Patrilines with Feedback; or, A Short Inquiry into the Y-Chromosome Bottleneck, Considered as a Control System

Rob Tow
Nova Lux, New Mexico High Desert, USA, Sol III
24 March 2026

Post-Neolithic DNA sequencing shows a steep collapse in Y-chromosome diversity; mitochondrial DNA shows no parallel collapse. The standard narrative reaches for war. A succession of cybernetic models—weak reproductive bias and thresholded fragility, an informal channel governed by female choice, hysteresis giving decline memory, then topology across partially coupled basins, then nested internal admissibility—shows that massacre-warfare is not necessary to produce the broad compression pattern. Punctuated pruning emerges from weak coupled filtering applied long enough. War is not required. Culture is.

1. Two Dark Thirty

Five nights ago, or two-dark-thirty, while Brenda slept beside me, I was on my iPad in night mode looking into the well-documented post-Neolithic collapse in Y-chromosome diversity, reconstructed from DNA sequencing of preserved corpses. Nature papers, and more. This gets paraphrased for the public as though 95% of men somehow vanished in a firestorm of Bronze Age melodrama. What actually vanished were male lineages. Not women. Not population as such—male lines.

The standard narrative reaches for war almost immediately. Patrilineal clans, slaughter, conquest, extinguished bloodlines, men on horseback setting everything on fire. Perhaps. Sometimes men on horseback *did* in fact set things on fire. But as explanation this strikes me as suspiciously blunt, a latter-day Kiplingesque just-so story for those who like their causality noisy, the sort of thing one says when one has not yet understood the machinery and prefers a gratifying

bang to an actual model. A kind of paleo-sociological physics envy.

The interesting asymmetry, as I read, is that mitochondrial DNA, passed down the maternal line, carries a separate information channel about lineage and diversity, and it does *not* show the same collapse. That alone should give one pause before galloping off toward the usual bloody male bloodbath.

That was the point at which I stopped wanting a story and started wanting a machine. The point was simple enough; ask what actually collapsed, what kind of machinery could do it, and whether the obvious story had mistaken noise for causality.

So I built it.

2. What Actually Collapsed

Before going any further, it is worth stating the problem cleanly, because public retellings of this episode are so routinely sloppy that they smuggle confusion in at the door. What collapsed was not the male population as such. It was not some global culling in which ninety-five percent of actual living men were swept off the board in one lurid Bronze Age bonfire. What narrowed, according to the sequencing literature, was Y-chromosome lineage diversity: the number and distribution of surviving paternal lines. That is a different claim, and not a minor bookkeeping distinction. A population may remain substantial while the effective diversity of male descent lines contracts sharply. If one cannot keep that straight, one is not even arguing about the right phenomenon.

What figure 1 gives me is the observational anchor, not a template to be copied point for point. My later plots are not fitted to reproduce this exact skyline shape as though I were tracing it with a steadier pen. The skyline reconstruction is already an inferential summary of the sequencing literature, smoothed and compressed

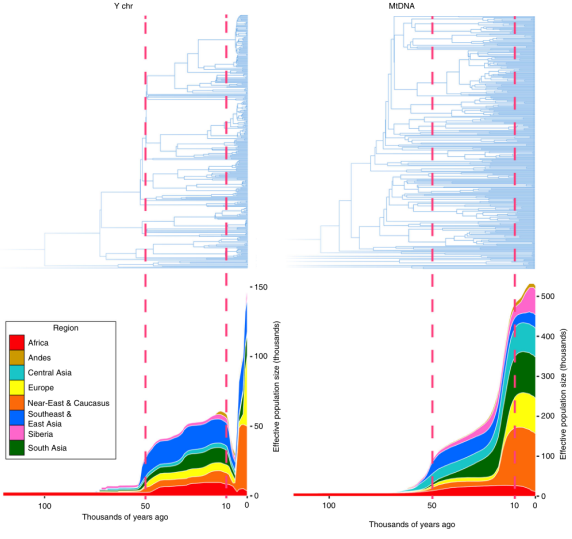

Figure 1. Cumulative Bayesian skyline plots of Y-chromosome and mtDNA diversity by world regions. Reprinted from Karmin et al. under a Creative Commons License (Attribution-NonCommercial 4.0 International). Note the very sharp dip in the left plot starting circa 10KY ago, followed by a rebound.

by its own machinery. The burden on my model is different and narrower: to recover the same class of observable facts, namely (1) a bounded contraction in Y-line diversity, (2) no comparable collapse in the mitochondrial channel, and (3) later expansion of the surviving male lines. Resemblance in exact profile would be pleasant, but it is not the epistemic standard here. The standard is whether the machine can generate the same kind of world.

The asymmetry is what makes the whole thing interesting. The Y chromosome carries one record of descent, strictly paternal. Mitochondrial DNA carries another, strictly maternal. They are not the same channel and should not be expected to tell identical stories, but that is precisely why the divergence matters. The post-Neolithic signal, in broad outline, is a strong contraction in Y-line diversity without a corresponding collapse in mitochondrial diversity. Women did not disappear in parallel. Population as a whole did not simply blink out and reappear. Something acted differentially on male lineage persistence. That immediately shifts the problem from melodrama to mechanism. If one inheritance channel tightens sharply while the other does not, then one is looking not merely at death, but at structured reproductive throughput.

This is the point at which the usual newspaper paraphrase becomes actively misleading. "Most men died"is a cruder statement than the data warrant, and in a way that matters. One can lose many lines without killing most bodies. If, generation after generation, a narrower subset of men leaves descendants while many others leave few or none, the Y pool will compress even in the absence of some singular exterminating event. The historical imagination, especially in its popular register, prefers cavalry, screaming, and lit fires. Fine. History has offered plenty of all three. But the empirical signal here is not merely that men died. Men have always died. The signal is that paternal descent became much more concentrated than maternal de-

scent, and that this concentration appears bounded, structured, and followed in many accounts by expansion of the surviving lines.

That is the observational problem I care about. Not "why were men violent," which is too theatrical and too cheap, and not "which tribe killed whom," which leaps far ahead of the evidentiary warrant, but rather this: what kind of social and reproductive system causes strong compression in male-line diversity while leaving the maternal channel comparatively broad? Put in engineering terms, what process reduces the effective entropy of the Y-lineage pool without requiring equivalent collapse elsewhere? Once stated that way, the problem stops being a saga and becomes something modelable. And once it becomes modelable, the massacre-first story begins to look less like an explanation and more like an impatient refusal to ask what machinery might have been doing the work.

3. Why the Standard Story Is Too Blunt

The massacre-first story is not merely common because it is possible. It is common because it is narratively irresistible. It offers chiefs, banners, horses, fires, towers, screaming, trophies, and the great explanatory narcotic of visible violence. It flatters the historical imagination because it presents causality in a form that can be staged. One sees the raid, the conquest, the broken walls, the smoking settlement, and imagines one has therefore seen the mechanism. Very often one has seen only the pageant. Violent explanation enjoys a prestige wildly out of proportion to its actual analytic precision. It feels hard-headed, unsentimental, and serious. Quite often it is merely explanatory laziness wearing a helmet with Viking horns (which they did not).

None of this is to deny violence. Men on horseback did in fact set things on fire. Raiding, conquest, slavery, seizure of women, extin-

guished lineages—history offers an inexhaustible inventory. But the existence of violence does not grant it explanatory monopoly over every pattern that can be made to resemble aftermath. A bounded compression of paternal lines, especially one not mirrored in the maternal channel, is not automatically best understood as a saga of singular events. It may instead reflect a long, structured filtering of reproductive throughput: legitimacy, inheritance, prestige, exclusion, taboo, alliance, and the many other small gates through which societies quietly do their work. What is wanted here is not more event-history, but control-system history.

That distinction matters because most people are educated to think historically in events. Battles, kings, migrations, collapses, treaties, decisive turns, and the scenery of consequence. The machinery, when it appears at all, is treated as backdrop. But a great many historical outcomes are not event-produced in that theatrical sense. They are system-produced: the cumulative effect of repeated constraints, incentives, permissions, exclusions, and feedbacks operating over many generations. Once one sees the Y-bottleneck as a problem of lineage throughput rather than merely one of male mortality, the old taste for cavalry and bonfires begins to look less like realism than like a refusal to ask what sort of machinery could have been doing the work the whole time.

And that, to my eye, is the real threshold. Once one admits the possibility of control-system history, verbal preference is no longer enough. One needs a machine. Not because machines are magically objective, but because a runnable model will at least force the question into a form where rates, thresholds, channels, feedbacks, and topology have to declare themselves. The alternative is to keep rewarding whatever story arrives wearing the loudest boots.

4. Why I Built a Model

So I built a model. Not a historical pageant with costumes and heroic moustaches, but a stripped cybernetic sketch, in runnable code shared on GitHub. The point was not to simulate prehistory in all its decorative clutter with a Bronze Age puppet theatre of arrows, raids, and men shouting in translated subtitles. The point was simpler, and harder: to ask what sort of system compresses male-line diversity, and whether war is actually required to get there? Once the problem is stated against the evidentiary ground-truth, and once one has resisted the reflex to confuse event-history with explanation, a runnable machine with tuneable knobs becomes preferable to glorious tales. One can then ask a bounded question—if I design into the machinery weak reproductive bias, thresholded fragility, a secondary informal channel, and some representation of cultural filtering, does the broad behavioral profile begin to emerge? Or does it not? At the outset, this was unapologetically a spherical cow—the old physicist's joke about simplifying a beast into tractable geometry for computing its trajectory when it jumps over the moon. The point is not zoological fidelity. The point is to see whether even a stripped machine could be made to do real work.

What mattered to me from the beginning was that the machine should be mechanistic without pretending to omniscience. I was not trying to rebuild the Bronze Age in silicon, or to conjure some poor ghost of actual prehistory out of a few variables and a GPU. I wanted a model austere enough that one could see its moving parts, yet rich enough to test whether the usual massacre-first story was in fact necessary. So the first versions tracked lineages, reproductive throughput, informal leakage, threshold effects, and diversity over time. Shannon entropy entered because what I cared about was not merely the raw count of surviving lines, but the changing infor-

mational structure of the pool. A system can lose diversity before it loses everything, and the loss of diversity changes what happens next. Once redundancy falls, fragility rises. The machine needed to be able to show that.

This is also the place where I part company with the habit of treating culture as scenery. Religion, inheritance, legitimacy, status, purity stories, marriage rules, all the memetic apparatus, are not commentary. *They are control surfaces.* Culture does not merely supply the filtering machinery. It edits it; and there is feedback and recursion. Religion, legitimacy, inheritance, purity rules, and status are control surfaces acting on lineage throughput, and control surfaces do not remain fixed. They are revised, hardened, softened, renamed, sacralized, and redistributed—sometimes explicitly, by law, priesthood, conquest, or reform, and sometimes emergently, by prestige, drift, imitation, taboo, and opportunism. I have in mind certain events in the Middle East about twenty-three centuries ago, where what was being edited was not merely power, but the control surfaces themselves, and where that editing became unusually legible. This is why I keep insisting on machinery. The reproductive filter is not merely applied through time; it is itself historically revised through time. Of course I simplified. I arrived at matrices.

The first model was never meant as a final account. It was an opening wedge. A way to test whether weak, repeated, structured filtering could do more work than the loudest historical folklore usually grants. If it failed, good; failure would at least tell me something. If it succeeded even crudely, then the next question would no longer be whether such machinery could matter, but how to represent it less falsely. That, in fact, is exactly what happened. And then I proceeded further.

5. Model Family I: Weak Bias and Thresholded Fragility

The first family of models was built to test a question much narrower than "what happened in the Bronze Age," and therefore much more useful. If one imposes only weak, repeated selective bias on male-line reproduction, does lineage diversity begin to compress in a serious way? And if one adds fragility below a threshold, does that compression begin to acquire shape rather than merely slope? This was still a spherical cow, of course—a knowingly simplified beast, useful precisely because one can see the machinery without losing it in decorative historical upholstery. But it was a spherical cow built to answer a real question. Could weak coupled filtering, sustained over generations, do much more work than the massacre-first reflex usually allows?

At this stage the machine was simple enough to state plainly. It carried a population of paternal lineages forward in discrete generations. Each lineage had reproductive weight. That weight was not wildly distorted; the whole point was to see what weak bias, applied repeatedly, could do. Diversity was tracked both by raw surviving line count and by Shannon entropy. I use Shannon entropy here as a collapsed instrument—not the native metric of the DNA sequencing papers I had read, but a principled compression of their richer observable pattern into something the machine can track. The system then needed one further insult: fragility. A lineage above threshold could persist with ordinary reproductive luck, but below threshold it became increasingly vulnerable. That was the first move away from simple thinning and toward actual pruning.

In each generation, the machine did something like this:

1. Carry forward the current paternal lineages and their reproductive weights.

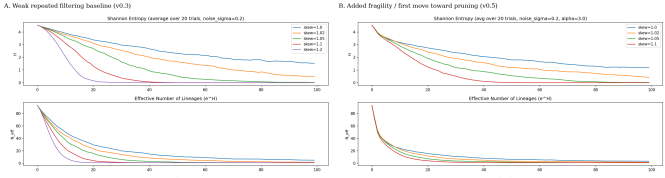

Figure 2. Early scalar compression under weak repeated filtering. **A:** Weak bias alone substantially compresses paternal-line diversity. **B:** Adding fragility shifts the behavior away from mere thinning and toward pruning.

2. Apply weak formal bias so that some lines have slightly better reproductive throughput than others.
3. Generate the next generation's lineage distribution from that weighted reproduction.
4. Measure diversity in the resulting pool, both as surviving line count and as Shannon entropy.
5. Apply thresholded fragility to lines that had fallen into low-support states.
6. Repeat, and watch whether the system merely thins or begins to prune.

That may look almost embarrassingly austere. Good. Embarrassment is often a sign that one has not yet hidden the moving parts under upholstery. And the moving parts here mattered. Even weak bias, if applied long enough, could indeed compress male-line diversity substantially. One did not need an exterminating climax. A low-gain filter, operating generation after generation, did real work all by itself. More interestingly, once fragility entered, the system no longer behaved as a smooth decline alone. Lines could hang on for a while and then fail more abruptly once support fell far enough. The

beginnings of clustered pruning started to appear. Not yet a full account of punctuation in the historical sense, but enough to suggest that the machine was beginning to acquire temporal shape rather than merely direction.

That result mattered because it broke the monopoly of the obvious story. If even this stripped machine, still very much a spherical cow, could produce strong compression under weak repeated filtering, then massacre-first explanation had already lost its exclusive claim to seriousness. But the model was still lying in a useful way. It treated the space too smoothly. It had lineages and thresholds, but not yet much structure in admissibility, not yet a real distinction between formal and informal channels, not yet the topology that would later turn out to matter so much. Useful, then, but still a spherical cow. It had shown that weak coupled filtering could matter. It had not yet shown how the social space in which that filtering lived might itself be shaped.

6. Informal Channel, Female Choice, and Signal Separation

Up to this point, the machine was still too obedient. All reproductive throughput passed through a single formal channel, which is already enough to show that weak coupled filtering can do real work, but is not how human systems actually behave. Institutions do not exhaust biology. They merely pretend to. Formal order may allocate marriage, inheritance, legitimacy, and status according to one logic, but actual reproduction has always had other ways of slipping through. That secondary channel need not be especially large to matter. It need only persist. And in human terms there is an obvious biological suspect here: female choice, in the old Darwinian sense, stripped of perfume and undergraduate sentiment. Not ro-

mance, not a greeting-card anthropology of liberated desire, but a secondary selective process acting within, around, and sometimes against the formal order.

Once that possibility is admitted, the first machine ceases to be enough. One must distinguish between formal and informal reproductive signals. The formal channel may reward lineage, status, legitimacy, inheritance position, and all the official sorting machinery by which a social order claims to know who counts. The informal channel may sample some of the same things, or partly different ones: attractiveness, vigor, charisma, protection, proximity, local dominance, opportunism, sheer mammalian fact. The exact content matters less, at this stage, than the structural point: the two channels do not have to be identical. They may be aligned, partly aligned, or substantially independent. And that distinction turns out to matter a great deal.

So the next family of runs split the machine accordingly. Formal reproduction still tracked the official signal, but a secondary informal channel now redistributed some reproductive success by a different rule. In rough terms, the model cycle became:

1. Carry forward the surviving paternal lines and their formal reproductive weights.
2. Apply formal reproductive bias according to lineage/status structure.
3. Apply a secondary informal channel governed by a separate attractiveness-like signal.
4. Combine the effects of the two channels into the next generation's lineage-share distribution.
5. Measure diversity again, both by surviving line count and by Shannon entropy over those lineage shares.
6. Repeat, and vary the degree of alignment between the formal and informal signals.

The interesting thing was not merely that leakage existed, but that its structure mattered. If the informal channel was tightly aligned with the formal one, it amplified concentration. The same lines that already did well under the official order gained additional reproductive advantage through the unofficial one. Compression accelerated. But if the informal channel sampled a more independent signal, it softened the collapse. Weak lines that might otherwise have vanished altogether could still receive some reproductive share through the secondary path. Diversity persisted longer. The machine, in other words, no longer behaved as though leakage were just noise. Leakage had a topology of its own, even before I had properly earned the right to call it that.

This was an important gain. It meant that the argument had moved beyond the crude claim that "culture filters reproduction." Yes, it does. But biology is not fully housebroken by ideology, and the degree to which informal reproduction echoes or departs from formal order changes the behavior of the system in visible ways. Once that became clear, the earlier scalar models began to look both more useful and more inadequate. Useful, because they had already shown that weak repeated filtering could matter. Inadequate, because they still treated the space too smoothly, as though both channels were flowing through the same vessel. They were not. That realization would become decisive later. For the moment, it was enough to show that the path outside the formal order was not an afterthought, but part of the machinery.

7. Punctuated Pruning and Hysteresis

By this point the machine was no longer merely thinning lineages. It was beginning, in a crude but recognizable way, to prune them in time. That distinction matters. A smooth decline is one

thing; a system in which lines can hang on, wobble, and then disappear in clusters is another. The early runs were already suggesting that the bottleneck should not be thought of only as a uniform narrowing, but as a process with punctuation of its own. Not punctuation in the full historical sense yet—certainly not enough to pretend one had captured the actual temporal grain of the ancient world—but enough to see that the machine could generate episodes of more abrupt loss once support fell past certain points.

The key addition here was hysteresis. A lineage does not merely occupy a state; it carries some memory of how it got there. Support can erode gradually while the line still persists, and recovery is not necessarily symmetric with decline. Once a line has fallen into a weakened regime, it may require much more than a slight reversal of conditions to regain robustness. This matters because historical systems are full of one-way thresholds masquerading as continuity. A lineage can remain visible after it has ceased to be truly secure. It can appear to endure while already living on borrowed time. Hysteresis gives the model a way to express that sort of lagged vulnerability, where the effect of prior erosion remains in force even when the immediate conditions have softened.

In rough terms, the machine was now doing something like this:

1. Carry forward the surviving paternal lines and their current reproductive weights.
2. Apply the relevant reproductive filtering.
3. Measure the resulting lineage-share distribution.
4. Update each lineage's support state, not only by its present condition but by its recent trajectory.
5. Apply fragility differently depending on whether a line is merely low, or has been low for long enough to enter a more vulnerable regime.
6. Repeat, and watch for the emergence of clustered losses rather

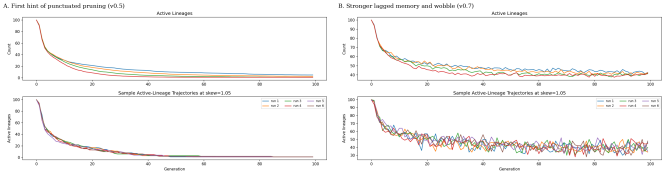

Figure 3. Punctuated pruning and hysteresis in the early single-basin branch. **A:** In **v0.5**, the first shaped-fragility run, decline begins to move away from smooth thinning toward more clustered, thresholded loss. **B:** In **v0.7**, stronger lagged memory produces more visible wobble and persistence in weakened states, showing how prior erosion can continue to shape the time series even after immediate conditions soften.

than mere smooth decline.

What this produced was not a melodramatic cliff, which would have been too easy and probably too stupid, but something more interesting. Lines could persist for a time in weakened states, then fail in bunches once enough prior erosion had accumulated. The system acquired temporal asymmetry. It had a history, in the narrow but important sense that current state no longer told the whole story. The machine began to remember, and that memory changed the shape of the decline. This is what I later took to calling punctuated pruning: not because I wished to borrow prestige from other fields, but because the model really was beginning to show clustered, history-dependent loss rather than featureless attrition.

This was another place where the model became more useful and more inadequate at the same time. Useful, because it showed that one did not need a single exterminating event to get something that looked temporally uneven and locally abrupt. Inadequate, because the punctuation was still being generated inside a space too smooth

and too understructured. Hysteresis improved the machine, but it did not yet solve the deeper problem of social topology. It gave the decline memory. It did not yet give the social space enough shape.

8. The Spherical Cow Problem

By this point the scalar family had earned its keep. It had shown that weak coupled filtering could compress male-line diversity without requiring massacre-warfare as the only story in the box. It had shown that fragility could turn thinning into pruning, that hysteresis could give decline a kind of memory, and that formal and informal channels need not behave the same way. That is already a respectable amount of work for a toy machine. I do not sneer at the spherical cow. The old joke survives because it names something real in scientific and engineering practice: sometimes one must simplify brutally in order to see anything at all. A perfectly faithful fuzzy mess is often more informative than a deliberately false machine whose falsehoods are precise. The problem begins when the simplification continues to pay rent rhetorically after it has ceased to pay rent analytically.

That was where I had arrived. The scalar models were useful, but they were useful in a particular and narrowing way. They still treated the social field too much as a common vessel. Bias could vary. Fragility could vary. Informal leakage could vary and hysteresis could remember. But the space itself remained too smooth, too shared, too polite. The lineages were still being filtered inside something that behaved, for modeling purposes, too much like one continuous medium. That is exactly the kind of lie one tells early and must later outgrow. A spherical cow is no disgrace at the beginning of inquiry. It becomes a disgrace only when one keeps admiring its smoothness after the problem has made roughness causal in the

wind tunnel.

The roughness here was no longer merely story-telling decoration. It had become part of the machinery. By then I could already see that "outside the formal order" did not mean nowhere in particular. Reproductive leakage had a topology of its own. Informal reproduction was not magical diffusion through the same homogeneous social soup. It was traversal—partial, structured, constrained traversal—through a field whose permissions and blockages mattered. Likewise, "single-basin" had stopped being just a manageable simplification and had started to become a real distortion. If the social field is broken, contoured, and only partly connected, then one does not get to keep modeling it as though everyone were effectively swimming in the same pool, merely with slightly different weights attached to their ankles.

This is the point at which one has to be a little more explicit about what kind of lie the scalar family was telling. The problem was not merely that it had too few variables. Anyone can add variables until the machine resembles a Victorian attic. The problem was subtler. The variables it did have were living in the wrong sort of space. They were being treated as though they occupied a field that was too smooth, too low-dimensional, and too politely orthogonal. That is to say, the model behaved as though the important directions of variation were few, fairly independent, and globally available. But the real social field, even in toy form, was beginning to reveal itself as something else: folded, co-varying, and only partly connected.

That distinction matters enough to belabor. In a low-dimensional and nearly Cartesian picture, one imagines a few reasonably clean axes: more status, less status; more legitimacy, less legitimacy; more fragility, less fragility; perhaps a knob for leakage and another for memory, and off one goes feeling wonderfully scientific. But actual social machinery is rarely so courteous. Legitimacy is not merely

one axis. It covaries with inheritance, cult, law, prestige, kin recognition, and access. Basin membership is not merely another axis. It changes who can reach whom, under what permissions, through what bridges, and at what cost. Informal traversal does not simply add a little noise to a formal rule. It changes the meaning of adjacency. Some dimensions matter only locally. Some become visible only at certain scales. Some that look independent from afar turn out, up close, to be folded into one another like bad bureaucracy and old theology.

So when I say the scalar family was beginning to live in the wrong coordinate system, I do not mean merely that it needed a few more knobs. I mean that the machine was representing the social field as though it were a nice room with a few clean axes, when it was beginning to look more like a folded landscape whose effective dimensions changed with scale and connectivity. That is not ornament. It changes what counts as nearness, what counts as leakage, what counts as separation, and what kinds of persistence or collapse are even possible.

This is also why simply adding yet another scalar penalty, threshold, or memory term would have been the wrong move. One can always keep decorating a scalar machine with fresh punishments and compensations and tell oneself one is marching toward realism. Often one is merely protecting an inadequate representation from the consequences of its own inadequacy. There comes a point at which refinement becomes evasion. The point had ceased to be "how many knobs can I add before the curve looks better?" The point had become "what kind of space is this machinery actually living in?" Once that becomes the question, honesty requires more than refinement. It requires re-representation.

That is why I do not think the scalar family failed. It did something better than that. It succeeded while still lying in the wrong

coordinate system. That is a much more interesting outcome. The machine showed that weak, repeated, culturally mediated filtering could do real work. It showed that hysteresis could give decline memory. It showed that informal leakage was structured and consequential. It showed enough, in short, to make the next inadequacy impossible to ignore. The old representation had ceased to be innocently false. It had become actively misleading in the very place where the next questions were beginning.

So this section marks a methodological threshold. The earlier models were not discarded because they were toys. They were surpassed because they had taught me what sort of toy they no longer had permission to remain. They could still instruct, but they could no longer suffice. Once the social field starts to reveal itself as broken, contoured, folded, and only partly connected, one does not get to go on treating it as a common vessel with a few nice axes and a handful of tuneable insults. Time to do some hyper.

That is why the next turn had to be *topological*. The question was no longer simply how much filtering, or how much fragility, or how much memory. The question became: filtering across what shape? fragility inside what structure? memory carried through what partially connected field? Once asked that way, the old scalar family had done its work. The spherical cow had earned its retirement. It was time to let the beast acquire edges.

9. Basins, Nesting, and Admissibility in the Valley of the Matryoshka Dolls

Once the scalar family had taught me that the common vessel was lying, the next question was obvious. If the social field is not one pool, then what is it? The first answer was not graphs. It was *basins*. By basin I mean a partially bounded reproductive field: a

zone in which pairing, inheritance, legitimacy, and social recognition circulate more readily within than across its edges. Or a valley, if you prefer. Geography may matter. Ritual may matter. Language may matter. Rank may matter. Kin recognition may matter. The point is not that a basin is only one of these, but that it is an interaction field in which access is easier internally than externally. Once that much is admitted, the old single-vessel picture already begins to look suspiciously polite.

This was the Daisyworld turn—Daisyworld, after Watson and Lovelock, as a deliberately simple self-regulating toy world in which local reproductive dynamics feed back into global conditions. The point, in my hands, was not reviving botanical whimsy. It was the first honest admission that several weakly coupled basins might already do a better job than one universal bowl. And indeed they did. Once the model moved from one common pool to several partially coupled ones, the global story improved. Compression no longer looked as forced; survivor monopoly weakened; aggregate timing began to behave less like a cartoon. But local timing was still off, and that mattered. The machine had become more honest without yet becoming honest enough.

That shortfall turned out to be instructive. Several basins are better than one. They are not yet enough. A basin is not a smooth interior simply because it has an edge. This is where the matryoshka dolls enter. One opens a basin and finds that it is not one room but several nested chambers, shells, and inner jurisdictions. Membership in the larger container does not imply equal admissibility inside it. One may belong and still not reach. One may circulate locally and still fail to cross inward. One may have formal access without practical traversal, or practical traversal without formal standing. The more one opened the dolls, the less plausible the common vessel became.

That is still only the beginning. It is not enough to say that there are inner levels. One has to say what those levels do. The outer shell may grant broad basin membership while saying little about actual pairing reach. Inside that, one may encounter class-like chambers, local prestige strata, ritual interiors, kin-recognized corridors, and weak bridges that connect some shells but not others. Some routes are broad but shallow. Some are narrow but decisive. Some can be traversed formally but not practically. Others can be traversed practically while remaining formally disavowed. One "inside"is not the same as another. Not every chamber speaks equally to every other chamber. The whole point of the dolls is that containment is not homogeneity.

This is why a basin can no longer be treated as a blob with an edge. Its interior has to be thought of as itself structured. There are shells, choke points, inner rooms, sacrificial margins, favored bridges, and pathways that are open only under particular conditions. A lineage can be secure in the outer doll and precarious in the inner one. It can possess nominal standing while lacking effective reach. It can travel sideways more easily than inward. It can cross one bridge and fail at the next. Once one sees that, "being in the basin"ceases to be a useful scalar summary. The interiors have become causal.

This is also where admissibility stops being sociological mush and becomes machinery. By admissibility I mean something quite concrete: who can traverse to whom, under what formal permissions, under what informal leakages, through what bridges, with what penalties, and with what history attached. It is adjacency under constraint. It is not enough to say that a basin contains people. The relevant question is how the basin is wired. Are there shells? Are there choke points? Are there high-status bridges that pass one way more easily than the other? Are there protected interiors, sacrificial margins, weak ties, ritual gates, local elites, or permeability gradi-

ents? These are not decorative details. They are the social equivalent of geometry.

The next step, then, was not merely "more basins." It was explicit internal structure inside them. This is where the essay has to stop speaking only in matryoshka metaphors and speak in machinery. Inside each basin the later eligibility-graph family eventually carried fixed internal eligibility classes, with distinct formal and informal traversal across those classes. In other words, the dolls were no longer only nested conceptually. Their interiors were being given a small matrix structure: who could reach whom, under what mode of passage, along which bridges, with which routes later tightening first. The matrix was not there for decoration. It was there because smooth interiors had already stopped being truthful.

This is why the nested turn here is not merely an arbitrary modeling complication. It reflects a more general Gaian cybernetic principle: regulation in complex systems often emerges through nested, partly permeable compartments rather than one homogeneous field. Large-scale behavior is not imposed from nowhere and not read off a single smooth surface. It emerges from local dynamics inside partially bounded compartments whose states feed upward and outward, while larger-scale conditions feed back down into them. That is true of ecologies. It is true of institutions. It is true of many technical control systems. And it appeared, with increasing insistence, to be true here.

So the next truthful representation was an eligibility graph. Not because graphs are fashionable, and not because once one has discovered network theory one must inflict it on every problem like a graduate student with a fresh hammer, but because the machine was now asking graph questions. Which lines are adjacent? Which edges are open, weak, forbidden, or contingent? Which bridges connect basins? Which interiors are nested? What happens when for-

mal order tightens not everywhere at once, but along selected edges and through successive admissibility changes? These are graph questions whether one likes the notation or not.

The progression through the codebase reflects exactly that realization. **v0.12** was the first real topological advance, the point at which several weakly coupled basins replaced the universal bowl. **v0.15** is, for present purposes, the best scalar/topological Daisyworld family, because it gets the three-basin global picture to behave much more plausibly even while its local timing remains too early. After that, the graph family begins in earnest. **v0.17** and **v0.18** establish the baseline intuition that nested admissibility can be represented explicitly rather than smeared into scalar penalties. **v0.19** adds staged tightening. **v0.20** then sharpens the question further by tightening edge-by-edge rather than by one-shot decree or matrix-wide administrative snowfall. I am getting ahead of myself. The point here is simpler: topology was no longer a flourish. It had become the next honest thing to do.

In rough terms, the representational shift now looked like this:

1. Stop treating the whole social field as one common reproductive vessel.
2. Represent several partially bounded basins with easier internal than external traversal.
3. Allow those basins to contain internal admissibility structure rather than pretending their interiors are smooth.
4. Open the dolls and represent nested inner levels rather than treating basin membership as sufficient description.
5. Represent permissions, barriers, bridges, and local traversals explicitly as edges rather than indirectly as scalar penalties.
6. Tighten or relax those edges over time according to formal order and informal leakage.
7. Measure not only aggregate compression, but how persistence,

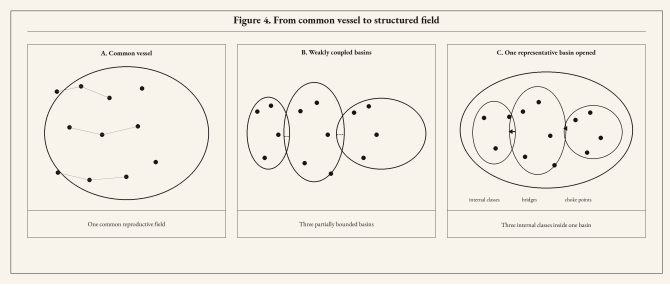

Figure 4. From common vessel to structured field. **A:** The early single-basin picture, with all lineages living in one common reproductive vessel. **B:** Three weakly coupled basins, improving the global story by allowing partial separation across reproductive fields. **C:** One representative basin opened to show nested admissibility: internal classes, bridges, choke points, and partial traversal, motivating the eligibility-graph family.

> collapse, and survivor ordering depend on the shape of the field itself.

This was, to my eye, the decisive turn. The problem was no longer merely how much filtering could be imposed, nor how much memory the decline could carry. The problem had become one of shaped social spaces. Across what compartments did filtering act? Through what bridges did leakage pass? Inside what nested jurisdictions did legitimacy, status, and access cease to align? Inside what inner shells were some lines effectively local while others remained traversable across the basin? Once the question took that form, the old scalar family could still instruct, but it could no longer govern. Basins were already more truthful than bowls. Nested admissibility was more truthful still. My machine investigation had ceased asking for more knobs. It was asking for a map.

10. Model Families, Truthful Lies, and Representational Bargains

Each of these model families is a kind of truthful lie. That is not an embarrassment peculiar to this project; it is the ordinary condition of engineering models. Put more politely, each family makes a different representational bargain. The question is not whether a model lies, but what sort of lie it tells, what it lets one see, what it hides, and when the bargain stops paying out. This matters here because the progression from **v0.9** to **v0.15** to **v0.20** is not a simple march from bad model to good model, still less a pious ascent from ignorance to truth. It is a sequence of increasingly disciplined distortions, each useful in a different way, each inadequate in a different way, and each forcing a different, and more detailed, representation of the social field. That is the comparison that matters.

For present purposes, the first of these is **v0.9**, the best current single-basin behavioral family. Its lie is plain and therefore unusually informative. It keeps the entire social field inside one common reproductive vessel, which is already wrong, but wrong in a way that makes the machinery starkly visible. Thresholded hysteresis acting on persistence gives the best current single-basin match to bounded compression plus a partial scar. Weak filtering can compress paternal-line diversity. Fragility can turn thinning into pruning. Hysteresis can give decline memory. Formal and informal channels can diverge. In that sense **v0.9** is the clearest single-basin behavioral demonstration I have. It also clarifies the uncomfortable part. Cleaner recovery-channel interpretations, as in **v0.10**, were more elegant and less adequate; the restrained hybrid in **v0.11** was informative and still inferior. That is an ugly result, and likely closer to history. So **v0.9** buys lucidity at the cost of spatial truth. It lies by keeping everyone in one vessel, but it lies clearly enough to teach.

Then comes **v0.15**, the best current scalar/topological Daisyworld family. Here the bargain changes. The universal bowl is abandoned in favor of several weakly coupled basins, Daisyworld-style, and the global story immediately improves. This is not a decorative borrowing. Daisyworld, after Watson and Lovelock, is a deliberately meager self-regulating toy world in which local reproductive dynamics feed back into global conditions. In this project the point was that several partially bounded reproductive fields already model the problem more finely and honestly than one universal competitive arena. The progression through **v0.12-v0.16** matters because it shows exactly what topology across basins fixes, and what it does not. **v0.12** was the first real topological advance: aggregate bounded compression looked less forced, pre-window global loss was lower, window-dominated loss was clearer, peak pruning moved materially rightward, survivor monopoly weakened, and diversity was preserved more naturally. But basin-specific pruning still occurred too early, and too much of the apparent lag improvement came from superposition rather than from cleanly staggered local transitions. **v0.13** showed that this was not mainly a metric artifact and could not be cured by a simple basin-slack gradient. **v0.14** introduced window-driven local closure and produced the first modest honest hint of local timing separation. **v0.15** then added window-gated local marginality and fragility gain, becoming the first Daisyworld version in which the sustained local pruning metric separated in the right direction, at least partially; closure mattered, local marginality mattered more, basin 1 moved later in a real way, basin 2 still eroded too early. **v0.16** added basin-specific pre-window buffering and improved coarse staggering without solving sustained late-basin delay. So **v0.15** buys archipelago realism at the cost of overly polite interiors. It corrects the lie of the universal bowl, but it still lies about what happens inside the basins. The innermost dolls needed to be painted in greater

detail.

A model can lie clearly, lie obscurely, or lie in exactly the wrong place. The eligibility-graph branch came to be because the wrong place had become harder to ignore. The problem was no longer merely topology across basins. It was internal admissibility inside them. A basin was no longer a smooth interior with an edge. It had become a nested field of permissions, exclusions, weak bridges, choke points, shells, and inner jurisdictions. That is where **v0.20** matters. But to say what **v0.20** adds, one has to say what **v0.17**–**v0.19** were already teaching.

The graph branch did not arrive as triumph, but rather emerged from considering informative failure. **v0.17**, the first eligibility-graph implementation, failed informatively. The graph proved potent, but the baseline world was too harsh, not open enough, and admissibility structure was being imposed too strongly from generation zero. Collapse was front-loaded. That was not a trivial failure. It showed that internal admissibility was a real mechanism family, but also that a graph branch can poison the baseline if it begins already too pinched. **v0.18** then softened the open graph and rehabilitated the branch. That was a real conceptual gain: the graph could matter structurally without poisoning the baseline, bounded compression could remain the main event, and the graph family could become a legitimate mechanism family rather than decorative machinery. **v0.19** improved the branch further by staging tightening inside the graph. One-shot graph switching had been too coarse; phased internal tightening produced more credible sustained local timing and showed that nested local topology matters. But **v0.19** was still too matrix-wide, too administrative, and too uniform in its closure schedule.

That is where **v0.20** earns its place. It does not merely graph harder. It changes what sort of time the graph lives in. The nested-topology eligibility-graph family already carries three basins, three

fixed internal eligibility classes per basin, class-based admissibility graphs inside each basin, and distinct formal and informal traversal of those graphs. What **v0.20** adds is a more truthful schedule of collapse inside that internal matrix. The issue is no longer merely whether the graph tightens, but which routes tighten first, which remain open longest, and how that changing internal connectivity shapes persistence, collapse, and survivor ordering. Matrix-wide staged tightening was still too coarse. Cross-class bridges had to collapse on different schedules. Edge-phased tightening is the first graph version to produce the correct mean sustained basin ordering. That is the gain. The innermost dolls now contain real structure, though one suspects they will not remain innermost for long: fixed classes, distinct traversals, bridges between classes, and edge families whose collapse clocks differ. The representational bargain therefore changes again: more truthful in one place, and perhaps easier to admire too long in another. **v0.20** buys structured admissibility and better behavioral ordering, but at the price of greater complexity and therefore greater temptation to admire subtle machinery for too long. More elaborate lies are not automatically more informative. At this point I became bemused: has the machine learned the world here, or only learned how to please me?

This all also means that the comparison is not merely between three versions, but between three different places where the lie resides. In **v0.9**, the lie is the single vessel. In **v0.15**, the lie moves inward: the larger field has become more truthful, but the basin interiors remain too smooth. In **v0.20**, the interiors are given real structure, and so the remaining lies become subtler. That is progress, but it is not innocence. Richer machinery does not absolve a model. It may simply give it more sophisticated ways to mislead.

That, not chronology, is the comparative point. **v0.9** remains valuable not despite its simplicity, but because it lies clearly. It gives

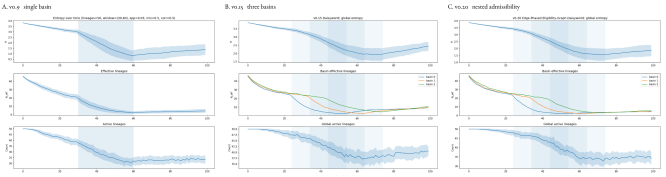

Figure 5. *Model families, truthful lies, and representational bargains.* **A: v0.9**, the best current single-basin behavioral family: one common reproductive vessel, strong mechanistic clarity, wrong room. **B: v0.15**, the best scalar/topological Daisyworld family: three weakly coupled basins, improved aggregate realism, interiors still too smooth. **C: v0.20**, the best nested-topology eligibility-graph family: explicit internal admissibility, edge-phased tightening, and correct mean sustained basin ordering.

the strongest current single-basin behavioral demonstration, and therefore the clearest view of weak filtering, fragility, hysteresis, and signal separation. **v0.15** remains valuable because it proves that topology across basins already changes the global story in nontrivial ways, even though local basin timing still comes too early and interior smoothness remains an unresolved lie. **v0.20** leads because it is the first family to make internal admissibility itself part of the machinery strongly enough to recover the correct mean sustained basin ordering, and because it shows that tightening may need to operate by phased narrowing of admissible connectivity inside basins rather than by scalar pressure alone. These are not three stations on a trolley to enlightenment. They are three different bargains, three different distortions, and three different ways of buying explanatory visibility.

So I do not read this progression as one of mere accumulation. I read it as one of changing representational bargains. The early scalar family told a lie that was still highly informative. The Daisy-

world family told a less rude lie about the larger field while still lying about the interiors. The graph family begins to tell a more truthful lie about the interiors as well, and for that reason earns greater authority, but also greater suspicion, because subtlety in richer machinery is easier to admire for too long. That is why the comparison matters. The question is never whether the model is pure. The question is whether the distortion still earns its keep in understanding the world.

11. Good Enough, and Where to Stop

By this point I had reached the threshold I actually cared about. The original burden was never to solve history in full drag, still less to reconstruct the Neolithic or Bronze Age as though I had been lurking invisibly behind every fence, shrine, and marriage bed. The burden was narrower, more focused, and more useful: to build a model sufficiently good to show that massacre-warfare is not necessary to produce the broad compression pattern in paternal-line diversity. That is a serious claim, and it deserves more than just-so stories. It deserves casual machinery.

I did not originate the thought that massacre might be unnecessary; that possibility was already on the table, most clearly in Zeng, Aw, and Feldman (2018). What I did was build a machine good enough to make that thought analyzable by *mechanism* rather than rhetoric. Weak filtering, thresholded fragility, hysteresis, topology across basins, internal admissibility, and phased narrowing together are now sufficient to show that one does not need massacre-warfare as a necessary precondition for strong lineage compression. The machine is not a proof of peace, and certainly not a denial of violence. Men on horseback did in fact set things on fire. But the older monopoly claim—that something like massacre-warfare must sit at

the causal center—no longer survives contact with a model good enough to do real work.

This is also why the model speaks to the original asymmetry between Y-chromosome compression and the absence of a corresponding mitochondrial collapse. The filter I am modeling acts primarily on male-line persistence. It is not killing populations symmetrically; it is narrowing which paternal lines remain durable under formal order, inheritance, legitimacy, status, and bounded informal leakage. Women remain fully inside the reproductive system, but the lineage channel being selectively pinched is the paternal one. Mitochondrial lines continue through daughters whether or not a particular male line retains status, legitimacy, or enduring patrilineal continuity. And female choice matters. And cheating. This is exactly why one can get strong compression in Y-line diversity without requiring a parallel collapse in mitochondrial diversity. The machine is not explaining away the asymmetry. It is giving it somewhere concrete to live, and runs a model of it.

That, for my original purpose, is enough. More than enough, perhaps. The danger changes at exactly this point. Early in the inquiry, the danger is crudity: the machine is too simple to challenge the obvious story, and so one remains trapped inside pageant by lack of mechanism. Later, once the machine begins to behave, a different danger appears. Once the mechanism begins to work, one is tempted to reward oneself with narrative. This is where epistemic hygiene matters. A functioning model is not permission to become Rousseau before lunch and Swift after dinner.

That temptation deserves contempt. A successful mechanism is not a license for antique ventriloquism. It does not authorize me to populate the Neolithic with invented motives, neat little moral psychologies, noble villagers, sly patriarchs, tragic hearth-scenes, or miniature satirical republics of appetite and fraud. I am under no

obligation to reward a working model by dressing speculation in skins and copper. The machine has earned the right to say something narrower and more important: that a mechanism class exists, that it is sufficient to challenge necessity claims, and that rhetorical certainty in discourse outran causal understanding.

At this point I became bemused in a different way than before. The earlier question was whether the machinery had become subtle enough to please its builder. The later question was what, exactly, I was now entitled to say. Not everything a machine can be made to mimic is something I am entitled to narrate. Not every plausible trajectory deserves a vignette. Not every internally coherent run deserves to be smuggled back into the world as historical personality. The machinery may have learned enough to weaken an older story. It had not thereby licensed me to become a novelist of the Chalcolithic. I snort at myself.

So this is also where heuristics for knowing when to stop begin to look like virtue. There is a point beyond which additional knobs, additional refinements, additional local realism, and additional descriptive ingenuity cease to answer the original burden and begin instead to feed the familiar human appetite for explanatory theater. A model can always be elaborated. The question is whether further elaboration is still buying understanding of the world, or merely buying more ways to admire one's own apparatus. I had already crossed the threshold I set out to cross. The machine was now sufficiently good to make massacre-warfare unnecessary as a premise. That was the point of departure, and I had reached it.

None of this means the work is finished in any grand sense. It means something narrower and much more honorable. The model is now good enough to discharge the original burden, and therefore good enough to justify a change in posture for those in the arcane polis that might care. The next task is not to inflate it into historical

omniscience. The next task is to ask more austere questions: what observables does it actually answer to; what counts as fit; where does it still lie; where should it be distrusted; what would falsify it; what kinds of refinement would still earn their keep? That is a colder light than the temptation of narrative, and very much the right one. But I shall pause.

So no, this does not solve history. It does something better than that. It narrows the field of honest argument—of discourse. It says that massacre-warfare is not necessary. It says that weak, structured, culturally mediated filtering is capable of doing much more than the obvious story had allowed. It says that topology and admissibility matter. It says that internal social structure matters. It says that a machine with enough exposed knobs can make a disputed causal claim analyzable instead of rhetorical. That is already quite a lot. It is also the place where I should stop congratulating myself and return to the world.

12. Observables, Fit, and Falsification

At this point my essay turns towards a colder light. Enough with plausibility theater, and enough with whether a curve feels suggestive after sunrise. The machine is not to be answerable to my mood, nor to the satisfactions of a story that arrives in costume and asks to be admired. It is answerable to *observables.* If it cannot be made to answer there, then it does not matter how elegant, intricate, or psychologically flattering it may become.

The observational burden was accepted from the start. Karmin et al. gave the bottleneck itself: a strong compression in paternal-line diversity, not mirrored by an equal collapse in the maternal channel. Poznik et al. added the later punctuated expansion of surviving male lines. Zeng, Aw, and Feldman then made explicit a sociocultural

mechanism class in which massacre is not required to sit at the causal center. Those are not interchangeable claims. Together, they define a disciplined ensemble. The model has to speak to strong Y compression, to sex asymmetry between Y and mitochondrial inheritance, to bounded timing rather than diffuse erosion everywhere, and to post-bottleneck survivor expansion. If it cannot speak to those, it is not merely uninteresting, it is wrong. From the observables.

This is also where I need to be explicit about Shannon entropy. The sequencing papers do not, of course, speak in that language. They publish the phenomenon in their own native terms: narrowing paternal-line diversity, sex-asymmetric demography, bounded contraction, and later expansion of surviving male lines. I use Shannon entropy here as a collapsed instrument—not the native metric of those papers, but a principled compression of their richer observable pattern into something the machine can track. More specifically, it tracks the distribution of reproductive shares across surviving paternal lines. As those shares concentrate into fewer lines, entropy falls. That is not their wording. It is my bridge from the published observable world to the cybernetic.

So what counts as fit? Not visual mimicry of the published skyline curves, and not the satisfactions of resemblance after a second look and a willing imagination. Those curves are already inferential summaries, produced by their own statistical machinery. My model is not a tracing device. The burden is narrower and cleaner: to recover the same class of observable facts. First, it must generate substantial compression in paternal-line diversity without demanding a corresponding collapse in the mitochondrial channel. The model is not explaining away the asymmetry; it is giving it somewhere concrete to live. Second, that compression should be bounded in time rather than smeared into a long, undifferentiated fog of decline. Third, surviving male lines should be capable of later expansion, because

merely crushing diversity is not enough; the post-bottleneck world matters too. Fourth, once topology enters, one has to care about more than aggregate fit. Local basin timing matters. Internal admissibility matters. Mean sustained basin ordering matters. A model that gets the broad global story roughly right while remaining wrong about the internal ordering may still be useful, but it should not be mistaken for having explicated the deeper problem.

This also means that fit is not whatever feels satisfying after the fact. A model does not get credit merely for producing an evocative curve, or for allowing me to tell a story with decent pacing and atmospheric bronze. Nor does it get credit merely for complexity. More knobs do not confer congruity. More graph does not confer truth. If additional machinery improves only the builder's confidence while leaving the observables slack, then it has not improved fit. It has improved flattery.

So *falsification* matters. In the modest engineering sense relevant here, a family fails if the only way it can produce strong Y compression is by dragging the mitochondrial channel down with it. It fails if it requires grotesque parameter settings to get the effect at all. It fails if the compression cannot be kept bounded and instead dissolves into slow fog. It fails if post-bottleneck survivor expansion cannot emerge. It fails if later topology and admissibility machinery are supposed to improve local ordering but do not. It fails if added structure ceases to improve relation to the observables and begins instead to improve only the machine's ability to please its builder. There are many ways for a model to be wrong. Some are loud. Some are elegant. The elegant ones can be no less wrong.

This is why I am unwilling to let this essay end in narrative intoxication. Once I have a machine that works well enough to challenge an older necessity claim, the next obligation is not to reward myself with historical embroidery. It is to ask what the machine must an-

swer to, how it can fail, and where its distortions still earn their keep in understanding the world. That is not a smaller ambition. It is the only honest one.

13. Reproducibility, the Repo, and What Others Can Now Do

At this point the next honest move is not another flourish of speculation. It is to put the machinery where other hands can inspect it. This is not a tale told over a lantern, asking to be admired for atmosphere. It is a runnable argument. The machine is there. The knobs are there; the distortions are visible; the failures are visible too, which matters just as much.

That is why the repository is part of the essay's ecology rather than an appendix to it. The README, the model files, the source files, the notes on parameter rails, the versioned summaries, the figures, and the entire git history snapshotting each stage are meant to let a reader skilled in the art do more than nod gravely at the prose. They are meant to let that reader rerun the machinery, inspect the progression, compare families, and decide where the model is earning its keep and where it is still cheating. The whole point of exposing the knobs is that one need not take my word for the behavior. One can vary the assumptions and see what survives. You can download it, improve it, and share it with the world.

You can find the repo at: https://github.com/robtow/Patrilines-with-Feedback

This matters especially because the project does not proceed by one immaculate model descending from heaven in finished form. It proceeds by informative failure, revision, comparison, and representational escalation. One version front-loads collapse too hard. Another improves aggregate timing while leaving local timing wrong.

Another rehabilitates a branch that had been born too pinched. Another discovers that interior admissibility matters. Another discovers that matrix-wide tightening is too coarse, and that different bridges must collapse on different clocks. That is not embarrassment. It is the work. The repo is where that work is visible as work rather than mythology. Its kimono is open.

So what can another reader, skilled in the art, now do? At minimum, reproduce the existing figures and inspect whether the claimed family differences are really there. More ambitiously, vary the parameter rails and see whether the reported behavior is robust or precarious. Push harder on the asymmetry between Y-line compression and mitochondrial persistence. Test whether the bounded timing can be maintained under stricter assumptions. Interrogate the basin ordering claims. Add finer internal classes, or different bridge schedules, and see whether they buy understanding of the world or merely buy more machinery. In short: extend, but with discipline.

That last clause matters. I am not handing across a toy chest for free historical embroidery. I am handing across an instrument. Instruments can be abused. One can always add more detail, more local realism, more explanatory furniture, more little historical murmurs in skins and bronze helmets. The machine will not stop you. The observables should. The repo is virtuous if it encourages a harder relation to evidence rather than a softer relation to fantasy.

This is also why the essay has stayed close to code structure and model family progression. The sections are not there merely to decorate the repo with hacker-swagger finery. They are there to help a serious reader orient inside it. Section by section, family by family, the prose says what changed, why it changed, what each bargain bought, and where each lie remained. That makes the code more legible. The code, in return, makes the prose. . . not "just-so".

The work now stands in a form that can be criticized, rerun, tight-

ened, or surpassed. That is as it should be. If the original burden has been met, then the next virtue is not possession but transfer. Hand the tool across the bench. Let the next pair of hands decide whether the machine has really learned something about the world, or only about the habits of the man who built it.

14. What the Machine Has Bought

So what has the machine bought? What have I wrought? Not history solved. Not a transcript of the Neolithic, nor permission to wander through the Bronze Age in borrowed sandals, assigning motives to the dead with the confidence of a man who has mistaken his own tinkered apparatus for omniscience. It has bought something narrower, and, to my mind, more honorable. It has made one disputed causal claim analyzable by mechanism rather than by rhetorically just *saying so*. It has shown that massacre-warfare is not necessary to produce the broad compression pattern in paternal-line diversity. That is a serious gain. Kinda novel.

It has instantiated more than a rebuttal, less than a world. It has shown that weak, structured, and culturally mediated filtering can do far more work than the trad story was willing to allow. It has shown that social topology matters, that basins matter, that internal admissibility matters, that the asymmetry between Y-line compression and mitochondrial persistence can be given somewhere concrete to live, and that the shape of the social field is not decorative but causal. It has shown, too, that I can move from one useful lie to another, from one representational bargain to the next, without pretending that the newest machinery has thereby become innocent.

That is enough for the effort I took up starting at two dark thirty five days ago. More than enough, perhaps. I began because I wanted to know whether the anti-massacre thought could be made to stand

on machinery rather than indignation, intuition, or mere contrariness. It can. I did not originate the thought; instead I built a machine with enough exposed knobs, enough visible failure, and enough hard-won structure to make the thought answerable—falsifiable—in a more serious register. My contribution is bounded, and I am content with its boundedness. So I publish, not in a journal, but via the communities of Open Source.

A larger lesson is methodological. I began provoked by evidence from graves tied to a story too blunt to trust, and the suggestion of a different story. I built a machine simple enough to explore structure. I found that the simplification lied, but lied usefully. I refined, compared, escalated representation, designed the dolls, mapped the admissibility, and kept asking whether the distortions were still earning their keep in understanding the world. I was fortunate; I reached the point at which the original burden had been met, and stopped before the machine became merely an increasing instrument of self-flattery.

This is the conceptual gain I care about most. I have not merely argued against one monopoly story. I have helped move a disputed historical claim out of rhetoric and into mechanism. I have shown that a serious machine can narrow the field of honest argument, and that once it does so, the old certainty no longer gets to swagger about as though it had never been challenged. The point was never to replace one just-so story with another, quieter and more flattering to my own temperament. The point was to make the question answerable in a harder register. To explore kyberos.

That is also why stopping is not a failure of nerve. It is part of the method. Once the original burden has been met, the next temptation is theatrical. I might begin to imagine that the machine has purchased rights to antique interior decoration, that because it can generate a plausible trajectory I may now furnish the dead with mo-

tives, small domestic psychologies, and intimate bronze-lit scenes. No—the honorable act is to stop where the mechanism has earned its keep, hand the tool across the bench, and refuse the temptation to inflate explanatory gain into asserted historical omniscience.

I have no wish to turn that discipline into false modesty—I have done something fun, explicative, and am sharing it; and I enjoyed doing it. The repo exists. The code is there; the notes are there. The figures can be regenerated. Others can rerun it, criticize it, tighten it, extend it, or prove me wrong in some productive direction. That is how such work should stand. A machine that cannot survive the next pair of hands is not much of a machine. Let the polis at it.

I end; the machine has bought me a narrower field of honest argument, a better map of where the older certainty failed, and a disciplined reason to refuse one monopoly story—and that is enough. I turn my recondite studies elsewhere.

15. Appendix. Papers and References

This essay was prompted by a small cluster of papers on the post-Neolithic contraction in Y-chromosome diversity, the asymmetry between paternal and maternal lineage histories, later punctuated male-line expansions, and possible sociocultural mechanisms capable of producing such a pattern without requiring massacre-warfare as a necessary premise. The sources listed here are not decorative authorities; they define the observational and mechanistic problem to which the model is answerable.

The central observational anchor is **Karmin et al. (2015)**, which identifies a recent bottleneck in Y-chromosome diversity coincident with broad cultural change, while not reporting a corresponding collapse in the mitochondrial channel. This is the empirical problem the essay takes as its point of departure.

The second major observational anchor is **Poznik et al. (2016)**, which describes punctuated bursts in human male demography inferred from worldwide Y-chromosome sequences. This is important here not because my model reproduces those published curves literally, but because it helps define the broader observable pattern: contraction of paternal-line diversity followed by later expansion of surviving male lines.

A further regional anchor is **Wang et al. (2013)**, on the Chinese "super-grandfathers,"which offers an especially vivid example of how a small number of surviving paternal lines can later dominate very large descendant populations.

The most important mechanistic precursor is **Zeng, Aw, and Feldman (2018)**, which advances a sociocultural explanation involving patrilineal kin groups and intergroup competition. I did not originate the thought that massacre might be unnecessary. What I have done here is to build a different machine, with different exposed knobs and a more explicitly cybernetic framing, to make that broader possibility answerable by mechanism rather than rhetoric.

My own machine differs from Zeng et al.'s in both scope and emphasis. Their model is a specific anthropological mechanism centered on patrilineal corporate kin groups, female exogamy, and intergroup competition severe enough that group failure becomes lineage failure. Mine is broader, more cybernetic, and less tied to one story about violent group competition. I ask what kinds of coupled reproductive filters, formal bias, fragility, leakage, female choice, topology, and internal admissibility can compress male-line diversity strongly enough to answer the observable pattern. In that sense Zeng et al. show that massacre is not the only game in town; my model family tries to show more generally that one need not begin with violent competition at all in order to make strong Y-line compression mechanistically plausible.

References

Karmin, M., Saag, L., Vicente, M., Wilson Sayres, M. A., Järve, M., Talas, U. G., Rootsi, S., Ilumäe, A.-M., Mägi, R., Mitt, M., Pagani, L., et al. **A recent bottleneck of Y chromosome diversity coincides with a global change in culture.** *Genome Research* 25, no. 4 (2015): 459–466.

Poznik, G. D., Xue, Y., Mendez, F. L., Willems, T. F., Massaia, A., Wilson Sayres, M. A., Ayub, Q., McCarthy, S. A., Narechania, A., Kashin, S., Chen, Y., et al. **Punctuated bursts in human male demography inferred from 1,244 worldwide Y-chromosome sequences.** *Nature Genetics* 48, no. 6 (2016): 593–599.

Wang, C.-C., Yan, S., Qin, Z.-D., Lu, Y., Ding, Q.-L., Wei, L.-H., Tan, J.-Z., Fu, W.-Q., Li, S.-L., Jin, L., and Li, H. **Late Neolithic expansion of ancient Chinese Y chromosomes.** *PLoS ONE* 8, no. 11 (2013): e80514.

Zeng, T. C., Aw, A. J., and Feldman, M. W. **Cultural hitchhiking and competition between patrilineal kin groups explain the post-Neolithic Y-chromosome bottleneck.** *Nature Communications* 9 (2018): 2077.

Originally published on GitHub, 24 March 2026.