it is just a story about ourselves, and we are not writing one of those. We are writing down a method, with its edges marked, so that the next hands — ours or someone else's — can see exactly where the solid ground ends.

Chapter 16 · The economics of a whole kingdom — *the matrix of the buildable.*

Step back far enough from the roadmap and a different kind of picture resolves — not a biology but a combinatorics, the sheer arithmetic of how many distinct things a portable method can make. We have spent the book on mechanism and organism; this chapter is about *multiplication*, because the deepest consequence of a transferable method is that it turns a linear catalog into a combinatorial one, and the difference between those two is the difference between a product line and a kingdom. The numbers here are the real ones from our own ledger, held to the single canonical set the whole library runs on, and they are worth laying out plainly because they are the quantitative shape of the thesis.



The bench where the combinatorics become real. Every construct in the matrix is, in the end, a sequence read, designed, and written here — the multiplication made physical.

Here is the matrix, counted:

SURFACE	COUNT	WHAT IT IS
Engineered constructs	~19,000	the validated library across all organisms
Live and purchasable today	1,000	available now
Total SKUs	78,800	the full addressable catalog
Host organisms (live)	42	the species edited today
Sensory axes	8	the receptors we read and write
Organism × axis intersections	336	42 × 8, the base grid of single-axis books and products
M2 dual-axis constructs	8,400	two axes composed in one plant
M3 triple-axis constructs	2,800	three axes composed at once
Full-spectrum (IFC) products	300	all eight axes, nanoparticle-delivered
Patents filed	3	the protected constructs
Peer-reviewed papers	21	Zenodo, CC BY 4.0 — the open method

Read the multiplication from the bottom up, because that is how the method actually builds value. Start with the base grid: forty-two organisms times eight axes is **336 single-axis intersections** — three hundred and thirty-six distinct answers to the question "this plant, tuned on that axis," each one a potential product and a potential chapter. That alone would be a large catalog. But axes compose, and composition is where the combinatorics ignite. The **M2 dual-axis library holds 8,400 constructs** — plants expressing two axes at once, a swicy pepper that is hot *and* sweet, a tea that is savory *and* less bitter — and the **M3 triple-axis library holds 2,800** three-axis builds. And at the apex sit the **300 full-spectrum products**, plants tuned on all eight axes simultaneously, delivered through nanoparticle formulation — the Phoenix principle generalized into a product line. Each layer up is not an addition but a multiplication, because a method that can combine axes does not add products one at a time; it multiplies the ones it already has by the ones it can compose them with.

The reason this combinatorial explosion is *achievable* rather than merely *imaginable* is the entire argument of Part One, now cashed out economically. If every one of these constructs required its own decade of basic biochemistry, the matrix would be a fantasy — 78,800 SKUs is a number no amount of conventional breeding reaches in a century. But the five portable modules mean the expensive thinking is done once and amortized across the whole matrix. Designing the ten-thousandth construct costs a fraction of designing the first, because the first thousand paid for the platform and the rest ride on it. This is the same unlock, in a different currency, that lets our sister library generate sixty-four *books* from one shared index: build the library once, render it through many lenses. A book is a lens on the index; a product is a lens on the construct library; and both are cheap to mint once the underlying thing is built, which is why both can run to numbers — sixty-four books, 78,800 SKUs — that look implausible until you see that they are views on one asset rather than one asset each.

Put a rough shape on the amortization, because it is the number that makes the catalog credible. The platform — the five portable modules, built once — is the fixed cost, and it is enormous: years of model-building, pathway reconstruction, assay development, the whole informational spine of the company. Every construct after the first draws on that spine without rebuilding it, so the marginal cost of *design* falls steeply as the library grows, bending toward the two things that never fully amortize: the per-species transformation of Module 6 and the per-batch agronomy of Module 7. This is why the catalog's structure mirrors the difficulty ladder so exactly. Within a single well-solved organism — a tomato, say — minting the next cultivar or the next axis is nearly free, which is how one plant spawns a five-rung Volcanic ladder and an eight-axis full-spectrum build without a proportional five-fold or eight-fold cost. Across organisms, the cost resets partway each time, because each new species re-presents the two non-portable modules. The 78,800 SKUs are therefore not 78,800 independent inventions; they are a modest number of hard-won organism beachheads, each multiplied internally by the near-free combinatorics of axes and cultivars that the portable modules make cheap. Seen that way, the implausible number becomes an arithmetic one: a few dozen tissue problems solved, times the thousands of flavor combinations each solved tissue can then carry. The kingdom is large because the method is multiplicative — and it stays affordable because the multiplication happens on the cheap side of the ledger.

But a combinatorial catalog raises a question a linear one never has to: not *can we build it* but *which of the buildable things should we build.* When the matrix holds tens of thousands of possible constructs, the scarce resource stops being capability and becomes *judgment* — the decision of which cells in the grid are worth a real plant and a real season. This is why the roadmap is sixty-four organisms and not six hundred, why the book program is sixty-four volumes chosen from a matrix of hundreds, why the live catalog is a curated thousand and not the full 78,800. The method's gift is that almost anything is buildable; the method's discipline is that we build only what is worth building — the organism whose flavor is genuinely new, the axis whose restoration genuinely matters, the staple crop whose nourishment genuinely feeds people. A kingdom of the buildable is not a mandate to build all of it. It is a responsibility to choose well, and choosing well — guided by the guardrails of the last chapter and the values of the one before it — is the real work that the arithmetic frees us to do. The multiplication is the easy part. The selection is the whole art.

Chapter 17 · The drawn and the undrawn map — *everything the alphabet can spell.*

We began this book looking at a tomato that had forgotten how to taste like itself, and we have ended at a combinatorial catalog of a whole kingdom's worth of flavor and nourishment. It is a long way between those two images, and the distance is the measure of what the method is. So let us close the way we opened — not with the catalog, but with the plant in the hand and the family at the table, because that is where the whole enterprise began and where it is finally for.



The hour the question was first asked, and the hour we still like best. The map has a drawn part and an undrawn part; standing at the edge of both is the whole of the work.

Here is the map as it stands, drawn plainly one last time. **Forty-two organisms are live**, carrying roughly nineteen thousand constructs, every one measured on the same eight axes against the same eight receptors, every one grown to within seven percent of itself batch after batch. **Sixty-four more are on the roadmap**, sorted into five tiers of honest difficulty: the Solanaceae core that is next season's planting, the extended nightshades one ring out, the cross-family leap where *beyond peppers* becomes literally true, the frontier synthetic where the grasses make us earn every year, and the fifth tier we have left blank because we have not yet learned to draw it. **Eight axes**, each a receptor and a pathway and a number. **A matrix of 78,800 possible SKUs**, of which we build only the thousand-and-growing that are worth building. That is the drawn part of the map, and we have tried in these chapters to show you every road on it with the arithmetic standing behind the lines.

And here is the undrawn part, which matters just as much. There is a fifth tier with no organisms in it yet, waiting on de novo enzyme design and architect-grade flux models and frameworks of care we do not yet have. There are hundreds of cells in the organism-by-axis grid we have not reached. There is the whole question of nourishment — the axis beyond flavor — that we have only begun to point the method toward. The map's edge is not a wall; it is a

shoreline, and shorelines are where the next work always happens. We said at the start that the honest blank is the most important part of the map, and we will end on the same conviction: a method is trustworthy exactly to the degree that it knows where its own confidence ends, and ours ends in specific, nameable places that we have marked rather than hidden.

Look back along the whole climb and the symmetry of the two generalizations is worth seeing whole. Across molecules, we began with one compound on one receptor — capsaicin on TRPV1 — and arrived at eight axes on eight receptors, the full span of what a human tongue can be made to feel, each axis its own pathway, its own EC_{50}, its own measured loudness, with Phoenix standing as proof that all eight can live in one fruit at once. Across organisms, we began with one genus and arrived at forty-two live species and sixty-four on the map, from the pepper's own nightshade kin out to the grasses that feed the world. Eight axes by forty-two organisms is already a grid of three hundred and thirty-six cells, most of them still empty, each one a plant waiting to be asked a question in the same units as every other. That grid, and the far larger one the roadmap opens, is the real inventory of what the method can reach — not a line of products but a *space* of them, measured on one Index, grown toward one standard, chosen one worthwhile cell at a time.

The thesis, then, stated one final time, now that the whole book stands behind it. Capsaicin was never a pepper problem. It was a worked example of a method — read the pathway, score the mutations, find the bottleneck, write the edits, measure the loudness at the receptor — and the method generalizes in two directions that turn out to be one. It generalizes across *molecules*, because heat was only the loudest of eight axes and the loop that designs capsaicin designs umami and menthol and tannin with equal indifference to which receptor it is speaking to. And it generalizes across *organisms*, because the pathway's expensive universal trunk runs in every vascular plant and only its last four paving stones are the pepper's private work — four enzymes, few enough to be given away, portable enough to be installed in a strawberry or a grain of rice. Five of seven modules move for free. The plant kingdom is the program. The pepper was only the proof.

One honest word about who built this, since it is a book written as *we* and should end as one. The computing was one pair of hands — the models, the edits, the long nights against the sequencer. The garden, the table, the palate that knew when a thing tasted *right* and not merely correct, the question that came up one night and would not leave: those belonged to all of us, and they are the half of the work no model performs. A flavor, in the end, is not a number. The number is how we learned to design it; the point of designing it was always to hand a plant to someone and watch their face. We built a way to read flavor like a text and write it like one. What we wanted, the whole time, was a better thing to put on the table — more delicious and more nourishing at once, grown with intention, offered with care.

The map has a drawn part and an undrawn part, and we are standing exactly at the line between them, which is the best place a maker can stand. Behind us: forty-two plants that prove the method travels. Ahead of us: a whole kingdom of the garden, and the discipline to build only what is worth building. The alphabet we learned from the pepper can spell far more than fire. This book was the first long sentence. The rest of the library — and the rest of the garden — is the paragraph we are still writing, together, at the table where it started.

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