

Inheritance Entropy: A Model-Independent Method to Probe the Hereditary Structure of Cell Lineage Trees

Alessandro Allegrezza ¹, Riccardo Beschi ¹, Domenico Caudo ^{1,2}, Andrea Cavagna ^{3,1,4,*}, Alessandro Corsi ⁵, Antonio Culla ³, Samantha Donsante ^{5,6}, Giuseppe Giannicola ⁷, Irene Giardina ^{1,3,4}, Giorgio Gosti ^{2,8}, Tomás S. Grigera ^{9,10,11,3}, Stefania Melillo ^{3,1}, Biagio Palmisano ^{5,3}, Leonardo Parisi ^{3,1}, Lorena Postiglione ^{3,12}, Mara Riminucci ⁵ and Francesco Saverio Rotondi ¹

¹Dipartimento di Fisica, *Sapienza Università di Roma*, Rome, Italy

²Center for Life Nano & Neuro Science, *Italian Institute of Technology*, Rome, Italy

³Istituto Sistemi Complessi, *Consiglio Nazionale delle Ricerche, UOS Sapienza*, Rome, Italy

⁴Istituto Nazionale di Fisica Nucleare, *Sezione Roma 1*, Rome, Italy

⁵Dipartimento di Medicina Molecolare, *Sapienza Università di Roma*, Rome, Italy

⁶Tettamanti Center, *Fondazione IRCCS San Gerardo dei Tintori*, Monza, Italy

⁷Dipartimento di Scienze Anatomiche, Istologiche, Medico Legali e dell'Apparato Locomotore, *Sapienza Università di Roma*, Rome, Italy

⁸Istituto di Scienze del Patrimonio Culturale, *Consiglio Nazionale delle Ricerche, Montelibretti*, Italy

⁹Instituto de Física de Líquidos y Sistemas Biológicos, *CONICET and Universidad Nacional de La Plata*, La Plata, Argentina

¹⁰CCT CONICET La Plata, *Consejo Nacional de Investigaciones Científicas y Técnicas*, La Plata, Argentina

¹¹Departamento de Física, *Facultad de Ciencias Exactas, Universidad Nacional de La Plata*, La Plata, Argentina

¹²Dipartimento di Ingegneria Chimica, dei Materiali e della Produzione Industriale, *Università degli Studi di Napoli Federico II*, Napoli, Italy



(Received 23 October 2025; accepted 24 April 2026; published 2 June 2026)

Quantifying the impact of hereditary transmission within lineage trees remains a fundamental challenge universal to a wide array of biological domains. Here we introduce the new concept of inheritance entropy, a quantity designed to gauge the hereditary structure of inactive cells across a lineage. We measure this entropy in 32 human stem cell clonal colonies, obtained from high-definition single-cell lineage tracing experiments, and show that in the greatest majority of clones the entropy is decisively smaller than that of the corresponding nonhereditary ensemble, hence proving that variations in the proliferative power of stem cell lineages are determined by hereditary epigenetic factors that regulate cell-cycle exit. The method can also be employed to locate the specific node of the tree where a mutation in the probability of inactivity occurs, together with a determination of the lag between the mutation ultimately leading to inactivity and its actual expression. This framework can be used to assess in a robust, simple, and model-agnostic way the hereditary origin of differential growth in any type of lineage trees.

DOI: [10.1103/p2mj-q682](https://doi.org/10.1103/p2mj-q682)

I. INTRODUCTION

The role of inheritance in cell replication has been a subject of intense study for a long time [1–7]. It is believed that cells switch between molecularly and phenotypically distinct states that are passed on to the descendants, thus regulating their fate [5,8]. One way to assess inheritance is therefore to characterize cell states and track how they evolve along a lineage. Single-cell profiling allows to measure the whole transcriptome, as well as proteomes and metabolic signatures, hence characterizing the multidimensional space of cell states (the “state manifold” or “transcriptional landscape” [9]); but this type of analysis is hard to combine with classic lineage

tracing techniques, which are, however, important to precisely track the hereditary ramifications. A powerful alternative is provided by barcoding lineage tracing, which allows us to acquire simultaneously cell state information and partial lineage relationships from bulk measurements [10–13]. However, despite their great potential, barcoding techniques still present limitations in terms of applicability to human cells [14] and accuracy in lineage reconstruction [9].

Epigenetic inheritance, namely the transmission of phenotypic traits across generations without changes to DNA sequence, is particularly elusive and difficult to assess: Non-genetic transmission can occur through a great variety of molecular mechanisms, among which DNA methylation, histone modifications, and noncoding RNA regulation [15,16], which makes it hard to pin down its hereditary nature in a sharp way. Epigenetic inheritance, though, is also hugely relevant, unfolding its full potential in the realm of stem cells, whose high replicative capacity and promising regenerative potential make them the ideal arena to study how different phenotypic traits emerge and evolve within a clonal population. Besides, differentiation and cell-fate—which are

*Contact author: andrea.cavagna@roma1.infn.it

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

core issues of stem cell biology—strongly rely on epigenetic mechanisms [17,18]. Exploring epigenetic inheritance in cell lineages is therefore a further urgent reason to try and find robust methods of hereditary assessment.

A particularly general trait, defining the hereditary structure of lineage trees at a very fundamental level, is *inactivity*, namely the state the cell enters when it stops dividing. This is primarily for three reasons. First, because in the context of cellular aging, inactivity essentially means senescence, whose epigenetic origin has huge clinical impact [19–22]. Second, because inactivity is in general a more complex state than just senescence: An inactive cell can be reversibly so (quiescent) or irreversibly inactive (differentiated or senescent), a difference that impinges greatly on the fate of cells [23–26]. Finally, inactivity is a very general template for any mechanism differentially cutting branch proliferation within a lineage tree, irrespective of the specific nature of the tree. Hence, studying the hereditary ramifications of inactivity has a potentially broad impact on the more general scope of gauging inheritance in *any* kind of tree.

Inactivity, however, has some peculiarities that make its hereditary character—if any—particularly hard to pin down. Normally, inheritance implies that when a mutation emerges in a cell, thus giving rise to a new phenotype, this mutation is inherited by its progeny, thus amplifying the detectability of that phenotype with increasing generations: When enough doublings have been reached, there is sufficient bulk information in the culture to map the hereditary ramifications of that mutation through bulk methods. The tricky thing about inactivity, though, is that its very impact on the tree development—namely the obliteration of entire branches—is also what makes it hard to characterize it at an hereditary level: Once inactivity emerges in one cell, this phenotype is *not* amplified in its progeny, as the effect of inactivity is rather to delete the descendants of that one cell, thus erasing all the information they potentially carried. Bluntly put, it is impossible to extract information from cells that do not exist.

Since all downstream branches of the inactive cell are missing, it seems instead that one might try to extract information from its *upstream* branches—or progenitors—which have been alive at some point in the tree development. This strategy, though, can be pursued only if inactivity manifests itself a few generations *after* the mutation causing it; otherwise, neither the future nor the past generations would contain any information about the emergence of inactivity. At least in the case of senescence, there is growing evidence that this is indeed the case: Senescence is characterized by markers that accumulate in a gradual manner *prior* to its expression [19–22]; hence, it is possible that this happens more generally for the expression of inactivity. And yet, the very absence—by definition—of the inactivity trait in the progenitors of the inactive cell, makes it unclear how to phenotypically retrace upstream the origin of the mutation and its hereditary unfolding along the lineage.

Finding a new method to back-track the ramifications of inactivity, hence establishing beyond doubt its hereditary character, is what we do here. The central idea behind this result is to extract information about the hereditary distribution of inactive cells through the calculation of the Shannon entropy associated to the topology of the entire lineage tree. We test this method through experiments of

high-resolution single-cell lineage-tracing of bone marrow stromal cells clonal colonies and find that indeed in the greatest majority of lineages the entropy is decisively smaller than that of the corresponding nonhereditary ensemble. Moreover, through our method, the epigenetic change giving rise to inactivity can be traced upstream, thus revealing a map of the mutation events across the colony development. This will allow us to measure the mutation lag, namely the generation gap between the mutation leading to inactivity and the actual expression of its phenotype.

II. LINEAGE TOPOLOGY

A. The system

Human bone marrow stromal cells (BMSCs) are a prominent source of skeletal stem cells with ground-breaking therapeutic potential [27]. However, BMSC colonies have very heterogeneous *in vivo* behavior, due to their different potency; this unpredictability is the greatest hurdle to the development of bone regeneration therapies [28].

Colony-level heterogeneity urges a fundamental question: How is it possible that one colony as a collective unit behaves differently from another one? If cell-to-cell variability were just an uncorrelated random process, then a million cells in a transplant-bound colony would be enough to yield statistical homogeneity, hence washing out any colony-level traits. A possible answer is that the differences between two originating cells are transmitted to their progenies and collectively persist through an epigenetic hereditary mechanism.

If hereditary epigenetic factors play a major role in determining cell-cycle exit, then they may be involved also in the regulation of BMSCs potency, thus shedding some light on that elusive holy grail of stem cell research that is differentiation potential [29,30]. Controlling the hereditary character of inactivity could therefore help developing protocols to control intercolony heterogeneity, thus advancing the clinical potential of skeletal stem cells. For these reasons, BMSCs are a very important testbed system for the development of new methods to detect inheritance in lineage trees.

B. Experimental data

Human BMSCs were seeded at very low density in order to produce clonal colonies derived only from single, spatially isolated cells; our dataset consists of 32 such single-cell-derived clones. The isolated nature of each colony is very important: If different colonies merge, then their clonal nature is destroyed, severely confounding the data analysis [31,32]. Growth was studied with phase contrast high-resolution time-lapse microscopy. Instead of adopting a fixed duration of the experiments, we fix the maximum generation of the colonies, keeping each experiment going until all replicating cells have arrived to the seventh generation, $k_{\max} = 7$; in this way we do not need to correct for any type of statistical bias [33,34]. Details of the experiment and a full characterization of the biological parameters of all colonies, can be found in Ref. [35].

Samples of four lineage trees are reported in Fig. 1, while all trees in the dataset can be found in Sec. 9 in Supplemental Material (SM) [36]. If no branches are interrupted, then the




FIG. 1. Lineage trees of experimental BMSC colonies. Red segments represent active cells and black disks represent mitosis; the length of each segment is fixed, not related to the division time. Inactive cells (which are all G_0 in these samples) are represented as black segments ending in a cap. The first cell of the colony is not represented, as we do not record its birth, but only its division (the central disk of the tree). Ovals separate different generations. Cells at the last generation (the k_7 leaves) are tracked up to their birth, not up to their division, and hence they are neither labeled as active nor inactive and we represent them in gray.

number of cells at the seventh generation is $2^7 = 128$, and a handful of colonies do indeed reach this state (3 of 32). But often branches *are* interrupted, which may happen for two reasons: Either the cell commits apoptosis and dies, or the cell stops dividing, yet remaining alive (black capped segments in Fig. 1); in this last case the cells enters the G_0 phase. We classify a cell as G_0 if it does not divide for 84 h (3.5 days), at which point we stop tracking it; this threshold is to be compared to the mean division time of BMSCs, which is 20 ± 4 h (the robustness of this G_0 criterion is thoroughly tested in Ref. [35]). Only 17 cells die in our whole dataset against 310 G_0 cells; given such paucity of apoptosis and given that the effect of a dead cell on the lineage topology is exactly the same as a G_0 cell, we will call *inactive* the whole set of G_0 and dead cells.¹ Conversely, we will call *active* those cells for which we do observe a mitosis. The final “leaves” of the tree, i.e., cells that belong to the last generation ($k_{\max} = 7$), are only tracked up to their birth, not up to their division, and hence they are neither labeled as active nor inactive; we call these cells the k_7 leaves of the lineage tree. As we shall see later on, two useful characterization of a tree are the total number of inactive cells, N_{inact} , and the total number of cells missing at the last generation, N_{miss} , that is the number of missing gray leaves in Fig. 1.

¹This is a rather standard (and venerable) convention in the literature on cellular senescence: Any difference between dead cells and cells that have ceased to divide is immaterial when dealing with population growth, see for example Refs. [37–39].

C. The peculiar distribution of inactive cells

Inactive cells determine the topology of a lineage through both their number and their position within the tree. The position of an inactive cell can be described by two “coordinates”: its generation k and its radial position. The first impacts greatly on the number of missing k_7 leaves, N_{miss} : The earlier the generation at which a cell becomes inactive, the larger the number of leaves that are cut; for example, the effect of an inactive cell is very different if it is located at generation $k = 2$ ($N_{\text{miss}} = 32$) or at $k = 6$ ($N_{\text{miss}} = 2$). At first sight, this effect seems related to the Luria-Delbrück argument about the uneven impact of hereditary mutations [40]: If the probability of a mutation is homogeneous throughout the lineage, then the fluctuations in the occurrence at different generations causes very large heterogeneities in the final number of cells carrying the mutation, which is the conceptual basis for all Luria-Delbrück-inspired fluctuation assays [41]. However, because in this case the phenotype is “inactivity” and because the very definition of inactive cell is that it does not have any progeny, the fact that the inactive phenotype is inherited by the (nonexistent) descendants of the inactive cell is completely trivial and uninformative. Therefore, to check whether the mechanism that ultimately leads to inactivity has an hereditary nature, we cannot use a standard fluctuation assay. Instead, we have to employ the lineage information to check whether some kind of mutation occurred *before* the emergence of the inactive cell, namely upstream in the tree.

To make progress, we notice that trees with many inactive cells display something more than the mere amplification of fluctuations due to the different intergeneration placement,

namely the fact that the radial distribution of inactive cells across different branches has quite a nontrivial pattern. Consider, for example, lineage 20241017_11 in Fig. 1: There is clearly something peculiar about the locations of the inactive cells in this tree, as they seem more frequent within the same few branches on the southwest side in a definitely nonrandom way; a similar observation could be made for lineage 20241203_03, where inactive cells are mostly concentrated in the west wing. This is a very general trait in our dataset: In most lineages we observe significant nonrandom heterogeneities in the number of inactive cells between branches, indicating that the emergence of these cells is more likely in some branches than in others. If we follow a branch with many inactive cells upstream along the lineage, then an hereditary mechanism would suggest that at some point we must find a mitosis at which some factor linked to the probability of inactivity changed, making one of the two daughters more prone to producing inactive cells in its progeny than the other. We need to put this hypothesis to a quantitative test; to do that we will use the concept of *entropy*.

III. QUANTIFYING INHERITANCE

A. Inactivity imbalance

As we have noted, the most conspicuous effect of the presence of an inactive cell is that all cells that would have been born as its progeny are instead missing, so that the number of cells at the final generation, $k_{\max} = 7$, is smaller than what it could have been. Hence, we will use the number of missing k_7 leaves as an indicator of the impact of inactive cells within that branch. We proceed as follows.

Each mitosis in the lineage generates two branches, left and right; in a mitosis connecting generation $(k - 1)$ to generation k , each branch can potentially produce up to $2^{(7-k)}$ descendants at generation 7. However, if inactive cells emerge downstream along these branches, then the *actual* number of k_7 descendants will be smaller (some of these leaves will be missing). Given a mitosis m , we will indicate with M_m^{left} the number of missing k_7 leaves in its left branch and similarly for M_m^{right} (see Fig. 2). We stress that the count of missing cells is done only among the cells of the last recorded generation (in our experiments, the k_7 leaves).

We then define the *inactivity imbalance* of mitosis m as the difference between the number of missing k_7 cells within its left and right branches (Fig. 2),

$$I_m = |M_m^{\text{left}} - M_m^{\text{right}}|, \quad (1)$$

where we use the absolute value because the actual left-right labeling of the branches after a division is arbitrary.

The definition of imbalance can be generalized to any value of $k_{\max}$. Moreover, the imbalance can be defined also for trees where all cells become inactive before $k_{\max}$: In this case, one can define an effective $k_{\max}^{\text{eff}}$, as the generation of the latest inactive cell; then the definition proceeds as before, with $k_{\max} \rightarrow k_{\max}^{\text{eff}}$. Such incomplete trees are not present in our experimental dataset, but they may be generated by the null models that we will consider in the next section.

The interesting thing about the inactivity imbalance is that it provides a sort of “mutation” map, giving a specific




FIG. 2. Inactivity imbalance. We illustrate here how the inactivity imbalance is calculated, using a portion of a real lineage (20230503_13). White dotted segments represent k_7 leaves that are missing because of the presence of some inactive cells upstream in that branch. For each mitosis or node, the inactivity imbalance is defined as the modulus of the difference between the number of missing k_7 leaves in the right branch and the number of missing k_7 leaves in the left branch.


information about what happens at each mitosis along the tree (see Fig. 3, left). A large imbalance on a certain node means that the two branches generated by that mitosis are *ultimately* very different from each other, suggesting that a mutation in the probability of generating inactive cells occurred there. From a topological point of view, the mutation is most easily attributed to the node, i.e., to the mitosis, whose large imbalance marks the difference between the two sub-branches born out of that mitosis; from a biological point of view, though, it is more appropriate to attribute the mutation to the one of the two sister cells born out of that mitosis, whose progeny is riddled with inactive descendants (this is also the most natural interpretation in the context of hereditary models). The main point is that the inactivity imbalance provides information about a possible change in the propensity to inactivity *before*—namely upstream in the tree—the inactivity phenotype is actually expressed. The imbalance is therefore a promising observable in the search of an hereditary test of cell-cycle exit.

B. Entropy

At variance with other measures of tree imbalance, such as the Colless index [42,43], we do not consolidate the inactivity imbalance (1) into a sum over all nodes of the tree. This is a crucial point: When trying to characterize inheritance, a large inactivity imbalance at a few nodes is *not* the same as a buildup of small inactivity imbalances scattered over many nodes.

Instead, inheritance is signalled by an exceptional event, namely a large imbalance concentrated on one or a few nodes, revealing a difference between the two cells generated at that mitosis, a difference that is passed on to their progenies (Fig. 3, left). Therefore, we need a quantity able to reveal when there are atypically large values of the inactivity imbalance in the lineage, and when—on the other hand—small imbalances are scattered all over the tree. The most promising tool to do this is the Shannon entropy. From the imbalance,

(a) Inactivity Imbalances - Biological Lineage



(b) Inactivity Imbalances - Scrambled Lineage




FIG. 3. Inactivity imbalance: BMSC vs scrambled. The distribution of the inactivity imbalance across the different mitosis is very different in biological vs randomly scrambled trees. (a) For lineage 20230503_13, the inactivity imbalance, I_m , is graphically represented as a purple vertical brushstroke on each mitosis or node (zero imbalances are not marked). In this lineage most of the imbalance is concentrated on just two nodes; of these two mitosis, the one connecting generation 1 to generation 2 (imbalance 22) is most likely that at which there has been a mutation in the probability of emergence of inactive cells. This large heterogeneity in the distribution of the imbalances across the lineage is the clearest symptom of inheritance. The entropy S measures exactly this heterogeneity; in this lineage $S_{\text{biological}} = 2.17$, which is rather low compared to the entropy of the tree in panel (b), which is one instance of the 10^6 randomly reshuffled versions of lineage 20230503_13 [random reshuffling (RR) ensemble; see text]; the number of inactive cells and their generation is the same as in the biological lineage, but their radial positions across the tree have been randomly scrambled, hence severing all potential hereditary relationships. In this tree the imbalance is more homogeneously distributed across the nodes than in the biological tree, thus giving a much lower inheritance signal: The entropy is indeed significantly larger than the biological one, $S_{\text{scrambled}} = 3.21$. If we repeat the scrambling 10^6 times, then we find $S_{\text{scrambled}} \leq S_{\text{biological}}$ in only 7 cases, proving that the evidence of inheritance is very strong in this clone.

we first derive the normalized weight of mitosis m ,

$$w_m = \frac{I_m}{\sum_n I_n}, \quad (2)$$

where the sum is extended over all mitosis in the lineage. Then, using the weights w_m , we define the lineage entropy as

$$S = - \sum_m w_m \ln w_m. \quad (3)$$

The entropy S is high when the inactivity imbalances I_m all have approximately the same value, forming a random pattern across the tree, so that we receive no particular information about the occurrence of an inherited change at any mitosis in the lineage; on the other hand, entropy is low when there are few large imbalances concentrated on a handful of mitosis, as in that case we have reliable information that an inherited change occurred at those mitosis (Fig. 3, left). In short—and very coarsely—high entropy means that inheritance is unlikely, while low entropy means positive evidence of inheritance. As we shall see in the next section, though, “high” and “low” entropy will only be significant in a comparative way.

We stress that the entropy in (3) is *not* the entropy associated to the probability distribution of the inactivity imbalance, because the weight w_m is not the probability of a certain value of I_m but is proportional to the imbalance itself; indeed, the sums do not run over all possible values of the imbalance but over all possible nodes (i.e., mitosis) over which the imbalance is defined. For the same reason, the normalized weights w_m do not have an obvious frequentist interpretation; therefore, we should not blindly rely on the standard properties of S in the context of Shannon information theory. It is also worth noting that similar concepts have been employed in different contexts: In astronomy an analogous quantity has been

defined from the intensity field of galaxy radioastronomical data for image reconstruction [44]; in economics, the entropy in (3) is related to Theil’s index of economic inequality [45]. In all these contexts, the key idea is to pinpoint the anomalous values of some observable (imbalance, intensity, inequality) across a certain “space” (lineage, image, country) rather than across a certain probability distribution. This is what the entropy defined in (3) does.

C. Inheritance test

The entropy of a lineage is designed to depend on the specific arrangement of the inactive cells, i.e., on their hereditary correlations. But S also depends on the total number of inactive and missing cells, irrespective of their hereditary distribution, and on the overall size of the tree. Therefore, it is hard to establish the absolute degree of inheritance of a tree by simply measuring its entropy S . Instead, the correct way to proceed is to compare the entropy of a specific biological lineage to that of an ensemble of *nonhereditary* trees.

We proceed as follows: The inheritance entropy is measured for a large number $\mathcal{R}$ of trees belonging to the nonhereditary (or null) ensemble, thus producing a nonhereditary probability distribution of S ; the hereditary significance of the biological tree, namely its P value, is given then by the fraction of nonhereditary trees that happen to have $S \leq S_{\text{biological}}$, namely by the probability that the null ensemble produces a tree with evidence of inheritance stronger than or equal to the biological one. The lowest P value (highest significance) is $P < 1/\mathcal{R}$, which happens when none of the $\mathcal{R}$ nonhereditary trees has entropy smaller than the biological one.

Such inheritance test requires the definition of an appropriate null ensemble of nonhereditary lineages. One way to

proceed is to generate this ensemble by means of a non-hereditary null-model; given a biological lineage we want to test, the null model must be fitted to that lineage, which is generally feasible by using the maximum likelihood method, considering that the null model will not be overly complicated (otherwise it would hardly be a *null* model). This apparently straightforward procedure, though, is less trivial than it might seem: As we shall see in the next section, the choice of the null model will not be unequivocal. Moreover, we will discover that the null ensemble that is most effective to assess inheritance, i.e., the one with the highest sensitivity, is not generated by any specific null model.

IV. NONHEREDITARY ENSEMBLES AND TEST VALIDATION

In this section we will discuss three different types of null ensembles that can be used to implement the inheritance test. Two of them are generated by their relative null models, while the third null ensemble requires no generating model, as it simply derives from the data. In the end, after a careful benchmarking of the three methods, only one inheritance test will be selected. The reader not interested in this benchmarking tour de force can skip this section.

A. The Galton-Watson ensemble

The simplest nonhereditary model of inactivity is one in which each cell is inactive with probability q . This is equivalent to a Galton-Watson branching process (GW) [46], with probabilities to have either zero or two children given by $p_0 = q$ and $p_2 = 1 - q$. Fitting GW to a biological tree is simple: The parameter q which maximizes the likelihood is the empirical value of q , i.e., the total number of inactive cells divided by the total number of cells of that tree (see Sec. 1 in the SM [36]).

Note that some trees of the null ensemble may not arrive at $k_{\max} = 7$, because all cells become inactive before; this is why in Sec. III we have defined the imbalance also for incomplete trees. Moreover, a very small fraction of trees of the ensemble have imbalance zero at each site, in which case the entropy is not defined; we discard these trees from the ensemble.

Even within a nonhereditary tree the imbalance has a mean trend with the generations: Early generations expect a higher imbalance than late ones, just because the randomness in the longer subtrees adds up. When using a null model it is important to normalize the imbalance of the tree to be tested and of the trees in the null ensemble by the expected imbalance of the null model, which can be calculated exactly for GW (see Sec. 1 in the SM [36]). Once the imbalance has been thus normalized, the calculation of the entropy proceeds as in (2) and (3). Finally, the P value of a certain lineage, P_{GW} , is given by the fraction of GW trees with entropy smaller than or equal to the entropy of that lineage.

Despite the clarity and elegance of the GW model, it is possible to foresee a problem: This simple version of the GW ensemble *does not* encompass all possible nonhereditary trees; the probability to be inactive could very well be dependent on the generation, $q = q_k$, producing trees that are still nonhereditary but unlikely to be produced by a GW model

with constant q , thus exposing the method to the risk of creating false positives. The response to that is to use a more complete null ensemble, which can be done by introducing a GW model where the probability of inactivity depends on generation.

B. The nonhomogeneous Galton-Watson ensemble

The nonhomogeneous Galton-Watson model (NHGW) is an uncorrelated branching process where the probability to be inactive depends on the generation k , so that the model is defined by a vector q_k [47,48]. There are sound biological reasons for using NHGW instead of GW, irrespective of any consideration of inheritance: There is plenty of evidence [49–51] showing that the fraction of senescent or dead cells can change significantly with the age of the cell culture, and therefore with the generations. Hence, it is unwise to commit to a null model that excludes this possibility.

As in the case of GW, also for NHGW the maximum likelihood method prescribes that the correct vector q_k to fit a specific tree to the model is simply given by the empirical set of inactivity probabilities, i.e., the total number of inactive cells divided by the total number of cells, both calculated at fixed generation k for that tree (see Sec. 2 in the SM [36]). Moreover, even though the calculation of the expected NHGW imbalance is significantly more complicated than in GW, it can still be done exactly (see Sec. 2 in the SM [36]), so we can normalize the imbalance and calculate the entropies, finally producing a P value of this ensemble, P_{NHGW} .

C. The randomly reshuffled ensemble

The third nonhereditary ensemble that we consider does not rely on a null model, but it operates directly on the lineage tree to be probed. The idea is to produce an ensemble of trees that have as much as possible the same structure in terms of inactive cells and missing cells as the original tree, but where all potential inheritance relations between mother and daughter have been severed. This ensemble is easily achieved by perform a random reshuffling (RR) of the input tree: At each generation k all cells are scrambled *within that same generation*; when two cells are exchanged, so are their progenies. Then we move to generation $k + 1$ and we scramble again, and so on (see Sec. 3 in the SM for a detailed description of the reshuffling procedure [36]). Notice that in this way all potential hereditary connections are randomized and yet the number of inactive cells at each generation remains the same (see Fig. 3, right). We shall see that this structural stability gives to the RR ensemble a decisive advantage over the other two ensembles. The P value of the RR method, P_{RR} , is finally given by the fraction of scrambled trees that happen to have entropy lower than or equal to the one of the lineage we are probing.

D. Benchmarking the methods

To evaluate the performance of the three entropy-based inheritance tests we must compare their output to some ground truth, i.e., to run the tests on a large set of lineage trees of known inheritance properties. This can be done by using a testbed of trees generated by a nonhereditary model (in which




FIG. 4. Benchmarking the inheritance tests. We plot the fraction Φ of trees that pass the inheritance test ($P < 0.05$) as a function of the size $k_{\max}$, for lineages of different hereditary nature. (a) The three methods are benchmarked on a set of homogeneous nonhereditary trees. All inheritance tests give a satisfactory result in this case, i.e., a very small fraction of trees passing the test (below 0.05, dotted line); we report a sample tree in the inset. (b) On the set of nonhomogeneous nonhereditary trees the GW inheritance test produces a large number of false positives, $\Phi_{\text{GW}} \approx 0.6$. The tree in the inset is a vivid example: It has a probability of inactivity that decays rapidly with the generations, which GW is unable to account for, thus giving a low P value. Instead, both NHGW and RR perform well on this set. (c) On the set of hereditary trees, the RR method decisively outperforms the other two: Φ_{RR} reaches the exact asymptotic value quite rapidly, while Φ_{NHGW} saturates at a much smaller value and Φ_{GW} never grows. For each tree the P value is calculated from $\mathcal{R} = 10^3$ extractions; for sets (a) and (b) we generate 10^3 trees; for (c)—given its relevance—we generate 10^4 trees. In compliance with

case the test should be negative) and by an hereditary model (in which case the test should be positive). We must be careful, though: At *finite size*, a tree generated by a nonhereditary model has a finite probability to have an hereditary structure and *vice versa*; any stochastic model can produce these cross-cases when the trees' size is small enough. Therefore, evaluating the methods only on trees of one size (for example, $k_{\max} = 7$, as in our experimental dataset) is unwise. It is only in the *infinite size* limit, $k_{\max} \rightarrow \infty$, that the hereditary character of the generating model must emerge. Hence the correct way to proceed is to generate a testbed of lineages using both nonhereditary and hereditary models, submit them to the different inheritance tests, and check how the results scale with the size of the trees.

We first benchmark the methods on a testbed of homogeneous nonhereditary trees, which we generate through the GW model with $q = 0.1$. For each method we measure the fraction Φ of trees that pass the inheritance test, i.e., such that $P < 0.05$; we then plot Φ as a function of the trees' size, $k_{\max}$. Results are reported in Fig. 4(a): A sound inheritance test must produce a small fraction of trees passing the test; more precisely, because we are using 0.05 as a significance threshold for the P value, any correct method should give $\Phi \leq 0.05$. The results show that all three methods perform fine on homogeneous nonhereditary trees.

Things start to change when we benchmark the methods on a set of nonhomogeneous nonhereditary trees, which we generate by using the NHGW model with a generation-dependent probability of inactivity, $q_k = q_0 \exp(-k/k_0)$, with $q_0 = 0.8$ and $k_0 = 1.1$. Because this is still a nonhereditary set of trees, we expect all methods to give $\Phi \leq 0.05$; however, we see from Fig. 4(b) that while this is true for the NHGW and RR methods, the inheritance test based on the GW null ensemble creates a very large number of false positives (around 60%). This is what we anticipated: Many nonhereditary but nonhomogeneous trees [as the stump tree in the inset of Fig. 4(b)] are very rare in the homogeneous GW ensemble, thus giving a low P value, and therefore a false positive. Other trees of this type, giving false positives for the inheritance test based on the GW ensemble, can be found in the tree zoology reported in Sec. 9 in SM [36].

Finally, we benchmark the methods on a set of hereditary trees, which we generate through a simple model: A wild cell has probability 0 to be inactive and probability p to mutate; a mutated cell *and all its descendants* have probability q to be inactive. As we shall see in Sec. VI, this model captures several features of our biological dataset, and hence we generate the trees by using the mean parameters fitted to the biological lineages, $p = 0.1$ and $q = 0.46$ (see Sec. VI). Results are reported in Fig. 4(c). Both the NHGW and RR methods give a fraction of hereditary trees which increases with size; however, Φ_{RR} grows dramatically faster than Φ_{NHGW} . More importantly, we can compare these curves to the ground truth: The asymptotic limit of Φ for $k_{\max} \rightarrow \infty$ can be calculated exactly for this hereditary model (see Sec.

the experimental dataset, only trees in which at least two cells arrive at the last generation have been used to benchmark the methods.




FIG. 5. Reshuffling BMSC lineages. Two examples of how the hereditary structure of the inactive cells is erased by the RR procedure and how this impacts on the inheritance entropy. Top: Lineage 20241203_03 (left) has a strong concentration of inactive cells, and therefore of missing cells, in the west wing of the lineage; in this tree the imbalance is generated at the very first mitosis (the central node). Instead, a randomly scrambled version of the biological tree (center) has inactive cells distributed evenly across all branches. The entropy of the biological tree is $S_{\text{biological}} = 0.81$, while the entropy of one of its scrambled versions is $S_{\text{scrambled}} = 2.32$. When we produce 10^6 scrambled trees of 20241203_03, we very rarely find $S_{\text{scrambled}} \leq S_{\text{biological}}$; this means that the probability that the entropy of the biological case is so small by pure chance is extremely low. The distribution of the randomly reshuffled entropies (right) is shown as a violin-plot on the right: It is evident that the biological value of the entropy is way below the bulk of the distribution of the reshuffled entropies. The P value is simply the integral of this distribution below $S_{\text{biological}}$, which gives $P = 8.2 \times 10^{-5}$, making the inheritance test for colony 20241203_03 highly significant. Bottom: A similar situation holds for lineage 20230503_13, which also gives a very strong inheritance signal.

5.6 in SM [36]); for these parameters we have $\Phi_{\infty} = 0.985$. Figure 4(c) shows that already at $k_{\text{max}} = 14$ the RR method is remarkably close to this exact value, $\Phi_{\text{RR}}(14) = 0.965$, a 2% error with respect to the ground truth. On the contrary, the data show that the NHGW method plateaus at a value much smaller than the exact limit, $\Phi_{\text{NHGW}}(14) = 0.214$. We have repeated this benchmarking at different values of the parameters of the hereditary marking ($p = 0.08$, $q = 0.4$) and found a very similar scenario: Φ_{RR} converges to the ground truth with a 0.3% error, while Φ_{NHGW} remains off by more than 50%. Finally, the GW method always underachieves dreadfully.

The reasons why GW and NHGW perform poorly on the set of hereditary trees are discussed and illustrated at length in Sec. 4 in SM [36]; here we just give a succinct version of the argument. Irrespective of the potential hereditary correlations between inactive cells, the entropy S grows with their total number, N_{inact} . While sampling their respective ensembles, the GW and NHGW models explore many trees with numbers of inactive cells significantly different from that of the input hereditary tree they have been fitted to; such large fluctuations of N_{inact} produce in turn large fluctuations of S , introducing many small entropies in the ensemble, hence biasing to the left the probability distribution $P(S)$, ultimately raising the P value. Instead, all trees in the RR ensemble have the same

number of inactive cells at each generation as the input hereditary tree, hence the entropy S changes only due to the *bona fide* changes in their hereditary correlations. An extensive zoology of trees generated with all models, together with their relative P values, can be found in Sec. 9 in SM [36].

In conclusion, the RR method emerges as the most effective and sensitive entropy-based test to assess the degree of inheritance in a lineage tree, irrespective of the type of tree. Luckily, RR is also by far the simplest to implement and the fastest to run of the three methods we scrutinized: It does not rely on any model, it does not require any calculation, it simply employs a large number of random reshufflings of the lineage tree that needs to be tested. Therefore, in the following we will use the RR entropy-based inheritance test to analyze the biological database of BMSC clonal colonies.

V. RESULTS

We first illustrate how the inheritance test works on the two experimental BMSC colonies in Fig. 5: The random reshuffling clearly destroys all correlations between inactive cells, yet keeping fixed their number at each generation, as well as the number of missing cells. The bulk of the reshuffled distribution of the entropy in these two trees is way above




FIG. 6. Inheritance of the BMSC colonies. Result of the entropy-based inheritance test, using the RR null ensemble, for all colonies in the dataset for which the test is nontrivial (i.e., for lineages with more than one inactive cell). For each colony, the violin plot represents the probability distribution of the entropy for the set of $\mathcal{R} = 10^6$ randomly reshuffled trees, while the black line is the value of the entropy of the original experimental tree; the P value is the integral of the probability below the black line. Green: Significant result of the inheritance test ($P < 0.05$). Gray: Nonsignificant result of the inheritance test ($P \geq 0.05$).

the actual biological value, hence giving a strong inheritance signal in both cases.

A. Inheritance in the experimental BMSC lineages

The outcome of the inheritance test on our entire BMSC dataset is reported in Fig. 6. We have 32 colonies in total; 4 colonies have either zero or one inactive cell, and hence the scrambling procedure always produces the same tree, trivially giving $P = 1$. Of the 28 colonies for which we can fruitfully run the test, 21 (75%) have lineages with a significant hereditary character ($P < 0.05$), while in 7 colonies (25%) the inheritance test gives a nonsignificant result ($P \geq 0.05$). The nonsignificant cases are mostly given by colonies with a very small number of inactive cells, for which the inheritance test—which is based on scrambling in many *different* ways the positions of the inactive cells—is not expected to be very effective: The four colonies with $P > 0.1$ have either 2 (20241112_02, 20250328_16) or 3 (20241017_16, 20250328_02) inactive cells; two colonies with 5 (20241203_09) and 7 (20230113_08) inactive cells have $0.05 < P < 0.1$; in just one colony in our dataset (20240627_06) the inheritance test gives a nonsignificant result ($P = 0.12$) despite having 22 inactive cells. In conclusion, the overall result of the inheritance test indicates that the

topology of the largest majority of human BMSC colonies cannot be attributed to random variations in the positions of inactive cells across different branches, thus proving that the processes responsible for regulating the probability of emergence of inactive cells have a strong hereditary character (in Sec. 7 in the SM we demonstrate that the structure of inactive cells cannot be due to physical proximity [36]). Given the short timescales of the experiment and the low number of generations, it is highly unlikely that such structure is determined by genetic mutations [6,32]; instead, our results are most likely due to the hereditary propagation of nongenetic, that is epigenetic, factors regulating cell-cycle exit.

We stress that a nonsignificant result of the test *does not* imply that there is no inheritance at work but merely that there is no detectable change in the probability of inactivity in the lineage: The lack of variation in the expression of a character makes inheritance undetectable, not necessarily absent; hence, the test is a sufficient, but not necessary, condition for the presence of an hereditary mechanism regulating inactivity. Were we to consider more complex factors, it could very well turn out that a lineage that did not pass the inheritance test at the bare topological level, still results hereditary according to other factors. But we cannot help noticing that calculating the entropy and its P value is easy and quick, as one needs nothing but the bare structure of the lineage. Hence, it seems lucky,

and interesting, that most BMSC colonies display a sharp hereditary character even in relation to the quite fundamental cell state of inactivity.

B. Navigating the imbalance map

As we have already remarked, the value of the inactivity imbalance defined on each node of the tree, i.e., on each mitosis of the colony, provides a sort of map of the most likely sites where mutation events occurred across the tree. We have to be careful about one point, though: When there is a chain, crossing different generations, of subsequent nodes all with similar anomalous imbalances, it is the node at the *latest* generation (largest value of k) that actually corresponds to the mutation event.

In order to understand this point let us go back to Fig. 3. In this highly hereditary lineage, the largest part of the imbalance is concentrated on just two nodes: (i) the central mitosis, which gives rise to the two cells of generation 1, let us call this node M1; (ii) and one of the two mitosis separating generation 1 from generation 2, let us call this node M2. We claim that the actual mutation occurred in M2, while the spike in M1 is simply a byproduct of that in M2. The reason is that the significance of a spike in the imbalance must be assessed going *upstream* in the lineage, starting from the terminal cells, i.e., the k_7 leaves. If we do this, then we see that the first anomalous node that we meet is indeed M2: Something clearly happened at this mitosis, because its right branch (no inactive, nor missing cells) is completely different from its left branch (which instead has 10 inactive cells, causing 22 missing cells), and this large left-right asymmetry occurs for the first time at this node, when we navigate the tree upstream. On the other hand, if no other mutation occurs at M1, then this node would still display roughly the same large imbalance as M2, up to small random fluctuations, which is exactly what happens in this colony. Reversing the order of the argument, i.e., starting from the central node and navigating the lineage *downstream*, each time we find a mitosis with a very large anomalous imbalance, we can only conclude that in one mitosis of its descendants there is a mutation.

C. The inactivity mutation lag

Once we have established a procedure to identify the mitosis where there is a change in the probability of inactivity, an inspection of the imbalance map in those lineages that pass the inheritance test shows that the mutation ultimately leading to the anomalous accumulation of inactive cells within a branch, typically happens a couple of generations *before* these cells start emerging in that branch (Fig. 7). Inactive cells may either gradually appear in the mutated branch across different generations (as in 20230503_13 or 20241203_03), or they may show up all at the same generation (20250328_11); moreover, in some clones there are two independent mutation events (20240627_03). Due to this complex scenario, there might be more than one way to quantify the inactivity mutation lag, l ; we stick to the simplest definition and count how many generations (i.e., mitosis, or nodes) separate the mutation node from the *first* inactive cell emerging after the mutation (see Fig. 7). The average value of the mutation lag over all BMSC clones passing the inheritance test is $l_{\text{BMSC}} \approx 2$ generations.

The existence of this mutation lag indicates that what we hypothesized in the Introduction, namely that the epigenetic change responsible for the emergence of inactivity occurs a few generations *before* inactivity itself is expressed, is actually true. As we have already noted, this phenomenon is consistent with results reported in the literature about cellular senescence [19]. Moreover, in a different context, the presence of generation delays between cell fate decisions and the onset of lineage markers has been found in hematopoietic stem cells [52]. We also notice that our method is able to capture the progression to inactivity (and quite possibly senescence) on the short timescales of the first proliferation stages of the clonal colony, contrary to most studies, which rely on mass population analysis across multiple passages and focus on the long timescales [53].

Finally, we stress once again that it is only thanks to the inactivity mutation lag that the hereditary nature of cell-cycle exit could be established: If the inactive cell was born *exactly* at the same mitosis where the epigenetic change responsible for inactivity occurred ($l = 0$), then it would be impossible to assess inheritance. This seems at odds with the fact that—lag or not—if the number of k_7 leaves cut by an inactive cell is large, then the imbalance will also be large, thus lowering the entropy, which should give a strong inheritance signal. In fact, the situation is more subtle. Imagine a tree in which there is just one mutation and imagine that that mutation occurs at the central node of the tree, the one producing the two $k = 1$ cells; imagine also that *there is no mutation lag*, $l = 0$, hence (say) the right cell immediately becomes inactive, thus cutting all its 64 potential k_7 leaves, while its sister—the left cell—generates a full half-tree, ultimately producing all its 64 k_7 leaves [this is exactly the tree represented in Fig. 4(b)]. In this case, a large inactivity imbalance, $I = 64$, is entirely concentrated on just that central mitosis, which gives a very small entropy, in fact the smallest, $S = 0$. However, when we attempt to randomly reshuffle this tree, we immediately recognize that the only move we can make is to switch the two $k = 1$ sisters, which produces exactly the same lineage; hence, *all* reshuffled trees have $S = 0$, so that the P value with the RR method is equal to 1 and we have zero inheritance signal. This is a general phenomenon: The absence of a mutation lag would imply that the locations of inactive cells coincide with the locations of the mutation events; if these events are uncorrelated from each other (and there are no strong reasons to think otherwise), then reshuffling the inactive cells creates the same statistical lineage and the entropy is not lower in the RR set, on average. This argument highlights very vividly that the *absolute* value of the entropy can trick us; only its value *relative* to the null ensemble of lineages is significant: A lineage is highly hereditary not because its entropy is low, but because it is *lower* than that of the corresponding nonhereditary ensemble, in our case RR. In this respect, the reshuffling procedure is as crucial an ingredient of the inheritance test as the entropy itself.

VI. A SIMPLE HEREDITARY MODEL OF THE DATA

The existence of a lag between the epigenetic change leading to inactivity and its expression, is very reminiscent of the *incubation period* introduced by the commitment theory




FIG. 7. Mutation lag in BMSC lineages. Each mitosis m is colored according to its normalized imbalance: (black) $w_m = 0$, (yellow) $0 < w_m \leq 0.15$, and (red) $w_m > 0.15$. The mitosis whose anomalous imbalance signals a mutation in the probability of inactivity is marked by an extra external black circle and the shaded area highlights the branch impacted upon by that mutation; the first (most upstream) link in this branch corresponds to the actually mutated cell. On the other hand, the first inactive cell in the mutated branch emerges a few generations downstream and it is marked by an extra external black rectangle. The mutation lag, l , is given by the number of generations, i.e., of nodes, separating the mutated cell from of the first inactive cell. In colony 20240627_03 there are two independent mutation events. Colony 20250328_11 is a more complex case: It seems that there are two concatenated mutation events, one at the very first mitosis in the lineage, differentiating the left and right wings of the tree, and another (stronger) event separating generation $k = 3$ from $k = 4$ (the one highlighted in the panel); a more in depth analysis is needed to work out whether these two events are correlated.

of cellular aging of Kirkwood and Holliday [37–39], whose goal was to model the evolution of senescence in human diploid fibroblasts. There is an important difference with our experimental findings, though: According to that theory, after the incubation period, *all* the descendants generated by the committed cell die out, while this is clearly not the case in our data. Instead, BMSC lineages suggest that at the mutation there is no deterministic commitment to inactivity, but rather a sudden increase in the probability to produce inactive cells; this probability is passed-on in a hereditary way and it acquires a phenotypic impact only a few generations later, when enough cells have been generated to produce some inactive descendants. This state of affairs suggests a simple stochastic generalization of the Kirkwood-Holliday hereditary model.

A. The stochastic Kirkwood-Hollyday model

As in the original Kirkwood-Holliday model [37], we assume that a wild cell has probability 0 to be inactive and probability p to mutate; once a cell mutates, all its descendants have probability q to be inactive, including the mutated cell itself (the descendants of a mutated cell can only become inactive, they cannot mutate again). Hence, at the mutation

there is a sharp change in the probability of inactivity, from 0 to q , and this change is inherited by the progeny of the mutated cell, whence the hereditary character of the model. In contrast with the original Kirkwood-Holliday model, not only the mutation, but also the emergence of inactivity is a stochastic process, which prevents the unrealistic case of perfect synchronization in the emergence of inactivity. We call this the stochastic Kirkwood-Holliday model (SKH).

The readers can easily convince themselves that for $q = 1$ the SKH model reduces to the GW model with probability p : This is exactly the case of *zero* mutation lag that we described at the end of the previous section; moreover, for $p = 1$ SKH reduces to GW with probability q : Once the root has mutated, all proceeds as in GW. Therefore, although SKH is in general an hereditary model, it contains also some nonhereditary sectors in the space of parameters. In fact, the nonhereditary region of SKH extends far beyond the two GW edges. For each pair (p, q) in a grid, we compute the fraction Φ of SKH trees that pass the inheritance test, hence drawing an hereditary map of the model (Fig. 8): This map shows that the most hereditary region of SKH is an oval-shaped zone around $p = 0.1$ and $q = 0.45$: Here the SKH model produces up to 35% of trees that pass the entropy-based inheritance test.




FIG. 8. Hereditary map of the SKH model. For each pair of parameters (p, q) we generate 10^3 SKH trees with $k_{\max} = 7$; for each tree we run the inheritance test with $\mathcal{R} = 10^3$ reshufflings and measure the fraction Φ of trees that pass the test; a larger value of Φ (green) represents a more hereditary region of the model, while a smaller value of Φ (red) represents a less hereditary one. Contour lines start at $\Phi = 0.01$ (up-right) and have spacing $\Delta\Phi = 0.01$; their steep density in the green zone shows how sharp is the inheritance sector of SKH. The star marks the position of the mean MLE fit of the SKH model to the experimental colonies.

The two GW nonhereditary edges of SKH are symmetric, ($p = 1 \forall q$) and ($q = 1 \forall p$); besides, the mean number of inactive cells is also symmetric, $N_{\text{inact}}^{\text{SKH}}(p, q) = N_{\text{inact}}^{\text{SKH}}(q, p)$, and the same happens for the mean number of missing cells, $N_{\text{miss}}^{\text{SKH}}$ (see Secs. 5.3 and 5.4 in SM [36]). However, the hereditary map is strongly *nonsymmetric*, hence suggesting that the two ways to recover GW from SKH are very different from each other. This is indeed the case. The nonhereditary limit $p \rightarrow 1$ is smooth: A tree with a very large probability of mutation, $p = 1 - \epsilon$, has a mutated root almost all the times, hence giving rise to a GW model, which has a very low probability to pass the inheritance test. On the contrary, the other nonhereditary limit, $q \rightarrow 1$, is *not* smooth, as any tree with $q = 1 - \epsilon$ is still hereditary in the infinite-size limit, because in a large-enough tree sooner or later inactivity will express not right at the mutation, giving rise to an hereditary signal; only for $q = 1$ the two models coincide. A finite-size echo of this interesting property shows up in Fig. 8, in which the strip $p = 0.9$ produces trees significantly less hereditary than the strip $q = 0.9$.

B. Fitting the SKH model to the biological lineages

Given an experimental tree, $\mathcal{T}$, we want to find the parameters (p_τ, q_τ) that give the best SKH fit to the data; we do this in three different ways. We first use a maximum likelihood estimate (MLE): we calculate the likelihood of the SKH model given the tree and check whether or not it has a well-defined maximum (p_τ, q_τ) ; this happens only for experimental lineages with a large number of inactive cells, which are also those trees that pass the entropy-based inheritance

test; BMSC trees with small N_{inact} have a flat SKH likelihood profile, which makes it hard to accurately locate the maximum (Sec. 5.7.1 in SM [36]). The second method uses the inactivity imbalance to locate the mutation within the experimental tree and employs MLE restricted to the mutated subtree to calculate q ; then, it fits p by matching the total number of inactive and missing cells to the analytic SKH expressions. The third method fits q by matching the mutation lag of the experimental tree to the analytic SKH expression, while p is fixed as in the second method. Also these last two methods work only on BMSC colonies that pass the entropy-based inheritance test, the sole for which it is possible to locate the mutations. Full analytic details of the three methods are reported in Sec. 5.7 in SM [36]. Since MLE is the gold standard of model inference, it is reassuring that the other two fitting methods – which rely heavily on our definitions of imbalance, mutation and mutation lag – give results consistent with MLE (see Table 1 in SM [36]). Finally, if we average the log-likelihood over all BMSC trees, then we get $p^* = 0.1$, $q^* = 0.46$, the parameters used in Fig. 4(c) to benchmark the inheritance tests; (p^*, q^*) falls right at the center of the hereditary SKH zone.

The reader may wonder why we do not report on the SKH inheritance map the positions of the individual BMSC lineages; if we did, then we would find that all experimental colonies comfortably tuck in the hereditary region of SKH, which could be interpreted as further proof of their hereditary nature. In fact, doing this would be a mistake: assessing the inheritance of a tree by fitting it to the SKH model and comparing the fit to the hereditary region of SKH, may lead to wrong conclusions. We show in Sec. 6 in the SM [36] that there are ensembles of nonhereditary trees that are nevertheless located by an MLE fit in proximity of the hereditary sector of SKH. The fact is that, fitting the SKH model to a tree does not necessarily prioritize inheritance as the main fitted trait. An example of nonhereditary tree that sneaks into the hereditary zone of SKH is the following: consider a *nonhomogeneous* tree Ξ with inactive cells all concentrated at the last generations (see Fig. S5 in SM [36]). The nonhereditary region of SKH has probabilities of mutation and inactivity quite large and—because it is an *homogeneous model*—SKH produces in this zone many inactive cells also at the early generations, making it very unlikely for a tree to survive to the last generation. Hence, the likelihood finds it more appropriate to fit Ξ through SKH parameters able to generate complete trees, indifferent to the fact that these parameters fall in the hereditary region.

This is a deeper issue than it seems. If instead of naively gauging the MLE fit on the hereditary SKH map, then we employ the more serious parametric bootstrap likelihood-ratio test [54–56] to compare the SKH and GW likelihoods, we would still wrongly conclude that the nonhereditary tree Ξ is instead hereditary (SM-6 [36]). This is a stark reminder that any inheritance-assessment method based on comparing the likelihoods of different models, inevitably relies on the arbitrary choice of these models: There might always exist lineages so alien to the structure of the selected models, to fool the likelihood-based methods. Instead, the entropy-based inheritance test is *completely independent* from the choice of models, so it correctly classifies Ξ as nonhereditary.




FIG. 9. Sample lineages of the hereditary SKH model. The parameters (p, q) of each one of these trees are equal to those fitted to the four biological trees in Fig. 7, so to make as fair a comparison as possible between the two figures (for the actual parameters (p, q) see Table 1 in SM [36]). The hereditary structure of the SKH trees is quite evident even by eye and indeed all four lineages pass the inheritance test ($p^{\text{SKH1}} = 0.0261$, $p^{\text{SKH2}} = 0.0173$, $p^{\text{SKH3}} = 0.0003$, $p^{\text{SKH4}} < 10^{-6}$).

C. Insights from the SKH model

How do SKH trees *look* like? In Fig. 9 we represent four SKH trees, each one drawn with the MLE parameters (p, q) relatives to the four biological trees of Fig. 7, so to make a fair comparison. We see that the SKH trees look indeed quite hereditary, with clustering of inactive cells along the same branches, very much as in the biological dataset. If the reader checks the tree zoology reported in Sec. 9 in SM [36], then they will further appreciate the qualitative similarity between SKH and BMSC trees. To evaluate in a slightly more quantitatively fashion the quality of the SKH fit to the BMSC data, we can calculate the mutation lag correlation: after we fit the SKH model to an experimental lineage $\mathcal{T}$ through MLE, we use the fitted parameters to calculate the mutation lag, l_{SKH}^{τ} , for that tree, using the analytic relations for the distribution of the mutation lag in the SKH model that we calculate in Sec. 5.5 in SM [36]; besides, we determine the lag of the experimental tree, l_{exp}^{τ} , by using the peaks of the imbalance, as described in the previous section. We find that the two sets of lags have decent Spearman correlation, $\rho = 0.67$, and very significant P value, $P = 0.001$ (10^6 extractions).

Hence, it seems that SKH is a sound starting point to describe inheritance in our data; more specifically, our interpretation of BMSC imbalance maps, namely that sometimes along the lineage there is a mutation of probability to become inactive, seems very reasonable, given the SKH fit to the experimental data. In this respect, it is interesting to note that the mean MLE fit to SKH of the biological lineages indicates

that the probability of mutation is quite low, $p \approx 0.1$, but that—after the mutation occurs—the probability to become inactive is quite more substantial, $q \approx 0.45$. One could have envisaged an opposite scenario, in which mutations occurred more frequently, but the mutated inactivity probability was lower, $p \approx 0.45$ and $q \approx 0.1$; what the SKH inheritance map shows is that this scenario would have given a far smaller inheritance signature, so small to go probably undetected by the inheritance test.

The SKH model has also limitations, though. The inheritance level of the SKH model is lower than that of BMSC lineages; as we have seen in Fig. 8, even in the most hereditary sector of the parameters space, about 35% of SKH trees (with $k_{\text{max}} = 7$) pass the inheritance test, against the 75% of hereditary trees in the experimental dataset. Moreover, the correlation between the mutation lags (inferred with SKH versus experimental) is good, but not great. In fact, we believe that both these discrepancies could be overcome by introducing more articulated (and complex) versions of the model. Instead of having a constant probability of inactivity q after the mutation, which produces an exponential distribution of the mutation lag (see Sec. 5.5 in SM [36]), one can formulate a model in which the inactivity probability depends on the distance Δk (in generations) from the mutation. In this way one can easily produce a nonexponential probability distribution of the mutation lag, with the desired mean and variance; by lowering the variance one can obtain as large an inheritance signal as desired. Of course, more parameters means more

problems to fit the model, so one should balance simplicity versus accuracy, as always. Exploring such broader class of SKH hereditary models goes beyond the scope of this work, but our results seem already to suggest that SKH is a very reasonable class of models for describing the inheritance of inactivity in cell lineages.

VII. CONCLUSIONS

We introduced a new entropy-based method to assess inheritance in a lineage tree. The method is model agnostic, as it relies solely on the topological structure of the truncated branches, namely of those elements of the tree that have stopped proliferating. In the case of cell lineages, truncation of a branch is due to inactivity or death of a cell, but in general the interruption of a branch may be due to any phenomenon impacting on the growth of the tree, irrespective of the specific origin and context of the lineage. In this respect, our method—which we proved to be reliable on a great number of different synthetic trees—seems to have a strong potential for very broad applications, well beyond the assessment of epigenetic inheritance in stem cell colonies.

The entropy-based inheritance test shows that human BMSC clonal colonies have a strong hereditary structure of their inactive cells: Seventy-five percent of colonies pass the inheritance test, and those that do not are characterized by such low numbers of inactive cells to make the test void. It is important to note that such evidence of inheritance—which we obtain from the inactivity imbalance—is due to the existence of *intracolony* heterogeneities: It is only thanks to the emergence throughout the lineage of changes in the propensity to become inactive that we can detect the anomalous imbalances that contribute to lowering the entropy; on the contrary, colonies with no inactive cells do not give any signal; as we have already noted, inheritance is undetectable in absence of mutations. Once the presence of inheritance is established through the intracolony fluctuations, it is legitimate to infer that the same hereditary mechanism is responsible for the colony-level permanence of epigenetic mutations, thus giving rise to *intercolony* heterogeneities. Hence, our results support the hypothesis that the strong colony-level heterogeneities met when dealing with *in vivo* transplant of BMSCs—heterogeneities that strongly affects the effectiveness of transplants and thus of skeletal regeneration therapies—do indeed have an epigenetic hereditary origin.

The imbalance map of a colony allows to identify the mutation in the probability of inactivity, a mutation then inherited by the descendants, hence causing the hereditary structure of the lineage. We find that in BMSC colonies the mutation occurs, on average, two generations before the expression of inactivity. The observation of this mutation lag prompted the development of a simple stochastic variant of the Kirkwood-Holliday model (SKH), which has an interesting hereditary landscape in the space of parameters. The SKH model produces synthetic trees that—despite being less hereditary than

the BMSC dataset—have a strong topological similarity with the experimental lineages. An MLE fit of the experimental data to the SKH model suggests that the epigenetic hereditary structure of cell-cycle exit in BMSCs is due to a relatively low probability of mutation, associated to a relatively high probability of inactivity after the mutation. Generalizations of the SKH model seem easy to implement, hence paving the way to even more accurate characterizations of epigenetic inheritance in stem cells lineages.

The distribution of inactive cells across a colony determines the inheritance entropy. Active cells, on the other hand, constitute a separate population that enters the entropy-based hereditary test only in a passive way. Experiments, though, shows that the topology of BMSC lineages (determined by inactive cells) is strongly correlated to the statistics of the division times of the active cells [35]: Colonies with the larger populations of inactive cells are also characterized by the slower populations of active cells [35]. This result suggests that not only cell-cycle exit, but also cell-cycle duration might be regulated by epigenetic hereditary mechanisms. To confirm this point, though, it is necessary to study the correlations of the division times. This has been done for bacterial colonies in relation to genetic factors [33,57], and to a scarcer degree for epigenetic factors in eukaryotic cells [3,58], but never in the context of stem cells. Besides, in all past studies, the assessed correlation was not quite deep enough to trace in a robust way the hereditary ramifications of the division times across the lineage tree. This remains therefore an interesting open experimental and theoretical problem along the path to establishing epigenetic factors as the primary cause of the strong heterogeneity affecting stem cell colonies.

ACKNOWLEDGMENTS

We thank William Bialek, Fabio Cecconi, Edo Kessel, Leonid Mirny, Felix Ritort, Aleksandra Walczak, and two anonymous referees for very insightful advice on the manuscript. A.C. and T.S.G. thank Enzo Branchini for discussions within the CC. Andrea Cavagna acknowledges the support and advice of Francesco Cavagna and of the late Giovanni Cavagna. This work was supported by European Research Council (ERC) Grant No. RG.BIO (Contract No. 785932), MIUR Grant INFO.BIO (Protocol No. R18JNYMEY), and MIUR Grant No. PRIN2020 (Contract No. 2020PFCXPE-005).

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

- [1] R. K. Ng and J. B. Gurdon, Epigenetic inheritance of cell differentiation status, *Cell Cycle* **7**, 1173 (2008).
- [2] E. H. Zion, C. Chandrasekhara, and X. Chen, Asymmetric inheritance of epigenetic states in asymmetrically dividing stem cells, *Curr. Opin. Cell Biol.* **67**, 27 (2020).
- [3] O. Sandler, S. P. Mizrahi, N. Weiss, O. Agam, I. Simon, and N. Q. Balaban, Lineage correlations of single cell division time as a probe of cell-cycle dynamics, *Nature (London)* **519**, 468 (2015).
- [4] N. Mosheiff, B. M. Martins, S. Pearl-Mizrahi, A. Grünberger, S. Helfrich, I. Mihalcescu, D. Kohlheyer, J. C. W. Locke, L. Glass, and N. Q. Balaban, Inheritance of cell-cycle duration in the presence of periodic forcing, *Phys. Rev. X* **8**, 021035 (2018).
- [5] S. Hormoz, Z. S. Singer, J. M. Linton, Y. E. Antebi, B. I. Shraiman, and M. B. Elowitz, Inferring cell-state transition dynamics from lineage trees and endpoint single-cell measurements, *Cell Syst.* **3**, 419 (2016).
- [6] H. H. Chang, M. Hemberg, M. Barahona, D. E. Ingber, and S. Huang, Transcriptome-wide noise controls lineage choice in mammalian progenitor cells, *Nature (London)* **453**, 544 (2008).
- [7] M. Plambeck, A. Kazerooni, D. Loeffler, L. Kretschmer, C. Salinno, T. Schroeder, D. H. Busch, M. Flossdorf, and V. R. Buchholz, Heritable changes in division speed accompany the diversification of single T cell fate, *Proc. Natl. Acad. Sci. USA* **119**, e2116260119 (2022).
- [8] M. Yampolskaya, L. Ikonou, and P. Mehta, Finding signatures of low-dimensional geometric landscapes in high-dimensional cell fate transitions, *arXiv:2506.04219*.
- [9] D. E. Wagner and A. M. Klein, Lineage tracing meets single-cell omics: Opportunities and challenges, *Nat. Rev. Genet.* **21**, 410 (2020).
- [10] C. S. Baron and A. van Oudenaarden, Unravelling cellular relationships during development and regeneration using genetic lineage tracing, *Nat. Rev. Mol. Cell Biol.* **20**, 753 (2019).
- [11] C. Weinreb, A. Rodriguez-Fraticelli, F. D. Camargo, and A. M. Klein, Lineage tracing on transcriptional landscapes links state to fate during differentiation, *Science* **367**, eaaw3381 (2020).
- [12] B. Raj, Single-cell profiling of lineages and cell types in the vertebrate brain, in *Lineage Tracing: Methods and Protocols* (Springer, Berlin, 2025), pp. 299–310.
- [13] B. Spanjaard, B. Hu, N. Mitic, P. Olivares-Chauvet, S. Janjuha, N. Ninov, and J. P. Junker, Simultaneous lineage tracing and cell-type identification using CRISPR–Cas9-induced genetic scars, *Nat. Biotechnol.* **36**, 469 (2018).
- [14] R. J. Ihry, K. A. Worringer, M. R. Salick, E. Frias, D. Ho, K. Theriault, S. Kommineni, J. Chen, M. Sondey, C. Ye, *et al.*, p53 inhibits CRISPR–Cas9 engineering in human pluripotent stem cells, *Nat. Med.* **24**, 939 (2018).
- [15] U. Deichmann, Epigenetics: The origins and evolution of a fashionable topic, *Dev. Biol.* **416**, 249 (2016).
- [16] E. Jablonka and M. J. Lamb, The changing concept of epigenetics, *Ann. N.Y. Acad. Sci.* **981**, 82 (2002).
- [17] V. V. Lunyak and M. G. Rosenfeld, Epigenetic regulation of stem cell fate, *Hum. Mol. Genet.* **17**, R28 (2008).
- [18] A. Meissner, Epigenetic modifications in pluripotent and differentiated cells, *Nat. Biotechnol.* **28**, 1079 (2010).
- [19] M. Ogrodnik, Cellular aging beyond cellular senescence: Markers of senescence prior to cell cycle arrest in vitro and in vivo, *Aging Cell* **20**, e13338 (2021).
- [20] C. M. Koch, S. Joussen, A. Schellenberg, Q. Lin, M. Zenke, and W. Wagner, Monitoring of cellular senescence by DNA-methylation at specific CpG sites, *Aging Cell* **11**, 366 (2012).
- [21] J. Franzen, A. Zirkel, J. Blake, B. Rath, V. Benes, A. Papantonis, and W. Wagner, Senescence-associated DNA methylation is stochastically acquired in subpopulations of mesenchymal stem cells, *Aging Cell* **16**, 183 (2017).
- [22] J. Franzen, T. Georgomanolis, A. Selich, C.-C. Kuo, R. Stöger, L. Brant, M. S. Mulabdić, E. Fernandez-Rebollo, C. Grezella, A. Ostrowska, *et al.*, DNA methylation changes during long-term in vitro cell culture are caused by epigenetic drift, *Commun. Biol.* **4**, 598 (2021).
- [23] D. Nachtwey and I. Cameron, Cell cycle analysis, in *Methods in Cell Biology* (Elsevier, Amsterdam, 1969), Vol. 3, pp. 213–259.
- [24] T. H. Cheung and T. A. Rando, Molecular regulation of stem cell quiescence, *Nat. Rev. Mol. Cell Biol.* **14**, 329 (2013).
- [25] C. Behl and C. Ziegler, Cell cycle: The life cycle of a cell, in *Cell Aging: Molecular Mechanisms and Implications for Disease* (Springer, Berlin, 2013), pp. 9–19.
- [26] M. M. Mens and M. Ghanbari, Cell cycle regulation of stem cells by microRNAs, *Stem Cell Rev. Rep.* **14**, 309 (2018).
- [27] P. Bianco and P. G. Robey, Skeletal stem cells, *Development* **142**, 1023 (2015).
- [28] P. Bianco, X. Cao, P. S. Frenette, J. J. Mao, P. G. Robey, P. J. Simmons, and C.-Y. Wang, The meaning, the sense and the significance: Translating the science of mesenchymal stem cells into medicine, *Nat. Med.* **19**, 35 (2013).
- [29] S. Mareddy, R. Crawford, G. Brooke, and Y. Xiao, Clonal isolation and characterization of bone marrow stromal cells from patients with osteoarthritis, *Tissue Eng.* **13**, 819 (2007).
- [30] S. Mareddy, N. Dhaliwal, R. Crawford, and Y. Xiao, Stem cell-related gene expression in clonal populations of mesenchymal stromal cells from bone marrow, *Tissue Eng. Part A* **16**, 749 (2010).
- [31] T. Kuczek and D. E. Axelrod, The importance of clonal heterogeneity and interexperiment variability in modeling the eukaryotic cell cycle, *Math. Biosci.* **79**, 87 (1986).
- [32] D. A. Rennerfeldt and K. J. Van Vliet, Concise review: When colonies are not clones: Evidence and implications of intra-colony heterogeneity in mesenchymal stem cells, *Stem Cells* **34**, 1135 (2016).
- [33] E. Powell, Some features of the generation times of individual bacteria, *Biometrika* **42**, 16 (1955).
- [34] A. Genthon, T. Nozoe, L. Peliti, and D. Lacoste, Cell lineage statistics with incomplete population trees, *PRX Life* **1**, 013014 (2023).
- [35] A. Allegranza, R. Beschi, D. Caudo, A. Cavagna, A. Corsi, A. Culla, S. Donsante, G. Giannicola, I. Giardina, G. Gosti, *et al.*, Lineage topology, replication kinetics and cell cycle synchronization reveal regulated growth dynamics in human bone marrow stromal cell colonies, *Sci. Rep.* (2026).
- [36] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/p2mj-q682> for all technical details about the models and for an exhaustive illustration of the lineage trees, which includes Refs. [59–66].
- [37] T. Kirkwood and R. Holliday, Commitment to senescence: A model for the finite and infinite growth of diploid and transformed human fibroblasts in culture, *J. Theor. Biol.* **53**, 481 (1975).

- [38] R. Holliday, L. Huschtscha, G. Tarrant, and T. Kirkwood, Testing the commitment theory of cellular aging: The finite lifespan of human fibroblasts may be due to the decline and loss of a subpopulation of immortal cells, *Science* **198**, 366 (1977).
- [39] C. B. Harley and S. Goldstein, Retesting the commitment theory of cellular aging, *Science* **207**, 191 (1980).
- [40] S. E. Luria and M. Delbrück, Mutations of bacteria from virus sensitivity to virus resistance, *Genetics* **28**, 491 (1943).
- [41] W. A. Rosche and P. L. Foster, Determining mutation rates in bacterial populations, *Methods* **20**, 4 (2000).
- [42] D. H. Colless, Phylogenetics: The theory and practice of phylogenetic systematics, *System. Zool.* **31**, 100 (1982).
- [43] M. Fischer, L. Herbst, S. Kersting, A. L. Kühn, and K. Wicke, *Tree Balance Indices: A Comprehensive Survey* (Springer Nature, Berlin, 2023).
- [44] S. F. Gull and G. J. Daniell, Image reconstruction from incomplete and noisy data, *Nature (London)* **272**, 686 (1978).
- [45] H. Theil, *Economics and Information Theory* (Amsterdam, North-Holland, 1967).
- [46] T. E. Harris *et al.*, *The Theory of Branching Processes* (Springer, Berlin, 1963), Vol. 6.
- [47] J. C. D'Souza and J. Biggins, The supercritical Galton-Watson process in varying environments, *Stochast. Process. Appl.* **42**, 39 (1992).
- [48] G. Kersting, A unifying approach to branching processes in a varying environment, *J. Appl. Probab.* **57**, 196 (2020).
- [49] L. Hayflick, The limited in vitro lifetime of human diploid cell strains, *Exp. Cell Res.* **37**, 614 (1965).
- [50] E. Robbins, E. Levine, and H. Eagle, Morphologic changes accompanying senescence of cultured human diploid cells, *J. Exp. Med.* **131**, 1211 (1970).
- [51] H. Rubin, Cell aging in vivo and in vitro, *Mech. Ageing Dev.* **98**, 1 (1997).
- [52] M. K. Strasser, P. S. Hoppe, D. Loeffler, K. D. Kokkaliaris, T. Schroeder, F. J. Theis, and C. Marr, Lineage marker synchrony in hematopoietic genealogies refutes the pu. 1/gata1 toggle switch paradigm, *Nat. Commun.* **9**, 2697 (2018).
- [53] W. Wagner, P. Horn, M. Castoldi, A. Diehlmann, S. Bork, R. Saffrich, V. Benes, J. Blake, S. Pfister, V. Eckstein, *et al.*, Replicative senescence of mesenchymal stem cells: A continuous and organized process, *PLoS ONE* **3**, e2213 (2008).
- [54] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap*, Monographs on Statistics and Applied Probability (Chapman & Hall, New York, London, 1993).
- [55] A. C. Davison and D. V. Hinkley, *Bootstrap Methods and Their Application* (Cambridge University Press, Cambridge, UK, 1997).
- [56] S. G. Self and K.-Y. Liang, Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions, *J. Am. Stat. Assoc.* **82**, 605 (1987).
- [57] R. Staudte, M. Guiguet, and M. C. d'Hooghe, Additive models for dependent cell populations, *J. Theor. Biol.* **109**, 127 (1984).
- [58] S. Hormoz, N. Desprat, and B. I. Shraiman, Inferring epigenetic dynamics from kin correlations, *Proc. Natl. Acad. Sci. USA* **112**, E2281 (2015).
- [59] M. Tran, A. Askary, and M. B. Elowitz, Lineage motifs as developmental modules for control of cell type proportions, *Dev. Cell* **59**, 812 (2024).
- [60] J. Felsenstein, Evolutionary trees from DNA sequences: A maximum likelihood approach, *J. Mol. Evol.* **17**, 368 (1981).
- [61] R. A. Neher, C. A. Russell, and B. I. Shraiman, Predicting evolution from the shape of genealogical trees, *eLife* **3**, e03568 (2014).
- [62] A. Aragon, A. Lambert, T. Mora, and A. M. Walczak, Learning evolutionary parameters from genealogies using allelic trees, *Genetics* **232**, iyaf112 (2026).
- [63] T. Dieselhorst and J. Berg, Phylogenetics of somatic evolution: A likelihood-based approach for cellular reproduction, *Mol. Biol. Evol.* **43**, msag002 (2026).
- [64] S. S. Wilks, The large-sample distribution of the likelihood ratio for testing composite hypotheses, *Ann. Math. Stat.* **9**, 60 (1938).
- [65] G. Froese, The distribution and interdependence of generation times of HeLa cells, *Exp. Cell Res.* **35**, 415 (1964).
- [66] H. Kubitschek, Normal distribution of cell generation rate, *Exp. Cell Res.* **26**, 439 (1962).