

The Geometry of Cellular Possibility

Two Admissibilities: Lineage Restriction, Multimodal Witness, and the History of the Reachable

Flyxion

Independent Researcher

September 2026

Abstract

Two recent papers are ordinarily read as independent contributions to developmental biology and to cell-state profiling, respectively. Jokhai et al. (2026) report that anterior and posterior neural ectoderm (aNE and pNE) arise in parallel during mouse gastrulation and are already lineage-restricted, in the sense that they resist experimentally imposed fates their chromatin landscape does not already foreshadow [1]. Ma and Qu, reporting on Zhang et al.'s RamanOmics platform, describe a multimodal method that pairs label-free Raman chemical imaging with spatial transcriptomics in order to characterize senescent cells, an identity that no single molecular marker or observational modality has previously been able to pin down with confidence [2, 3]. The present essay argues, at greater length and with a more explicit mathematical apparatus than either source requires for its own purposes, that the two results share a common underlying structure: that of a reachable-state space whose geometry changes with accumulated history, indexed by a specified family of interventions, and whose boundary is only ever partially disclosed by any single observational projection. Neither paper proposes the formalism developed below, and neither is treated here as evidence for a theory it was not designed to test; the formalism is offered as an interpretive scaffold that the two results, taken jointly, make unusually natural to construct. The central distinction drawn throughout is between two contractions that are easily conflated in ordinary usage: development contracts a system's own reachable futures relative to a given intervention family, a dynamical or ontic form of admissibility, whereas multimodal observation contracts an observer's hypotheses about a system's present state, an epistemic form of admissibility governed by an entirely different mathematical object. Read together, the two papers motivate a shift in vocabulary and in formal apparatus from the notion of a cell *type*, which specifies a position, to that of a continuation *type*, which specifies a set of admissible trajectories conditioned on an accumulated history and an intervention family, together with a corresponding shift from asking what a single observational channel discloses to asking what remains once several independent witnesses have been reconciled with one another.

Introduction: The Ambiguity of “State”

A cell is routinely represented, for computational and statistical purposes, as a vector

$$x = (x_1, \dots, x_n) \in \mathcal{X}, \tag{1}$$

where the coordinates x_i might denote transcript abundances, chromatin accessibility at particular loci, metabolite concentrations, morphological measurements, or spectral intensities at particular wavenumbers, depending on the assay from which the vector is drawn. This representation carries with it an implicit assumption that is rarely stated explicitly and almost never tested directly, namely that if two cells occupy sufficiently similar observable coordinates at a given time, their relevant biological possibilities going forward are correspondingly similar, so that closeness in $\mathcal{X}$ under some natural metric can stand in for closeness of biological fate. Two recent papers complicate this assumption from what turn out to be opposite directions, and it is the tension between these two directions, rather than either result taken alone, that motivates the formal apparatus developed in what follows. Jokhai et al. show that a cell’s later developmental competence can already be restricted well before the fates that restriction forecloses become visible in any assay ordinarily used to distinguish cell types: anterior and posterior neural ectoderm are molecularly distinguishable from their earliest appearance by the expression of *Otx2* versus *Gbx2*, respectively, and the striking finding of that paper is not that these populations are indistinguishable, but rather that their competence to generate forebrain and midbrain identity, as against hindbrain identity, is already restricted far more sharply than their early molecular distinction alone would suggest, a restriction independently legible in chromatin accessibility well before the later phenotypic divergence it forecloses has occurred [1]. It is the paper’s perturbation experiments, in which each population is deliberately challenged with signals appropriate to the other’s fate, rather than its markers alone, that establish the functional significance of that early distinction, because inappropriate developmental signals are shown, under the tested conditions, to fail systematically at overriding the earlier restriction. Ma and Qu, reporting on Zhang et al.’s RamanOmics platform, establish a structurally different but formally analogous point: that a single observational modality, however precise within its own domain, systematically under-determines what a cell’s present state actually is, since neither Raman chemical imaging nor spatial transcriptomics alone is sufficient to capture senescence, while the conjunction of the two modalities yields improved classification performance relative to either taken in isolation [3]. Read together, the two papers suggest that the ordinary notion of cellular “state” conflates at least four distinct mathematical objects, namely an observed position, an accumulated history, a boundary delimiting what remains reachable from that position and history under a specified family of interventions, and a corresponding set of admissible continuations; the sections that follow separate these four objects formally, and indicate what each of the two papers contributes toward distinguishing them, in a synthesis that should be understood as supplied here rather than as claimed by either paper’s own authors.

1 State, Trajectory, and Two Kinds of Admissibility

The formalism requires, from the outset, a careful separation between an underlying physical state, an observation of that state, and a continuation of that state, since these three objects are of different mathematical type and much confusion follows from allowing a single symbol to do all three jobs at once.

Definition 1. *Let $X_t \in \mathcal{X}$ denote the underlying biological state of a cell at time t , understood as comprising every causally relevant physical degree of freedom, whether or not any of them happens*

to be measured by a given assay. Let $\pi : \mathcal{X} \rightarrow \mathcal{Y}$ denote an observational projection, so that $Y_t = \pi(X_t) \in \mathcal{Y}$ is the quantity an experimenter actually records. Let H_t denote the accumulated developmental history available to the system up to time t , and let U denote a specified family of experimentally realizable interventions. Define the intervention-relative continuation set

$$\Gamma_t(U) = \{ \gamma : [t, T] \rightarrow \mathcal{X} \mid \gamma(t) = X_t, \gamma \text{ is realizable under some } u \in U \text{ given } H_t \}, \quad (2)$$

a set of trajectories through the underlying state space $\mathcal{X}$, and define the corresponding intervention-relative reachable-state set at a later time $\tau \geq t$ as its projection,

$$\mathcal{A}_t(U; \tau) = \{ \gamma(\tau) : \gamma \in \Gamma_t(U) \} \subseteq \mathcal{X}, \quad (3)$$

together with the aggregate reachable set $\mathcal{A}_t(U) = \bigcup_{\tau \geq t} \mathcal{A}_t(U; \tau)$.

Under this definition, $\Gamma_t(U)$ contains continuations and $\mathcal{A}_t(U)$ contains states; a trajectory and a state are never conflated, since the former lives in a path space and the latter in $\mathcal{X}$ itself, and the map from the former to the latter is an explicit projection, evaluation at a point, rather than an implicit identification. The dependence on U is written explicitly rather than suppressed, for a reason that bears directly on how the epistemic and dynamical notions of admissibility developed later are to be kept apart: $\mathcal{A}_t(U)$ is a fact about what the biological system will actually do under each intervention in U , and is in this sense fully ontic, but the question being posed of the system, namely which interventions are being tried, is chosen by the experimenter, so that enlarging U can only enlarge or preserve, never shrink, the reachable set.

Proposition 1. *For intervention families $U_1 \subseteq U_2$ and any $\tau \geq t$,*

$$\mathcal{A}_t(U_1; \tau) \subseteq \mathcal{A}_t(U_2; \tau). \quad (4)$$

This is immediate from the definition, since every trajectory realizable under some $u \in U_1$ is realizable under some $u \in U_2$, but it is worth stating explicitly because it forecloses a natural but mistaken objection to the whole framework, namely that making $\mathcal{A}_t$ depend on an intervention family collapses the dynamical, or ontic, notion of admissibility into a merely epistemic one, on the grounds that what is reachable would then appear to depend on what an observer happens to try. The proposition shows this is not so: the dependence on U is analogous to the dependence of a controllability set on an allowed control set in control theory, where enlarging the admissible controls can only enlarge what is reachable, and there remains, for each fixed U , a perfectly determinate fact of the matter, independent of any observer's beliefs, about which states $T_u(X_t)$ actually occupies for each $u \in U$.

Jokhai et al. supply an unusually concrete empirical illustration of the distinction between X_t , Y_t , and $\Gamma_t(U)$ [1]. During mouse gastrulation, anterior neural ectoderm, marked as $Sox2^+Otx2^+$, and posterior neural ectoderm, marked as $Sox2^+Gbx2^+$, appear simultaneously and in parallel rather than sequentially, the former population being associated in vivo with forebrain and midbrain fate and the latter with hindbrain fate, a correspondence the authors confirm directly through genetic lineage tracing; here the marker expression itself constitutes an observation Y_t , already sufficient to distinguish the two populations under π . The methodologically important move in that work is not this initial marker-based distinction, which alone would say nothing about $\Gamma_t(U)$, but the subsequent perturbation experiments, in which the two populations, differentiated in vitro from human pluripotent stem cells, are deliberately and systematically challenged with an intervention family U consisting of extracellular signals that would ordinarily be sufficient to induce the alternative

regional identity in a naive population of neural progenitors. Under this tested U , posterior neural ectoderm largely fails to acquire forebrain or midbrain identity when supplied with forebrain- or midbrain-inducing signals, and anterior neural ectoderm reciprocally fails to acquire hindbrain identity when supplied with hindbrain-inducing signals; moreover, this mutual resistance persists even when the two populations are co-cultured together, which rules out a simple community-effect explanation in which lineage commitment would depend on being surrounded predominantly by cells of one’s own kind. This experimental design constitutes a direct probe of $\Gamma_t(U)$ and its projection $\mathcal{A}_t(U)$ rather than a mere redescription of Y_t , in that it asks which continuations the cell can undergo under a specified family of active challenges, rather than what identity the cell happens to display when left alone.

This licenses a definition of developmental commitment as a strict contraction of the reachable set over developmental time, for a fixed intervention family U :

$$\mathcal{A}_{t+1}(U) \subsetneq \mathcal{A}_t(U). \quad (5)$$

Differentiation, under this definition, need not correspond to any large displacement in Y_t at the precise moment the contraction occurs. The populations aNE and pNE already differ in identity from their first appearance, as measured by $Y_t = \pi(X_t)$ under any ordinary marker panel; what Jokhai et al. show is that a description confined to this early marker-level distinction alone considerably understates how sharply their developmental competence has already diverged, a divergence in $\mathcal{A}_t(U)$ that remains invisible to an observer relying on Y_t alone until it is specifically probed by an appropriately designed perturbation experiment. Absent such a probe, a cell can lose a substantial portion of its reachable futures well before that loss is detectable in any coordinate an experimenter happens, for reasons of convenience or convention, to be measuring.

The overall argument of this essay can be compressed, at the cost of some precision, into three successive non-equivalences, each of a genuinely different kind rather than three instances of a single failure mode:

$$Y_t \neq X_t, \quad (6)$$

$$X_t \neq \Gamma_t(U), \quad (7)$$

$$\Gamma_t(U) \neq \Gamma_{t+1}(U). \quad (8)$$

Relation (6) is an *observational* non-equivalence, and is the content of the RamanOmics result: an observation is a projection of its underlying object, and a projection need not be injective, so that what any single observational channel reports cannot in general be taken to be the cell’s actual underlying physical state. Relation (7) is a *categorical* non-equivalence, and is the content of the lineage-restriction result: a state and a set of admissible continuations are objects of different logical type entirely, an element of $\mathcal{X}$ against a set of trajectories through $\mathcal{X}$, so that even a fully accurate and complete reading of X_t would not, without further assumptions, disclose $\Gamma_t(U)$. Relation (8) is a *dynamical* non-equivalence, and is the claim this essay adds on top of both of the foregoing: the continuation structure itself is not fixed once and for all at the outset of development but instead possesses its own history, changing as a function of the trajectory the system has already traced through $\mathcal{X}$. None of the three relations requires the considerably stronger and more speculative claim that biological development in general consists of a monotonically and irreversibly shrinking sequence of possibility sets across all timescales and all interventions whatsoever; each instead licenses a comparatively modest and well-supported claim, namely that an observational representation adequate for one purpose, distinguishing cell types at a given moment, need not be adequate for another, predicting which continuations remain available under a specified family of future interventions.

2 Restriction Precedes Appearance

Jokhai et al. go further than merely establishing that $\mathcal{A}_t(U)$ contracts over developmental time; they additionally locate a molecular signature of that contraction well before the phenotypic consequences it forecloses have had the opportunity to arise [1]. Divergent accessible-chromatin landscapes between aNE and pNE, detected using OmniATAC sequencing, are found to be enriched for markedly different regulatory motifs: aNE-enriched accessible chromatin is enriched for OTX2 binding motifs, whereas pNE-enriched accessible chromatin is enriched for HOX, HNF1B, and MAFB motifs, together with motifs associated with the downstream transcriptional effectors of FGF signaling, such as ETV1, and of retinoic-acid signaling, such as RARB and RXRG. When each population is subsequently challenged in vitro with the inducing signal appropriate to the other, it is the chromatin landscape itself, and not merely transcript output as measured by ordinary RNA sequencing, that resists reprogramming: pNE supplied with forebrain-inducing signals fails to gain chromatin accessibility at forebrain-associated loci such as *SIX3* and instead gains accessibility preferentially at hindbrain-associated regulatory elements, and the symmetric relation holds when aNE is supplied with hindbrain-inducing signals. The restriction that Jokhai et al. document is therefore inscribed, at least in part, at the level of what the genome as a whole will admit becoming chromatin-accessible under a given signaling regime, rather than being confined to the level of which genes happen already to be transcribed at any particular time t ; chromatin accessibility here functions as a candidate present-time physical carrier of developmental history, a point taken up again in Section 5.

A conventional developmental model, of the kind implicit in most single-cell trajectory-inference methods, tracks a single sequence of observed states,

$$Y_0 \rightarrow Y_1 \rightarrow Y_2 \rightarrow \cdots, \quad (9)$$

representing the observations a given cell, or a given cell’s descendants, actually yield over the course of an experiment. This essay’s formalism instead accompanies that sequence with a second, coupled sequence of reachable-state sets,

$$\mathcal{A}_0(U) \supseteq \mathcal{A}_1(U) \supseteq \mathcal{A}_2(U) \supseteq \cdots, \quad (10)$$

a nested, decreasing family rather than, despite the superficial resemblance, a filtration in the technical sense of probability theory, where that term denotes an *increasing* family of σ -algebras; no measure-theoretic machinery is imported here, and none is needed. What the biology in fact establishes is simply that the trajectory $\{Y_t\}$ is accompanied by a second trajectory, not of points but of admissible sets, and that this second trajectory develops monotonically only in the comparatively weak sense Jokhai et al. actually demonstrate experimentally: not that every conceivable future is permanently and irrevocably foreclosed under every possible intervention family whatsoever, but that under every U the authors specifically tested, the earlier restriction was found to hold.

This motivates separating out three objects that a single observed state vector conflates by construction, namely

$$\text{observation: } Y_t, \quad \text{internal constraint: } C_t, \quad \text{reachable future: } \mathcal{A}_t(U), \quad (11)$$

where C_t denotes whatever internal regulatory or chromatin-level structure is responsible, at a given moment, for determining which continuations of X_t remain realizable under U . It is worth situating this move relative to the older metaphor it revises. Waddington’s epigenetic landscape pictures development as a ball rolling through a fixed topography of valleys and ridges laid down

once and for all at the outset [4], and later treatments of that picture have made explicit what the metaphor leaves implicit, namely that the landscape itself, and not merely the ball’s position on it, is what should be understood to change over developmental time, with cell fate and cellular potential correspondingly distinguished as separate objects rather than read off a single trajectory [5]. The admissible-set formalism developed here can be read as a further sharpening of that same revision: not merely a landscape whose contours shift, but a state space whose currently accessible region $\mathcal{A}_t(U)$ contracts, an object with an explicit set-theoretic and topological structure rather than a suggestive picture. Two cells can accordingly be close to one another under the observational projection π while already possessing sharply divergent reachable sets, a relation that may be written

$$\pi(X_t^{(1)}) \approx \pi(X_t^{(2)}) \quad \text{while} \quad \mathcal{A}_t^{(1)}(U) \not\approx \mathcal{A}_t^{(2)}(U), \quad (12)$$

where the approximate-equality symbol on the left denotes closeness under whatever metric the projection π carries, and its negation on the right denotes substantial divergence under the reachable-set metric made precise in Section 8. Under this view, differentiation is better described mathematically as the progressive deletion of allowable continuations from an initially larger set than as the displacement of a single point toward some fixed destination fixed in advance. The destination metaphor presupposes, often without acknowledging that it does so, that what is being tracked by a developmental model is observed position alone; the deletion metaphor instead tracks the shrinking of an option set, of which observed position is only one, comparatively impoverished, visible trace.

3 A Cell Has No Privileged Coordinate

Ma and Qu, describing Zhang et al.’s RamanOmics platform, address a problem that is complementary, in a precise sense to be made explicit below, to the one addressed by Jokhai et al.: namely that no single observational modality captures more than a fragment of a cell’s complex underlying state, and that no universal molecular marker unambiguously identifies senescent cells across the heterogeneous range of tissues in which senescence is studied [2, 3]. Increased expression of the cell-cycle inhibitors p21 and p16, positive staining for senescence-associated β -galactosidase, accumulation of unrepaired DNA damage, elevated markers of oxidative stress, and the acquisition of a senescence-associated secretory phenotype are each, individually, partial and imperfect indicators, precisely because senescence is not governed by a single molecular switch that any one assay could hope to read out directly, but instead emerges from coordinated alterations distributed across epigenetic, transcriptional, proteomic, and metabolic layers of cellular organization simultaneously. RamanOmics addresses this by combining hyperspectral, label-free Raman chemical imaging with single-nucleus RNA sequencing and STARmap-based spatial transcriptomics on the very same tissue sections, so that the modalities can be registered against one another at single-cell resolution [2]. Raman imaging supplies biochemical information, arising from intrinsic molecular vibrational modes, that transcriptomics does not and cannot directly measure, while conversely Raman spectra alone do not, by themselves, disclose which genes are being transcribed, nor do they establish that a given cell occupies a senescence-associated transcriptional state as conventionally defined. Each modality is genuinely informative within its own domain, and each is, taken alone, insufficient to establish senescence with confidence.

Formalize the two modalities as projections from the common underlying state X onto two distinct, and only partially overlapping, observational spaces,

$$\pi_R : \mathcal{X} \rightarrow \mathcal{R}, \quad \pi_T : \mathcal{X} \rightarrow \mathcal{T}, \quad (13)$$

corresponding respectively to Raman and to transcriptional observation. Neither projection is assumed to be injective, and indeed both are known empirically not to be, so that

$$\pi_R(x_1) = \pi_R(x_2) \not\Rightarrow x_1 = x_2, \quad (14)$$

with the analogous statement holding for π_T by an identical argument. Multimodal observation is therefore not merely “more data” in a naive additive sense in which two partial measurements are simply concatenated and left otherwise unrelated; rather, it reduces observational degeneracy in a specific and formally tractable way, by replacing either projection individually with the joint, or product, map

$$\Pi(x) = (\pi_R(x), \pi_T(x)) \in \mathcal{R} \times \mathcal{T}, \quad (15)$$

whose fibers $\Pi^{-1}(\Pi(x))$ are generically smaller than, and in fact contained within, the fibers of either π_R or π_T taken alone, since $\Pi^{-1}(y_R, y_T) = \pi_R^{-1}(y_R) \cap \pi_T^{-1}(y_T)$. This interpretation is consistent with the empirical result, reported by Zhang et al. and discussed by Ma and Qu, that a machine-learned multimodal classifier, or “barcode,” combining Raman spectral features with gene-expression patterns extracted from the same cells achieves improved classification performance in identifying senescent cells relative to models constructed from either modality taken in isolation [2, 3]. It should be emphasized that π_R and π_T are arbitrary observational maps rather than linear operators, so that the improvement is not evidence for any difference in “null space” in the technical sense that term carries for linear maps; what the improvement is consistent with, more cautiously, is that the two modalities resolve different ambiguities in cellular state, such that their joint observation discriminates states that either modality alone leaves less well separated, and that senescent cellular identity, as an object of biological description, is correspondingly better approximated by their conjunction than by either projection on its own.

4 Identity as an Intersection of Constraints

The preceding sections motivate a definition of cellular identity that does not depend on the existence of any single privileged marker function $m : \mathcal{X} \rightarrow \{0, 1\}$ separating one cell type or cell state from another. Instead, one may define identity through the intersection of the constraints imposed by a collection of distinct and mutually irreducible observational projections,

$$\mathcal{C}(x) = \mathcal{C}_R(x) \cap \mathcal{C}_T(x) \cap \mathcal{C}_E(x) \cap \cdots, \quad (16)$$

where $\mathcal{C}_R(x)$, $\mathcal{C}_T(x)$, and $\mathcal{C}_E(x)$ denote, respectively, the sets of candidate underlying states consistent with the observed Raman signal, with the observed transcriptional profile, and with observed chromatin-accessibility or epigenetic data, and where the ellipsis indicates that further independent observational modalities, such as lineage history, spatial tissue context, or experimentally measured developmental competence in the sense of Section 1, each further eliminate candidate descriptions of what the cell in question could actually be. Cellular identity, under this construction, becomes progressively constrained by the accumulation of independent projections rather than being read off directly from any single privileged coordinate axis of $\mathcal{X}$.

This construction yields a duality that the two papers under discussion illustrate from what are, in an important sense, opposite sides, and which the remainder of this essay is largely concerned with keeping carefully unconfuted. Write $\mathcal{H}$ for the set of biological hypotheses, understood loosely as candidate descriptions of a cell’s actual underlying state, that remain consistent with everything an observer has so far measured about a given cell, so that additional independent measurements contract $\mathcal{H}$ according to

$$\mathcal{H} \supset \mathcal{H}_T \supset \mathcal{H}_{T,R}, \quad (17)$$

where $\mathcal{H}_T$ denotes the hypotheses remaining after transcriptional observation alone and $\mathcal{H}_{T,R}$ denotes those remaining after the further addition of Raman observation; formally, $\mathcal{H}_T$ is the fiber $\pi_T^{-1}(\pi_T(x))$ narrowed by prior belief, and $\mathcal{H}_{T,R}$ the further narrowing $\Pi^{-1}(\Pi(x))$ of Section 3. By contrast, write $\mathcal{A}_t(U)$, as in Section 1, for the set of states that remain dynamically accessible to the biological system itself under intervention family U , independently of what any observer happens to have measured, so that developmental history contracts $\mathcal{A}_t(U)$ according to

$$\mathcal{A}_0(U) \supset \mathcal{A}_1(U) \supset \mathcal{A}_2(U) \supset \cdots, \quad (18)$$

exactly as in Section 2. RamanOmics, on this account, primarily improves an observer’s estimate of a cell’s present state by contracting $\mathcal{H}$, the set of biological descriptions still compatible with what has actually been observed of that cell so far. Lineage commitment, by sharp contrast, changes the cell’s own actual reachable futures by contracting $\mathcal{A}_t(U)$, a fact about the biological system itself that holds quite independently of whether any observer is present to measure it at all. This can be summarized by placing the two forms of admissibility on opposite sides of a single inference diagram,

$$H_t \longrightarrow X_t \xrightarrow{\pi} Y_t, \quad (X_t, U) \longrightarrow \Gamma_t(U), \quad (19)$$

in which RamanOmics acts on the rightward, observational, arrow by replacing a single projection π with a joint witness Π , contracting the epistemic object $\mathcal{H}$, while Jokhai et al. act on the downward, dynamical, arrow by perturbing X_t under a specified U and reading off the resulting $\Gamma_t(U)$, contracting the ontic object $\mathcal{A}_t(U)$. Treating these two contractions, of $\mathcal{H}$ and of $\mathcal{A}_t(U)$, as instances of a single underlying phenomenon produces precisely the category mistake this essay is constructed to dissolve: a well-witnessed cellular state, one for which $\mathcal{H}$ has been narrowed to a small set by extensive multimodal observation, is not thereby a state possessing more open developmental futures, and conversely a lineage-restricted cell, one for which $\mathcal{A}_t(U)$ has been substantially narrowed by developmental history, is not thereby a particularly well-witnessed one; the two forms of narrowing are logically and empirically independent of each other, a point to which the essay’s closing formulation returns.

5 History Is a Coordinate

The preceding sections should not be read as asserting the strong metaphysical claim that no conceivable instantaneous physical state could, in principle, encode its own history; that claim would be considerably stronger than the biology requires and is not the one defended here. If X_t is defined as an exhaustive state containing every causally relevant degree of freedom, then by construction

$$P(X_{t+\Delta} \mid X_t, H_t) = P(X_{t+\Delta} \mid X_t), \quad (20)$$

so that history need not be supplied as a separate coordinate once X_t is taken to be sufficiently rich. The more precise and considerably better-supported claim that the biology in this essay actually establishes concerns not X_t itself but the observational representation $Y_t = \pi(X_t)$ ordinarily available to an experimenter, namely that this reduced representation can fail to be a sufficient statistic for future response to intervention:

$$P(X_{t+\Delta} \mid Y_t, u, H_t) \neq P(X_{t+\Delta} \mid Y_t, u). \quad (21)$$

When this inequality holds for a given u , the measured Y_t is not sufficient for predicting the cell’s response to that intervention, and H_t carries genuine predictive information that the projection π has discarded. “History is a coordinate,” understood in this qualified sense, becomes a claim

about non-Markovianity relative to a chosen observational representation, rather than a claim that the underlying physical dynamics themselves fail to be Markovian; the world need not literally remember an abstract H_t as a separate variable, since chromatin structure and other persistent regulatory configurations can embody history within the present physical state X_t while remaining entirely omitted from a reduced observational coordinate $Y_t = \pi(X_t)$.

Jokhai et al. supply exactly the kind of present-time physical carrier this qualified claim requires [1]: cells exposed to developmentally inappropriate later signals do not simply become whatever those signals would ordinarily induce in an otherwise naive population, precisely because the cell’s prior developmental history has already altered, at the level of chromatin structure, what the same subsequent inductive cue is capable of accomplishing. Chromatin accessibility here is not a separate metaphysical variable H_t standing outside the physical state; it is part of X_t itself, one that an observational coordinate confined to gene-expression markers such as *Otx2* or *Gbx2* can fail to register even though it is fully present in the cell. In terms consonant with a HOLOcoord-style formalism, the biologically useful coordinate for describing a cell is accordingly not Y_t taken in isolation, but rather the extended tuple

$$\text{HC}(x_t) = (Y_t, H_t, \mathcal{A}_t(U), \Gamma_t(U)), \quad (22)$$

comprising an observation together with the history, the reachable set, and the continuation structure that history has, over developmental time, left behind, with the understanding that H_t here names predictive information omitted by π rather than a component of X_t that could not in principle have been included in a richer observational coordinate. The authors of that paper are, notably, appropriately cautious on the further question of permanence: they explicitly leave open the possibility that other interventions not included in their tested U , such as direct manipulation of the transcription factors or signaling pathways, including *Otx2* and *Fgf8*, that specify aNE versus pNE identity during their initial emergence at gastrulation, might yet permit interconversion at that earlier developmental stage, even though aNE and pNE are shown not to interconvert once fully established under the U actually tested. The set $\mathcal{A}_t(U)$, as used throughout this essay, is accordingly a claim about what has been demonstrated to be reachable under the specific intervention family actually tested, and not a claim about what could never, under any conceivable future intervention outside that family, be reached at all.

6 The Boundary Is the Biological Object

Suppose the cell’s nominal, or theoretically possible, state space is $\mathcal{X}$, while its currently realizable states, under intervention family U , form the proper subset $\mathcal{A}_t(U) \subseteq \mathcal{X}$ introduced above. Differentiation can then be represented, not primarily by the changing position of the cell within the ambient space $\mathcal{X}$, but by the deformation over time of the topological boundary

$$\partial\mathcal{A}_t(U) \subset \mathcal{X}. \quad (23)$$

A given developmental signal matters, on this account, according to whether it is capable of moving the system across that boundary, and not according to how far it moves the system through $\mathcal{X}$ as measured in some absolute or extrinsic sense.

Considerable care is required, however, in stating what a single perturbation experiment actually establishes about this boundary, since a failed transition constrains $\mathcal{A}_t(U)$ without, by itself, identifying any particular point of $\partial\mathcal{A}_t(U)$.

Proposition 2. *Let $x \in \mathcal{A}_t(U)$ denote the cell’s present realizable state, and let $u \in U$ denote an experimentally specified intervention intended to induce a target state y_u . If the transition to y_u succeeds, then y_u is evidence for membership of y_u in the empirically reachable set under the tested intervention family, that is, $y_u \in \mathcal{A}_t(U)$. If the transition reproducibly fails, the result constrains $\mathcal{A}_t(U)$ by excluding y_u as a member reachable via that particular tested transition under those specific conditions. Neither outcome alone identifies a point of the topological boundary $\partial\mathcal{A}_t(U)$; localization of the boundary requires a sufficiently resolved family of neighboring interventions or neighboring target states, together with appropriate continuity or regularity assumptions relating nearby interventions to nearby outcomes.*

This formulation distinguishes boundary *evidence* from boundary *localization*, a distinction the earlier and looser statement of this proposition collapsed. A single failed transition is genuine evidence, since it excludes a candidate member of $\mathcal{A}_t(U)$, but it does not by itself locate $\partial\mathcal{A}_t(U)$ any more than a single point known to lie outside a region locates that region’s boundary without some further assumption about how nearby points behave. This is nonetheless a biological instance of a considerably more general principle developed elsewhere in this author’s work under the heading of refusal as distinct from absence: if a cue u normally drives a transition $x \rightarrow y$ in an unrestricted population, but a lineage-restricted cell refuses that transition when the same cue is applied to it, the refusal does not constitute an absence of information about that cell, but rather constitutes a genuine constraint on $\mathcal{A}_t(U)$, one that a sufficiently dense and systematically sampled family of such refusals, together with their successful counterparts, can jointly triangulate into an increasingly precise picture of where $\partial\mathcal{A}_t(U)$ actually lies. A failed transition is accordingly not an absence of data to be discarded from subsequent analysis, but a positive, if individually incomplete, measurement, on exactly the same evidential footing as a successful one, and it is only the accumulation of many such measurements, appropriately structured, that converts evidence of this kind into a localized estimate of the boundary itself, a point developed further in Section 8.

7 The Residual Carries the History

It is at this point that RamanOmics becomes considerably more than merely an improved assay for detecting senescent cells. Lung and skin p21-positive cells are reported to display markedly tissue-specific transcriptional programs, related in the lung to extracellular-matrix remodeling and to transforming growth factor- β signaling, as reflected in altered expression of *Serpine1*, *Dab2*, and *Igfbp7*, and related in the skin instead to keratinization and to epidermal barrier function, as reflected in altered expression of *Krt10*, *Lor*, and *Sbsn*; yet Zhang et al. nonetheless found, as Ma and Qu report, a shared enrichment of lipid-associated Raman spectral features across these two otherwise transcriptionally divergent populations [2, 3]. Rather than treating this discrepancy between the two modalities as measurement noise to be averaged away in the course of ordinary data analysis, one can instead ask whether it reflects biological information systematically omitted by one representation and retained by the other.

Stating this precisely requires some care, since the full underlying state x is, by the argument of Section 3, not itself directly observable; only its two projections $\pi_R(x)$ and $\pi_T(x)$ are available. The residual is accordingly better defined relative to a specific pair of representations rather than relative to the unobserved x itself. Define the transcriptomically unpredicted portion of the Raman signal as

$$r_{R|T} = \pi_R(x) - \widehat{\pi_R(x) \mid \pi_T(x)}, \quad (24)$$

where $\widehat{\pi_R(x) \mid \pi_T(x)}$ denotes the best prediction of the Raman projection obtainable from the

transcriptomic projection alone, and define the converse residual symmetrically,

$$r_{T|R} = \pi_T(x) - \widehat{\pi_T(x) | \pi_R(x)}. \quad (25)$$

In general $r_{R|T} \neq r_{T|R}$, since each residual is defined relative to a different predicting representation; there is accordingly no single, privileged notion of “the” residual left over by observation, only a family of representation-relative residuals, and asking what a given representation leaves out is not a well-formed question until the representation doing the leaving-out has been specified. Raman chemistry, under this reading, functions as one possible instrument for measuring precisely what a transcriptome-only coordinate system leaves in $r_{R|T}$: not noise scattered around an otherwise adequate transcriptional prediction, but a partially independent observational channel carrying its own history, one that may reflect a form of broader lipid remodeling common to senescent cells across tissue types, with consequences for membrane properties, for signal transduction, or for the senescence-associated secretory phenotype itself.

A detected biochemical signature is not, however, thereby an identified molecule, and an identified molecule is not thereby a demonstrated causal mechanism; these three levels of description, though frequently elided in casual discussion of such findings, remain logically distinct from one another and require independent forms of evidence to establish. Ma and Qu are explicit on precisely this point, noting that Raman spectra alone do not straightforwardly identify individual chemical species without recourse to complementary techniques such as mass spectrometry or isotope tracing, and further noting that whether the biochemical alterations detected by RamanOmics merely accompany senescence as an epiphenomenon, or instead actively contribute to its downstream functional consequences, remains at present an entirely open empirical question. In terms of a witness-before-mechanism distinction developed more fully elsewhere, one may write

$$\text{detectable signature} \not\Rightarrow \text{identified molecule} \not\Rightarrow \text{causal mechanism}, \quad (26)$$

where each arrow’s negation records a genuine and, on present evidence, unclosed inferential gap rather than a merely rhetorical qualification. RamanOmics improves the available witness to senescent cellular states very considerably, without thereby licensing any automatic promotion of that improved witness into a demonstrated mechanism; a representation of this kind can faithfully transport a property that is genuinely available through its own measurement channel without thereby transporting every property, and in particular without transporting every causal fact, that pertains to its underlying source.

8 From Cell Type to Continuation Type

The formal apparatus developed across the preceding sections supports a cautious terminological, and ultimately conceptual, shift. A cell type, in the ordinary usage of that term, answers approximately the question “what is this cell?”, a question addressed to the observation Y_t or to some marker function defined upon it. A continuation type instead answers a structurally different question, already formalized in Section 1 as $\Gamma_t(U)$: what continuations remain available from this position, under a specified intervention family, given everything that has already happened to bring the cell to that position. Two cells with similar or even identical appearance under Y_t , but with substantially different associated continuation sets $\Gamma_t(U)$, are biologically different from one another in a way that ordinary state-based classification is structurally incapable of disclosing; conversely, two cells that differ considerably under Y_t might nonetheless share a common continuation structure,

and so be, in the relevant biological sense, more alike than any single-coordinate classifier operating on Y_t alone would ever report them to be.

This proposed shift does not discard the ordinary notion of cell type as a useful descriptive label, and it is not intended as a wholesale replacement for existing taxonomic practice; rather, it relocates what such a label is properly supposed to be tracking. A classification constructed from Y_t alone is, in the terms developed here, a snapshot of instantaneous observed position, and it should be said plainly that prior work has not simply ignored this limitation. Pseudotemporal ordering methods reconstruct a continuous trajectory from a collection of such snapshots taken across many cells [7], and RNA velocity goes further by estimating, from the same static single-cell measurement, a direction of future transcriptional change based on the ratio of unspliced to spliced transcripts [8]; more recent probabilistic approaches such as Palantir assign each cell a distribution over terminal fates and use an entropy-like quantity to measure cellular plasticity from population structure alone [9]. The continuation-type proposal made here needs to distinguish itself from this considerable body of prior work rather than presenting itself as a simple corrective to a snapshot view no one actually holds. The relevant distinction is that a fate probability of this kind is inferred from the observed structure of a population, whereas the continuation set $\Gamma_t(U)$ of this essay is constrained by direct experimental intervention:

$$\text{inferred fate probability} \neq \text{experimentally probed reachability}. \quad (27)$$

A classification that instead incorporates the continuation set $\Gamma_t(U)$, established this way, is a claim about trajectory and about admissible future trajectory in particular, and it is precisely this latter, more demanding claim that both of the papers under discussion, in their respective and rather different ways, are actually engaged in testing whenever they challenge cells experimentally with inappropriate signals, in the case of Jokhai et al., or combine independent observational modalities against one another, in the case of Ma and Qu’s account of RamanOmics.

The distinction between these two notions of classification has empirical, and not merely conceptual, consequences, which can be made precise by equipping continuation sets with a metric.

Definition 2. Let d_X be a metric on the observed-state space $\mathcal{X}$. For trajectories γ, η sharing a common domain $[t, T]$, define the uniform distance

$$d_\infty(\gamma, \eta) = \sup_{s \in [t, T]} d_X(\gamma(s), \eta(s)), \quad (28)$$

noting that if biologically meaningful trajectories can proceed at different rates, d_∞ can register two otherwise similar developmental paths as far apart merely on account of differing timing, in which case a reparameterization-invariant or Skorokhod-type distance would be the more appropriate choice; no such refinement is required for the purposes of this essay. Define the corresponding Hausdorff-type distance between two continuation sets $\Gamma_i = \Gamma_i(U)$ and $\Gamma_j = \Gamma_j(U)$ by

$$d_\Gamma(\Gamma_i, \Gamma_j) = \max \left\{ \sup_{\gamma \in \Gamma_i} \inf_{\gamma' \in \Gamma_j} d_\infty(\gamma, \gamma'), \sup_{\gamma' \in \Gamma_j} \inf_{\gamma \in \Gamma_i} d_\infty(\gamma, \gamma') \right\}. \quad (29)$$

Proposition 3. The metrics d_X and d_Γ are, in general, independent of each other, in the sense that neither

$$d_X(x_i, x_j) < \varepsilon \implies d_\Gamma(\Gamma_i, \Gamma_j) < f(\varepsilon) \quad (30)$$

for any monotonic function f , nor the converse implication, holds as a matter of biological necessity.

The empirical content of the aNE/pNE result may be read as an instance in which this proposition’s first implication fails: two populations that have only recently diverged from definitive ectoderm, and that already exhibit distinct early molecular identities under Y_t , nonetheless prove, once probed by perturbation, to be far more sharply separated in developmental competence than a description confined to their present regional markers alone would suggest, since their admissible continuation sets $\Gamma_t(U)$, established experimentally, prove almost entirely disjoint. A state-based taxonomy relying on Y_t alone would understate how different these populations already are; a continuation-based taxonomy correctly identifies them as importantly, and consequentially, different. The converse failure, in which populations differing considerably under Y_t nonetheless share a common continuation structure, remains at present a live empirical possibility rather than a demonstrated fact, and its investigation constitutes an empirical question in its own right: does classification by reachable-set geometry, in the sense formalized above, predict subsequent biological behavior more accurately than classification by instantaneous molecular state alone?

This question suggests a research program that extends somewhat beyond either paper as actually published, though considerable care is required in specifying exactly what such a program can and cannot directly measure. Rather than asking only what state a given cell currently occupies, one may apply a controlled family of interventions $u \in U$ and record, for each such intervention and each of a specified family of operationally defined target phenotypes $y \in Y$, whether the resulting transition succeeds, according to a response surface

$$\chi_t : U \times Y \rightarrow \{0, 1\}, \quad (31)$$

a matrix that an experimental design of the Perturb-seq type, combined with lineage or fate readouts, could in principle sample systematically. The directly measurable object arising from such a screen is the empirical response region

$$R_t = \{ (u, y) \in U \times Y : \chi_t(u, y) = 1 \}, \quad (32)$$

which, unlike the state-space boundary $\partial\mathcal{A}_t(U)$ itself, is an object living entirely within the observable product space $U \times Y$ and is therefore genuinely measurable by the experiment as designed. It is only after imposing a further model connecting interventions and observed target phenotypes to the underlying state space, of the kind sketched in Definition 1’s map $\gamma \mapsto \gamma(\tau)$, that one may infer anything about $\mathcal{A}_t(U)$ or its topological boundary from R_t ; the inference proceeds through a hierarchy of successive levels, from perturbation outcomes, to the empirical response geometry R_t they jointly trace out, to an inferred reachable set $\mathcal{A}_t(U)$ once a model of χ_t has been fit, to an admissibility boundary $\partial\mathcal{A}_t(U)$ only once that inferred set has itself been examined for its own topological boundary. This hierarchy mirrors, at the level of experimental design, the same discipline insisted upon in Section 7 with respect to detectable signature, identified molecule, and demonstrated mechanism: the framework refuses, at every stage, to infer its own theoretical boundary directly from raw empirical witness, requiring instead that each successive level of inference be earned by an explicit modeling step rather than assumed by definitional fiat.

9 Admissibility Relative to Intervention

Section 1 introduced $\mathcal{A}_t(U)$ as depending explicitly on an intervention family U rather than suppressing that dependence, and the consequences of that choice deserve to be drawn out more fully before the essay proceeds further. Reachability, properly stated, is never reachability simpliciter; it is always reachability under some specified family of allowed interventions. The question “can

this cell become y ?” is therefore incompletely typed as it stands; the properly typed question is whether the cell can become y under a stated U , since the answer can differ arbitrarily as U is enlarged or restricted. One may illustrate this with an informal hierarchy of increasingly powerful intervention families, such as

$$U_{\text{phys}} \subseteq U_{\text{signal}} \subseteq U_{\text{molecular}} \subseteq U_{\text{genetic}}, \quad (33)$$

ranging from passive physical perturbation, through the extracellular signaling manipulations Jokhai et al. actually employ, to direct molecular or genetic intervention on the transcription factors and signaling pathways governing lineage specification; this particular chain of inclusions is offered only as an illustration of the general point, not as a claim that these specific families are universally ordered in exactly this way across every developmental system.

The chief conceptual payoff of making U explicit is that it separates restriction from irreversibility, two notions the earlier, looser formulation of $\mathcal{A}_t$ risked running together. A cell can be lineage-restricted relative to ordinary developmental signaling, in the sense that $\mathcal{A}_t(U_{\text{signal}})$ has contracted sharply, while remaining transformable under direct genetic manipulation, in the sense that $\mathcal{A}_t(U_{\text{genetic}})$ has not; there is no contradiction here, only two different questions being asked of the same system, and the monotonicity result of Section 1 guarantees that $\mathcal{A}_t(U_{\text{signal}}) \subseteq \mathcal{A}_t(U_{\text{genetic}})$ whenever $U_{\text{signal}} \subseteq U_{\text{genetic}}$. Restriction, in this vocabulary, is always a defensible claim relative to an experimentally specified U ; irreversibility, by contrast, would quantify over an enormously stronger and in practice untestable class of interventions, one that neither paper under discussion claims to have exhausted. This is precisely the caution Jokhai et al. themselves express when they leave open the possibility that interventions beyond those tested, such as direct manipulation of the transcription factors specifying aNE versus pNE identity, might permit interconversion at an earlier developmental stage [1]; in the present vocabulary, their claim is exactly that $\mathcal{A}_t(U_{\text{signal}})$ has been shown to exclude interconversion, together with an explicit refusal to claim anything about $\mathcal{A}_t(U_{\text{genetic}})$.

10 Biological Memory Without a Historical Variable

Section 5 established that the qualified claim actually required by the biology is not that X_t must be supplemented by a metaphysically independent variable H_t , but that the observational representation $Y_t = \pi(X_t)$ can fail to be a sufficient statistic for future response. That claim can be sharpened into an operational criterion using conditional mutual information. History H_t is operationally relevant relative to a representation Y_t , for a given intervention u , whenever

$$I(Y_{t+\Delta}; H_t \mid Y_t, u) > 0, \quad (34)$$

where $I(\cdot; \cdot \mid \cdot)$ denotes conditional mutual information; a positive value here means that knowing H_t , in addition to Y_t and u , genuinely reduces uncertainty about the future observation $Y_{t+\Delta}$, and a value of zero means Y_t already screens off H_t from the future given u , so that history, relative to this particular representation and this particular intervention, carries no further predictive content. This gives biological memory a precise operational meaning without requiring the stronger and less defensible claim that memory is an ingredient the physical world must literally carry alongside its present state: memory, in this sense, is simply history rendered predictively relevant by the compression a given observational representation performs on the underlying state.

The chromatin results of Jokhai et al. supply a concrete illustration of exactly this phenomenon [1]. The divergent accessible-chromatin landscapes of aNE and pNE are associated with, and

plausibly causally responsible for, their later developmental competence, so that something which appears as “history dependence” under a coarse representation confined to a handful of marker genes corresponds, under a richer representation that includes chromatin accessibility, to perfectly present physical structure carrying no residual dependence on anything not already encoded in X_t . This gives the HOLOoord-style coordinate introduced in Section 5 a more precise role than merely appending H_t as an additional slot: a useful coordinate system is not necessarily one that remembers the entire developmental past in full, but one that retains enough of the present physical consequences of that past, such as chromatin accessibility, to recover the relevant continuation structure $\Gamma_t(U)$ without further reference to H_t as an independent object.

11 Observational Equivalence and Dynamical Inequivalence

The two papers under discussion can be unified at a somewhat deeper mathematical level by observing that every observational map induces an equivalence relation on the underlying state space. Given $\pi : \mathcal{X} \rightarrow \mathcal{Y}$, define

$$x \sim_\pi x' \iff \pi(x) = \pi(x'), \quad (35)$$

so that an observer never works directly with $\mathcal{X}$ but only with the induced equivalence classes $[x]_\pi = \{x' \in \mathcal{X} : \pi(x') = \pi(x)\}$. RamanOmic, in this vocabulary, refines a partition rather than simply adding data in the abstract: transcriptomic observation induces one equivalence relation $\sim_T$, Raman observation induces another, $\sim_R$, and their conjunction induces the finer relation whose classes are

$$[x]_{T,R} = [x]_T \cap [x]_R, \quad (36)$$

which is a precise restatement, in the language of equivalence classes rather than fibers, of the point made in Section 3, and which sits comfortably with what Zhang et al. and Ma and Qu describe as the complementary information the two modalities supply [2, 3].

A second equivalence relation can be defined not on appearance but on continuation, by declaring

$$x \sim_\Gamma x' \iff \Gamma(x, H, U) \cong \Gamma(x', H', U), \quad (37)$$

for a suitable notion of isomorphism, or sufficiently close resemblance under d_Γ , between continuation sets. This yields two partitions of the same underlying state space, $\mathcal{X}/\sim_\pi$ and $\mathcal{X}/\sim_\Gamma$, with no general reason for the two to coincide, and four logically distinct joint cases arise depending on whether a given pair of states is observationally close or far under $\sim_\pi$ and dynamically close or far under $\sim_\Gamma$:

	$\Gamma_1 \approx \Gamma_2$	$\Gamma_1 \not\approx \Gamma_2$
$Y_1 \approx Y_2$	observationally and dynamically similar	latent restriction
$Y_1 \not\approx Y_2$	different appearance, similar competence	full divergence

The upper-right cell, in which two systems appear observationally similar yet possess sharply different continuation sets, is the case Jokhai et al. document for aNE and pNE at their earliest, molecularly similar stage, and is the case this essay has been most concerned to make legible under the heading of latent restriction. The lower-left cell, in which two systems differ considerably in appearance yet share approximately the same continuation structure, remains a live and, on the present evidence, untested empirical possibility rather than a documented finding, but it follows from the same formal apparatus as a distinct and equally coherent prediction. Continuation type, on this reading, is properly understood as a quotient of $\mathcal{X}$ by future behavior, namely $\mathcal{X}/\sim_\Gamma$,

rather than as a quotient by present appearance, $\mathcal{X}/\sim_\pi$, and the entire argument of this essay can be restated as the claim that these two quotients need not agree.

12 From Failed Perturbation to Boundary Inference

Section 6 established, with appropriate care, that a single perturbation outcome constrains $\mathcal{A}_t(U)$ without thereby localizing $\partial\mathcal{A}_t(U)$, and Section 8 introduced the response surface $\chi_t : U \times Y \rightarrow \{0, 1\}$ and the empirically measurable region R_t as the object such a perturbation program can actually access. It is worth stating explicitly the inferential chain connecting these two observations, since each arrow in that chain introduces its own assumptions and is not to be taken for granted:

$$R_t \longrightarrow \widehat{\Gamma}_t(U) \longrightarrow \widehat{\mathcal{A}}_t(U) \longrightarrow \widehat{\partial\mathcal{A}_t(U)}, \quad (38)$$

where the hat denotes that each successive object is inferred, via some explicit statistical or mechanistic model connecting interventions and outcomes to the underlying state space, rather than measured directly. This licenses a principle that belongs naturally beside the witness-before-mechanism distinction of Section 7: a refusal witnesses a constraint on $\mathcal{A}_t(U)$; it does not, by itself, locate the cause of that constraint, nor does it, without further modeling, locate a specific point of $\partial\mathcal{A}_t(U)$. If neighboring interventions and neighboring target outcomes are sampled with sufficient density, and the response surface χ_t varies with an appropriate degree of regularity across that neighborhood, then the boundary between the successful and unsuccessful regions of R_t , itself an object in the observable product space $U \times Y$, can be estimated with increasing precision, and only through an explicit model relating that estimated decision boundary in $U \times Y$ to trajectories in $\mathcal{X}$ does the experimental finding become evidence about $\partial\mathcal{A}_t(U)$ itself. The resulting methodological symmetry is worth making explicit:

$$\text{Raman signal} \not\Rightarrow \text{chemical mechanism}, \quad (39)$$

$$\text{perturbation refusal} \not\Rightarrow \text{state-space boundary}. \quad (40)$$

Both are witnesses to something beyond themselves; neither is automatically identical to the thing it witnesses, and both require an additional, explicitly stated inferential step before the promotion from witness to underlying object can be justified.

13 Beyond Monotone Contraction

The essay has so far emphasized contraction, $\mathcal{A}_{t+1}(U) \subseteq \mathcal{A}_t(U)$, as though development were properly described as the monotone deletion of continuations from an initially larger set, and it is worth pausing to note explicitly where this emphasis may mislead. If $\Gamma_t(U)$ is equipped with a suitable measure μ , one may define a continuation volume $V_t(U) = \mu(\Gamma_t(U))$ and a corresponding loss

$$L_{t \rightarrow t+1}(U) = \log \frac{V_t(U)}{V_{t+1}(U)}, \quad (41)$$

though this expression should not be interpreted as an entropy in any thermodynamically or information-theoretically loaded sense without a specified probability structure over $\Gamma_t(U)$, which the biology as reported does not supply. A more cautious and empirically groundable quantity, for a finite family of experimentally sampled target outcomes $y \in Y$, is a continuation entropy defined directly from an empirical distribution $p_t(y | U)$ over reachable outcomes,

$$H_\Gamma(t) = - \sum_y p_t(y | U) \log p_t(y | U). \quad (42)$$

Developmental commitment need not reduce this quantity monotonically, and this is itself an interesting rather than a merely technical possibility: a cell might lose one broad family of futures, such as the capacity to generate an alternative brain region, while simultaneously gaining a finer-grained set of differentiated possibilities within the family it retains, such as the capacity to generate several distinct rhombomere-specific neuronal subtypes once committed to hindbrain identity. The “pages being ripped out” image used earlier in this essay is accordingly useful but incomplete: mathematically, the book being modeled may also acquire new pages conditional on having passed through certain earlier ones, so that a progenitor’s loss of broad lineage options is compatible with, and in typical developmental systems accompanied by, its gain of newly available specialized responses. The general theory should accordingly permit

$$\Gamma_{t+1}(U) = \mathcal{D}_t(\Gamma_t(U)), \quad (43)$$

where $\mathcal{D}_t$ is a deformation operator that may delete, split, merge, or newly create locally available continuations, rather than assuming the strictly monotone relation $\Gamma_{t+1}(U) \subseteq \Gamma_t(U)$ as a universal law; the aNE/pNE result establishes an instance of net contraction under the specific U tested, but the more general mathematical object this essay’s framework requires is deformation, of which contraction is only the simplest special case.

14 The Commutativity Test

A further consequence of taking developmental history seriously can be stated as a question about commutativity. Let D denote the operator of ordinary developmental evolution over some fixed time interval, and let u denote an intervention drawn from U ; one may then ask whether applying the intervention before development agrees with applying it after, that is, whether $D \circ u = u \circ D$. In biological systems generally it does not, and the discrepancy can be quantified directly, using a distance d_X on the state space, as

$$\Delta_{D,u}(x) = d_X(D(u(x)), u(D(x))), \quad (44)$$

a quantity that is large precisely when intervention and developmental history strongly fail to commute, and small when the order in which intervention and development are applied makes comparatively little difference to the outcome. This gives “history matters” a directly testable operational meaning distinct from the mutual-information criterion of Section 10: history matters, in this further sense, precisely when the effect of an intervention depends on when, relative to the system’s own developmental progression, that intervention is applied. The aNE/pNE result is again the relevant illustration: Jokhai et al. distinguish interventions capable of initially specifying which of the two populations a cell will become from later attempts, after that initial specification, to redirect an already-established population, and the sharp difference in outcome between these two orderings of intervention and development is exactly a large value of $\Delta_{D,u}$, while the authors’ own explicit openness about what sufficiently early or otherwise different interventions might accomplish is exactly the acknowledgment that $\Delta_{D,u}$ has been measured only at the specific developmental stages actually tested [1].

15 Implications for Longevity Research: Rejuvenation as Recovery of Admissible Futures

The framework developed above bears fairly directly on longevity research, though the connection should be stated carefully: RamanOmics itself supports only the multimodal-witnessing side of

the argument that follows, and the broader conclusions about intervention and reachability are theoretical extensions of the present framework rather than demonstrated longevity results.

The standard criterion in many rejuvenation experiments asks whether an intervention u moves an aged or senescent cell’s observed profile Y_{old} toward a youthful profile Y_{young} , typically operationalized as

$$d_Y(Y_{\text{treated}}, Y_{\text{young}}) < d_Y(Y_{\text{old}}, Y_{\text{young}}). \quad (45)$$

The two-admissibilities framework exposes a limitation in taking this criterion as sufficient for rejuvenation as such: a reduction in observational distance establishes movement in measured phenotype space, but it does not by itself establish restoration of the continuation structure characteristic of a younger cell, since

$$Y_{\text{treated}} \approx Y_{\text{young}} \not\Rightarrow \Gamma_{\text{treated}}(U) \approx \Gamma_{\text{young}}(U). \quad (46)$$

RamanOmics gives concrete reason to take this gap seriously, since senescence is not captured by any universal single marker, and Raman and transcriptomic measurement disclose complementary rather than redundant aspects of cellular condition [2, 3]; a purely transcriptomic rejuvenation signature is accordingly only one projection of the treated state, and a treatment might normalize that projection while leaving biochemical organization, chromatin-mediated competence, or the capacity to respond appropriately to subsequent stress and repair signals substantially unrestored. The word “rejuvenation” therefore compresses at least three logically distinct and only loosely correlated claims: marker reversal, multimodal state restoration, and functional continuation restoration.

Under this framework, senescence itself is more naturally characterized not merely as occupation of a recognizable region of molecular state space but as a deformation of continuation structure, $\Gamma_{\text{young}}(U) \rightarrow \Gamma_{\text{senescent}}(U)$, of the general kind $\mathcal{D}_t$ introduced in Section 13 rather than a simple monotone contraction, since some responses may disappear entirely, others may become abnormally accessible, and still others may merely shift in threshold or timing. Genuine rejuvenation, on this reading, involves some approximate inverse of that deformation,

$$\Gamma_{\text{senescent}}(U) \xrightarrow{u} \Gamma_{\text{rejuvenated}}(U) \approx \Gamma_{\text{young}}(U), \quad (47)$$

rather than merely the movement of a measured expression profile toward a youthful centroid. This suggests a considerably stronger experimental design than a simple before-and-after comparison: young, aged, senescent, and treated cells could each be subjected to a common standardized panel of interventions U , comprising stresses, repair signals, metabolic challenges, and differentiation cues, and the resulting empirical response relations $R_g = \{(u, y) : \chi_g(u, y) = 1\}$, one for each group g , compared directly via a distance $d_R(R_{\text{treated}}, R_{\text{young}})$ alongside conventional molecular-age distances. The experiment thereby asks not whether the treated cell resembles a young cell in appearance, but whether it responds to a common battery of possible futures the way a young cell does, a considerably stronger and more directly functional test of rejuvenation.

The framework further predicts that the biomarker most strongly correlated with chronological age need not be the biomarker most informative about retained biological competence, since an age estimator $B_{\text{age}} : Y \rightarrow \mathbb{R}$ and a continuation estimator $B_{\Gamma} : Y \rightarrow \hat{\Gamma}$ answer different questions, the former locating a system along an age-associated observational axis and the latter estimating what the system remains capable of doing; an ideal longevity assay would report both rather than treating the first as a proxy for the second. The same caution urged elsewhere in this essay against promoting a witness into a mechanism applies here without modification, since Ma and Qu explicitly note that Raman spectral changes do not by themselves identify particular molecular

species, nor establish that the detected biochemical alterations causally contribute to senescence rather than merely accompanying it [3], so that even a highly predictive multimodal senescence signature remains a witness rather than an established mechanism or a validated therapeutic target:

$$\text{age-correlated signature} \not\Rightarrow \text{senescence mechanism} \not\Rightarrow \text{rejuvenation target}. \quad (48)$$

A final caution belongs here against overcorrecting in the opposite direction. Complete erasure of developmental and physiological history would not itself be the appropriate objective of a longevity intervention, since biological history encodes useful specialization as well as accumulated damage, and a mature cell’s comparatively restricted continuation structure is partly constitutive of its functional identity rather than simply a deficit to be reversed; longevity, accordingly, cannot be equated with indiscriminately maximizing the size of $\Gamma(U)$, since an unrestricted cell is not thereby a healthy one. The more appropriate target is a continuation structure appropriate to a given tissue and age, $\Gamma_{\text{healthy, age}}$, rather than maximal developmental plasticity as such. Aging, under this framework, is best conceptualized not as movement away from a single youthful point but as a sequence of transformations of possibility, $\Gamma_0 \rightarrow \Gamma_1 \rightarrow \dots$, and rejuvenation correspondingly not as movement back toward that point but as the restoration of selectively chosen lost or distorted continuations, restoring desirable competence without indiscriminately reopening developmental possibilities that mature tissue has good functional reason to keep closed.

Conclusion

Neither of the two papers under discussion proposes the formalism developed across the preceding sections, and neither should be read here as offering confirmation of a theory it was not designed, or intended by its authors, to test. Ma and Qu, reporting on RamanOmics, demonstrate why a single observational projection, however sophisticated, can inadequately characterize cellular state as such [3]. Jokhai et al. demonstrate why present appearance, together with whatever environmental instruction an experimenter happens to choose to apply, can inadequately characterize developmental potential as such [1]. Neither result depends, in order to be stated correctly, on the vocabulary of reachable sets, topological boundaries, or continuation types developed in this essay; both results, however, become considerably easier to state with precision once that vocabulary is made available, and the resulting integration of the two papers’ contributions may be organized around the single diagram introduced in Section 4,

$$H_t \longrightarrow X_t \xrightarrow{\pi} Y_t, \quad (X_t, U) \longrightarrow \Gamma_t(U), \quad (49)$$

in which RamanOmics improves the rightward map from X_t to Y_t by supplying an additional, only partially redundant witness, thereby contracting the epistemic object $\mathcal{H}$ introduced in Section 4, while Jokhai et al. interrogate the downward map from (X_t, U) to $\Gamma_t(U)$ by experimentally challenging developmental competence, thereby contracting the ontic object $\mathcal{A}_t(U)$ introduced in Section 1. The former concerns what can be inferred about what presently exists; the latter concerns what the existing system can subsequently do, and the two concerns, though routinely elided under the single word “state,” answer to entirely different mathematical objects and entirely different forms of evidence.

Together, the two papers motivate a broader distinction among the four objects that a single, unadorned state vector conflates by construction, namely observed position, accumulated history, the topological boundary of what remains reachable from that position and history under a specified intervention family, and the corresponding space of admissible continuations. The resulting

formulation, stated as compactly as the preceding sections permit without recourse to any further notational apparatus, is that cellular identity is not position alone, but position together with constraint, history, and possible continuation, jointly. What a cell currently measures as, under any given assay or combination of assays, is only ever a partial report on what that cell could still, under some as yet untested intervention, become; and what a cell could still become is itself, in turn, a fact possessing its own history, rather than a static property that could in principle be read off from a single instantaneous snapshot, however comprehensive that snapshot might otherwise be. A better witness, in the sense of Section 3, does not thereby enlarge the future available to the system being witnessed, and a smaller future, in the sense of Sections 1 and 2, does not thereby make the system whose future has narrowed any better witnessed; these two operations, refining $\mathcal{H}$ and contracting $\mathcal{A}_t(U)$, remain, on the account developed here, mathematically and empirically distinct all the way down.

References

- [1] Jokhai, R. T., Dundes, C. E., Ahsan, H. S., et al. Two parallel neural ectoderm progenitors contribute to the developing brain. *Nature Neuroscience* (2026). <https://doi.org/10.1038/s41593-026-02433-7>
- [2] Zhang, K., Chen, X., Monticolo, F., et al. RamanOmics decodes the spatial vibrational–molecular architecture of senescence in aging and repair. *Nature Aging* (2026). <https://doi.org/10.1038/s43587-026-01219-7>
- [3] Ma, S. & Qu, J. The chemical fingerprint of cellular senescence. *Nature Aging* (2026). <https://doi.org/10.1038/s43587-026-01230-y>
- [4] Waddington, C. H. *The Strategy of the Genes: A Discussion of Some Aspects of Theoretical Biology*. George Allen & Unwin, London (1957).
- [5] Moris, N., Pina, C. & Martinez Arias, A. Transition states and cell fate decisions in epigenetic landscapes. *Nature Reviews Genetics* **17**, 693–703 (2016). <https://doi.org/10.1038/nrg.2016.98>
- [6] MacArthur, B. D. & Lemischka, I. R. Statistical mechanics of pluripotency. *Cell* **154**, 484–489 (2013). <https://doi.org/10.1016/j.cell.2013.07.024>
- [7] Trapnell, C., Cacchiarelli, D., Grimsby, J., et al. The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells. *Nature Biotechnology* **32**, 381–386 (2014). <https://doi.org/10.1038/nbt.2859>
- [8] La Manno, G., Soldatov, R., Zeisel, A., et al. RNA velocity of single cells. *Nature* **560**, 494–498 (2018). <https://doi.org/10.1038/s41586-018-0414-6>
- [9] Setty, M., Kiseliovas, V., Levine, J., Gayoso, A., Mazutis, L. & Pe’er, D. Characterization of cell fate probabilities in single-cell data with Palantir. *Nature Biotechnology* **37**, 451–460 (2019). <https://doi.org/10.1038/s41587-019-0068-4>
- [10] Rodriguez-Fraticelli, A. E., Weinreb, C., Wang, S.-W., et al. Single-cell lineage tracing unveils a role for TCF15 in haematopoiesis. *Nature* **583**, 585–589 (2020). <https://doi.org/10.1038/s41586-020-2503-6>
- [11] Chang, H. H., Hemberg, M., Barahona, M., Ingber, D. E. & Huang, S. Transcriptome-wide noise controls lineage choice in mammalian progenitor cells. *Nature* **453**, 544–547 (2008).
- [12] Garcia-Ojalvo, J. & Martinez Arias, A. Towards a statistical mechanics of cell fate decisions. *Current Opinion in Genetics & Development* **22**, 619–626 (2012).
- [13] Kutejova, E., Sasai, N., Shah, A., Gouti, M. & Briscoe, J. Neural progenitors adopt specific identities by directly repressing all alternative progenitor transcriptional programs. *Developmental Cell* **36**, 639–653 (2016).
- [14] Levine, M. & Davidson, E. H. Gene regulatory networks for development. *Proceedings of the National Academy of Sciences USA* **102**, 4936–4942 (2005).
- [15] Davidson, E. H. *The Regulatory Genome: Gene Regulatory Networks in Development and Evolution*. Academic Press (2010).