

Revisiting Xia' s Index of Substitution Saturation from a Mathematical Perspective

Zhi-Jie Xu

zjxmolls@outlook.com

Donghua University <https://orcid.org/0009-0003-3855-0921>

Research Article

Keywords: substitution saturation, index of substitution saturation, spectral decomposition, reversible substitution models, phylogenetic signal, statistical power, sequence length, distance matrices

Posted Date: October 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11338476/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Revisiting Xia's Index of Substitution Saturation from a Mathematical Perspective

Zhi-Jie Xu*

October 12, 2026

Abstract

The index of substitution saturation $I_{ss} = \bar{H}/H_{FSS}$ [Xia et al., 2003] decides, before any tree is built, whether an alignment retains phylogenetic signal. We characterize the index exactly, and test that characterization in new simulations and in re-analyses of two published data sets. The entropy enters $\mathbb{E}[I_{ss}]$ only through the spectrum of the rate matrix: the index is a linear functional of the coordinates $\langle e^{\mu d} \rangle$, collapsing onto mean pairwise p -distance exactly when each non-zero eigenvalue's multiplicity-to-composition-weight ratio is constant across eigenvalues. This holds for every reversible model at uniform composition and for degenerate spectra; it fails for non-uniform composition with a non-degenerate spectrum, the regime it was calibrated in. To second order — the order at which the entropy enters — the index adds nothing under the collapse about the number of substitutions realised beyond the uncorrected proportion of differing sites; outside it, what it adds is the model's composition, which its user has not fixed. The mean of λ reduces the same way.

The critical value fails separately: it carries no rate-heterogeneity coordinate, and near saturation the index is pinned within a few percent of its ceiling, a fifteen-fold change in the tree length at which recovery fails moving it by 1.2%. A four-taxon transfer of long branches at fixed patristic total makes this direct: the index is constant along the path while recovery spans its full range.

Real data occupy both: the saturated codon partition of the index's own data set falls between two readings of its threshold, and a genome-scale avian alignment sits five to six times below both. The index is not information-free — in a crossed simulation it predicts recovery where p -distance does not — but what it adds tracks the tree's shape and the higher-order composition of sites, not the substitution count; its threshold is a power threshold carrying its own sequence length. Passing the test is evidence against saturation, failing it is weak evidence, neither is evidence of resolution.

Key words: substitution saturation; index of substitution saturation; spectral decomposition; reversible substitution models; phylogenetic signal; statistical power; sequence length; distance matrices.

1 Introduction

The index of substitution saturation, $I_{ss} = \bar{H}/H_{FSS}$, was introduced by Xia et al. [2003] to answer a question that substitution models address only implicitly: have a set of sequences accumulated so many multiple substitutions that their phylogenetic signal is exhausted? Its construction is

*College of Biological Science & Medical Engineering, Donghua University, Shanghai 201620, China.
Correspondence: zjxmolls@outlook.com

deliberately minimal. $\bar{H}$ is the mean entropy of the base composition at a site, H_{FSS} the entropy expected under full saturation, so the ratio requires only an alignment — no tree, no model fit, no likelihood. That independence is why it is used: I_{ss} can be computed before any phylogeny is built, and compared against a tabulated critical value $I_{ss,c}$ to decide in advance whether a data set is fit for analysis [Xia et al., 2003, Xia and Lemey, 2009].

The two halves of the test are constructed differently, and the difference matters. The index itself is a statistic of the alignment; the critical value is a threshold calibrated by simulation under specified conditions [Xia et al., 2003]. Here we characterise the first exactly, and show that the second fails for reasons that have nothing to do with that characterisation.

We show that the entropy construction enters $\mathbb{E}[I_{ss}]$ only through the spectrum of the rate matrix, which reduces the question of what I_{ss} measures to a spectral one. Resolving the transition matrix on the eigenbasis of the symmetrised generator $S = \Pi^{1/2}Q\Pi^{-1/2}$ makes the index a linear functional of the coordinates $\xi_\mu = \langle e^{\mu d} \rangle$, and singles out the condition under which it collapses onto mean pairwise p -distance: the ratio τ_μ/w_μ must be constant across the distinct non-zero eigenvalues of Q . The condition is met for every reversible model under uniform base frequencies and for every composition when the spectrum is degenerate, and it fails — mildly, by a term we quantify — in the regime used to calibrate I_{ss} . Where it holds, I_{ss} and mean p -distance induce the same ordering of alignments, so the content of the test is carried entirely by the critical value. That value we show to be L -dependent, so that the test statistic grows as $\sqrt{L}$, and in the near-saturated region to be pinned close enough to the index’s upper bound that it cannot track the divergence it is meant to threshold.

2 Results

2.1 The statistic and its normalisation

Let N sequences of length L be related by a tree $\mathcal{T}$ with leaf states generated by a reversible rate matrix Q stationary with respect to $\boldsymbol{\pi}$, normalised so that $-\sum_j \pi_j Q_{jj} = 1$. Write $\hat{\mathbf{p}}^{(i)}$ for the empirical state frequencies at site i , $\hat{p}_j$ for their average over sites (the pooled frequencies), and

$$H_i = h(\hat{\mathbf{p}}^{(i)}), \quad h(\mathbf{p}) = -\sum_j p_j \ln p_j, \quad \bar{H} = \frac{1}{L} \sum_{i=1}^L H_i. \quad (1)$$

The normalising constant admits two readings, and both are needed below:

$$H_{FSS}^{(A)} = \mathbb{E}_{\mathbf{n} \sim \text{Mult}(N, \hat{\mathbf{p}})} [h(\mathbf{n}/N)], \quad H_{FSS}^{(B)} = h(\hat{\mathbf{p}}), \quad (2)$$

convention (A) being Xia’s [Xia et al., 2003, Xia and Lemey, 2009] and convention (B) the plug-in composition entropy. Since $H_{FSS}^{(A)} = h(\hat{\mathbf{p}}) - \frac{K-1}{2N} + O(N^{-2})$, the two are related on any data set by

$$I_{ss}^{(A)} = c I_{ss}^{(B)}, \quad c = \frac{h(\hat{\mathbf{p}})}{H_{FSS}^{(A)}}, \quad c \in (1, 1.08] \text{ for } N \geq 16. \quad (3)$$

Because c depends on $\hat{\mathbf{p}}$ and not on the tree, it is a constant within any one alignment: the convention rescales the index but does not change whether the affine relation holds (Section 2.4), nor which of two alignments is more saturated when they share a composition. Across data sets with different base frequencies c differs, so the two conventions need not even agree on the ordering; the ratio moves by at most 8% over the nucleotide range for $N \geq 16$, but by 51% at $N = 4$, where $c = 1/0.6620$.

2.2 Spectral resolution

Throughout, $\Pi = \text{diag}(\boldsymbol{\pi})$ and

$$S = \Pi^{1/2} Q \Pi^{-1/2} = U \Lambda U^\top, \quad \Lambda = \text{diag}(\lambda_1, \dots, \lambda_K), \quad 0 = \lambda_1 > \lambda_2 \geq \dots \geq \lambda_K. \quad (4)$$

Because Q is reversible, S is symmetric, so its eigenvectors are orthonormal and the spectral projectors are orthogonal:

$$P_\mu = \sum_{k: \lambda_k = \mu} u_k u_k^\top, \quad \tau_\mu = \text{tr } P_\mu \in \{1, \dots, K-1\}, \quad w_\mu = \text{tr}(\Pi P_\mu) = \sum_j \pi_j (P_\mu)_{jj}. \quad (5)$$

Lemma 1 (Spectral resolution). $Q = \Pi^{-1/2} U \Lambda U^\top \Pi^{1/2}$ and $P(t) = e^{Qt}$, hence

$$P_{ij}(t) = \pi_j + \sum_{k \geq 2} A_{ij}^{(k)} e^{\lambda_k t}, \quad A_{ij}^{(k)} = \pi_i^{-1/2} u_{ki} u_{kj} \pi_j^{1/2}, \quad A_{jj}^{(k)} = u_{kj}^2. \quad (6)$$

Proof. $P(t) = \Pi^{-1/2} U e^{\Lambda t} U^\top \Pi^{1/2}$ gives the entries directly. For $k = 1$, $\lambda_1 = 0$ and $\mathbf{u}_1 = \sqrt{\boldsymbol{\pi}}$ — indeed $S\sqrt{\boldsymbol{\pi}} = \Pi^{1/2} Q \mathbf{1} = \mathbf{0}$ — so $A_{ij}^{(1)} = \pi_i^{-1/2} \sqrt{\pi_i} \sqrt{\pi_j} \pi_j^{1/2} = \pi_j$. $\square$

Because S is symmetric its eigenvectors are orthonormal, so $\tau_k = \sum_j u_{kj}^2 = 1$ for every k : the grouped coefficient τ_μ in Eq. (5) is the algebraic multiplicity, with no separate argument needed. In particular $\text{tr } P(t) - 1 = \sum_{\mu \neq 0} \tau_\mu e^{\mu t}$, so the phylogenetic term $\Psi = \langle \text{tr } P \rangle - 1$ weights each eigenvalue by how many times it occurs. This is the reason the criterion below is a statement about *multiplicities* rather than about individual eigenvalues.

Lemma 2 (The three scalars). *With $\langle \cdot \rangle$ averaging over ordered leaf pairs and rate classes, and $\xi_\mu = \langle e^{\mu d} \rangle$,*

$$\Psi = \sum_{\mu \neq 0} \tau_\mu \xi_\mu, \quad \wp = \left(1 - \sum_j \pi_j^2\right) - \sum_{\mu \neq 0} w_\mu \xi_\mu, \quad \Delta = \sum_{\mu \neq 0} \left(w_\mu - \frac{\tau_\mu}{K}\right) \xi_\mu, \quad (7)$$

where $\wp$ is the mean pairwise p -distance and $\Delta = \sum_j (\pi_j - 1/K) (\langle P_{jj} \rangle - \pi_j)$ is the composition-leakage term of Theorem 13.

Proof. By Lemma 1, $\langle P_{jj} \rangle = \pi_j + \sum_{\mu \neq 0} (P_\mu)_{jj} \xi_\mu$. Summing over j gives $\Psi = \sum_j \langle P_{jj} \rangle - 1 = \sum_{\mu \neq 0} \tau_\mu \xi_\mu$; weighting by π_j gives $\sum_j \pi_j \langle P_{jj} \rangle = \sum_j \pi_j^2 + \sum_{\mu \neq 0} w_\mu \xi_\mu$, and 1 minus that is $\wp$. Finally $\langle P_{jj} \rangle - \pi_j = \sum_{\mu \neq 0} (P_\mu)_{jj} \xi_\mu$, so $\Delta = \sum_{\mu \neq 0} (w_\mu - \tau_\mu/K) \xi_\mu$ on substituting $\sum_j (P_\mu)_{jj} = \tau_\mu$ and $\sum_j \pi_j (P_\mu)_{jj} = w_\mu$. $\square$

Everything the index can see about the tree enters through the single vector $\boldsymbol{\xi}$; what distinguishes the three scalars is only how the spectrum is weighted — unweighted for Ψ , composition-weighted for $\wp$ and Δ . Section 2.4 shows that the index collapses onto $\wp$ exactly when those two weightings agree up to scale.

2.3 Exact and second-order closed forms

The index has two closed forms, one asymptotic and one exact on star trees.

100 **Lemma 3** (Entropy expansion). For $\hat{\mathbf{p}}$ with $\mathbb{E}[\hat{\mathbf{p}}] = \boldsymbol{\pi}$, lying in a compact neighbourhood of $\boldsymbol{\pi}$ on
 101 which $\hat{p}_j \geq \delta > 0$ for every j ,

$$\mathbb{E}[h(\hat{\mathbf{p}})] = h(\boldsymbol{\pi}) - \frac{1}{2} \sum_j \frac{\text{Var}(\hat{p}_j)}{\pi_j} + R_3, \quad |R_3| \leq \frac{1}{6\delta^2} \sum_j \mathbb{E}[|\hat{p}_j - \pi_j|^3]. \quad (8)$$

102 The bound is governed by the smallest frequency rather than by the divergence: it is small whenever
 103 every state occupies a non-negligible fraction of the sites, and degrades only as some $\hat{p}_j$ approaches
 104 zero.

105 *Proof.* $\partial h / \partial p_j = -(\ln p_j + 1)$, $\partial^2 h / \partial p_j^2 = -1/p_j$, $\partial^3 h / \partial p_j^3 = 1/p_j^2$, and the mixed partials vanish.
 106 Taylor expansion at $\boldsymbol{\pi}$ with Lagrange remainder; the linear term vanishes because $\mathbb{E}[\hat{p}_j] = \pi_j$. $\square$

107 **Lemma 4** (Variance aggregation). $\text{Var}(\hat{p}_j) = \frac{\pi_j(1 - \pi_j)}{N} + \frac{N-1}{N} \pi_j (\langle P_{jj} \rangle - \pi_j)$, hence $\sum_j \frac{\text{Var}(\hat{p}_j)}{\pi_j} =$
 108 $\frac{K-1}{N} + \frac{N-1}{N} \Psi$.

109 *Proof.* $\Pr(X_i = j, X_{i'} = j) = \pi_j P_{jj}(d_{ii'})$ for leaves $i \neq i'$ by reversibility and Chapman–Kolmogorov.
 110 Summing the $N(N-1)$ ordered covariances $\pi_j P_{jj}(d_{ii'}) - \pi_j^2$ and dividing by N^2 gives the first display;
 111 dividing by π_j and summing gives $\sum_j (1 - \pi_j)/N = (K-1)/N$ and $\frac{N-1}{N} \sum_j (\langle P_{jj} \rangle - \pi_j) = \frac{N-1}{N} \Psi$. $\square$

112 **Theorem 5** (Second-order closed form). Under the assumptions above, in convention (B), whose
 113 normalising entropy is the composition term $h(\boldsymbol{\pi})$ alone,

$$\mathbb{E}[I_{ss}^{(B)}] = 1 - \frac{1}{2h(\boldsymbol{\pi})} \left[\frac{K-1}{N} + \frac{N-1}{N} \langle \text{tr } P(rd) - 1 \rangle \right] + O(3). \quad (9)$$

114 Under convention (A) the normaliser carries its own finite-sample deficit of the same order, $H_{FSS}^{(A)} =$
 115 $h(\boldsymbol{\pi}) - \frac{K-1}{2N} + O(N^{-2})$, so the two leading terms of the bracket cancel and

$$\mathbb{E}[I_{ss}^{(A)}] = 1 - \frac{N-1}{2Nh(\boldsymbol{\pi})} \langle \text{tr } P(rd) - 1 \rangle + O(3) + O(N^{-2}). \quad (10)$$

116 Under JC69 in particular, Eq. (9) reads $\mathbb{E}[I_{ss}] = 1 - \frac{1}{2\ln 4} \left[\frac{3}{N} + \frac{N-1}{N} 3 \mathbb{E}_r e^{-8rd/3} \right] + O(3)$. The two
 117 are not interchangeable and the difference is not cosmetic: the coefficient of the phylogenetic term is
 118 $(N-1)/2Nh(\boldsymbol{\pi})$ under (A) against $(N-1)/2h(\boldsymbol{\pi})$ under (B), which is 7.8% at $N = 16$ and 51% at
 119 $N = 4$. Everything below that varies a parameter at fixed model — the coefficient $(N-1)/(2Nh(\boldsymbol{\pi}))$
 120 of Section 2.4, the sensitivity Eq. (17), the collapse section — is convention (A), which is the
 convention the observed values are reported in and the one DAMBE effectively uses. Equation (9) is
 (B), and is meant to be read as the expansion of the mean before normalisation.

121 *Proof.* Sites are independent, so $\mathbb{E}[\bar{H}] = \mathbb{E}[h(\hat{\mathbf{p}})]$ for one site. Apply Lemma 3 and substitute
 122 Lemma 4; divide by H_{FSS} . For JC69, $h(\boldsymbol{\pi}) = \ln 4$ and $\Psi = 3\mathbb{E}_r e^{-8rd/3}$ by Lemma 1 with $t = 2rd$.
 123 The (A) form follows on expanding the ratio in $(K-1)/2N$. $\square$

124 The remainder is small whenever every state occurs in a non-negligible fraction of sites, which is
 125 the saturated regime the index is used in, and degrades at low divergence or extreme composition.
 126 The expansion is also unnecessary on a star tree, where the leaves are conditionally i.i.d. given the
 127 root.

128 **Lemma 6** (Exact multinomial entropy). For $\mathbf{n} \sim \text{Mult}(N, \mathbf{q})$,

$$\mathbb{E}[h(\mathbf{n}/N)] = \ln N - \frac{1}{N} \sum_j \sum_{m=1}^N m \ln m \binom{N}{m} q_j^m (1 - q_j)^{N-m}. \quad (11)$$

129 *Proof.* $h(\mathbf{n}/N) = \ln N - \frac{1}{N} \sum_j n_j \ln n_j$ since $\sum_j n_j = N$. Take expectations with $\Pr(n_j = m) =$
 130 $\binom{N}{m} q_j^m q_j^{N-m}$, the $m = 0$ term contributing 0. $\square$

131 **Theorem 7** (Exact closed form on a star tree). For a star with N leaves, root-to-leaf length d , rate
 132 distribution $\{(r_c, w_c)\}$ and any reversible Q with stationary distribution $\boldsymbol{\pi}$,

$$\mathbb{E}[I_{ss}] = \frac{1}{h(\boldsymbol{\pi})} \sum_c w_c \sum_\rho \pi_\rho \left[\ln N - \frac{1}{N} \sum_j \sum_{m=1}^N m \ln m \binom{N}{m} q_{\rho j}^m (1 - q_{\rho j})^{N-m} \right], \quad (12)$$

133 where $q_{\rho j} = P_{\rho j}(r_c d)$, the (ρ, j) entry of the transition matrix — a finite sum of K exponentials by
 134 Lemma 1. Under JC69 this reads $q_{\rho\rho} = \frac{1}{4} + \frac{3}{4}e_c$ and $q_{\rho j} = \frac{1-e_c}{4}$ for $j \neq \rho$, with $e_c = e^{-4r_c d/3}$; under
 135 rate homogeneity the outer sum degenerates to $r = 1$.

136 *Proof.* Given the root state ρ and rate class c the N leaf states are i.i.d. with law $\mathbf{q}_\rho = P_\rho(r_c d)$,
 137 which by Lemma 1 has the stated entries. Apply Lemma 6, average over ρ and c , divide by $h(\boldsymbol{\pi})$.
 138 Nothing in the argument uses JC69; the specialisation enters only through the entries of $\mathbf{q}_\rho$. $\square$

139 The theorem is stated for JC69 in the literature this paper is about, but its proof never asks for it,
 140 and the general statement is worth having because it answers what the index can be written as when
 141 the model is not degenerate. Take a REV model: six exchangeabilities $t_{AC}, t_{AG}, t_{AT}, t_{CG}, t_{CT}, t_{GT}$
 142 and four frequencies $\boldsymbol{\pi}$, with $Q_{ij} = t_{ij}\pi_j$ and $c = \sum_{i < j} 2\pi_i\pi_j t_{ij}$ so that branch lengths are in
 143 expected substitutions. The parameters enter Eq. (12) only through $S = \Pi^{1/2} Q \Pi^{-1/2}$, which for a
 144 REV model is the *explicit* symmetric matrix

$$S_{ij} = \frac{\sqrt{\pi_i \pi_j} t_{ij}}{c} \quad (i \neq j), \quad S_{ii} = -\frac{1}{c} \sum_{j \neq i} \pi_j t_{ij}, \quad (13)$$

145 with $S\sqrt{\boldsymbol{\pi}} = \mathbf{0}$ and therefore $S = U\Lambda U^\top$, U orthonormal, $\mathbf{u}_1 = \sqrt{\boldsymbol{\pi}}$. Nothing else is needed: the
 146 three non-zero eigenvalues are read off S , and $q_{\rho j} = \pi_j + \sum_{k \geq 2} \pi_\rho^{-1/2} u_{k\rho} u_{kj} \pi_j^{1/2} e^{\lambda_k r_c d}$. So Eq. (12)
 147 is the closed form for the general reversible model, with no matrix exponential and no iteration;
 148 the archived implementation evaluates it against the pattern enumeration and agrees to machine
 149 precision under both conventions.

150 What it is not is a formula in $\bar{p}$, and that is the whole point of Section 2.4. The eight free
 151 parameters survive into S , hence into the eigenvalues and eigenvectors, and reach $\mathbb{E}[I_{ss}]$ through the
 152 full vector $\boldsymbol{\xi}$; the reduction to a function of the mean p -distance requires $\tau_{[\mu]}/w_{[\mu]}$ to be constant
 153 across the distinct non-zero eigenvalues. Uniform composition delivers that for *every* exchangeability
 154 vector, by Corollary 14, and a degenerate spectrum delivers it for every composition, by Corollary 15;
 155 a generic REV model is in neither case. For such a model Eq. (12) is therefore not a convenient
 156 way of writing a function of one variable but the shortest exact expression available — a sum whose
 157 every term is a distinct product of exponentials.

Corollary 8 (The ceiling is a function of N in one convention only). *In convention (B), where Eq. (9) holds as written, $\mathbb{E}[I_{ss}^{(B)}] \leq 1 - \frac{K-1}{2Nh(\boldsymbol{\pi})}$, since $\Psi \geq 0$. In the saturation limit every $\xi_\mu \rightarrow 0$, so $\mathbb{E}[\bar{H}] \rightarrow \mathbb{E}_{\text{Mult}(N, \boldsymbol{\pi})}[h]$ and*

$$\lim_{d \rightarrow \infty} \mathbb{E}[I_{ss}^{(A)}] = 1 \quad \text{for every } N, \quad \lim_{d \rightarrow \infty} \mathbb{E}[I_{ss}^{(B)}] = \frac{\mathbb{E}_{\text{Mult}(N, \boldsymbol{\pi})}[h]}{h(\boldsymbol{\pi})}, \quad (14)$$

the latter equal to $1 - \frac{K-1}{2N \ln K} + O(N^{-2})$.

Proof. $H_{FSS}^{(A)} = \mathbb{E}_{\text{Mult}(N, \hat{\mathbf{p}})}[h]$ and $\hat{\mathbf{p}} \rightarrow \boldsymbol{\pi}$ by the law of large numbers, so numerator and denominator coincide in the limit. For (B), $H_{FSS}^{(B)} = h(\boldsymbol{\pi})$ does not depend on N while the numerator retains the finite-sample deficit. The expansion is the Miller–Basharin term. $\square$

The bound is second order, and its slack cannot account for the excursions above 1 that the literature reports (1.49): those exceed any $O(3)$ or $O(N^{-2})$ term by two orders of magnitude, so they are evidence of state-space mismatch rather than of a bound that is merely loose. A non-asymptotic version would compare $\mathbb{E}[h]$ under the exchangeable, positively correlated site law with the multinomial law in convex order, which is a stronger statement than the argument uses.

Numerically, the exact ceiling under (B) is 0.661995 ($N = 4$), 0.840643 ($N = 8$) and 0.927377 ($N = 16$). The $O(N^{-2})$ bookkeeping matters: at $N = 16$ the exact multinomial entropy at uniform composition is 1.285617 nats against the second-order $\ln 4 - 3/32 = 1.292544$, and the normalised ceilings are 0.927377 against 0.932374. The discrepancy is $O(N^{-2})$, not $O(3)$, but at $N = 4$ it is large — exact 0.6620, second order 0.7295.

Corollary 8 is the first place the two conventions separate. It is not confined to the limit: at $\bar{p} = 0.599$, where Eq. (9) applies with a remainder under 10^{-3} , the exact star values are 0.7770, 0.7810 and 0.7983 under (A) for $N = 4, 8, 16$, against 0.5144, 0.6566 and 0.7403 under (B). Between $N = 4$ and $N = 16$ convention (A) moves by 2.7% of its value and convention (B) by 44%. Xia’s normalisation is not merely a rescaling: it removes the leading finite-sample dependence on the number of taxa, which the plug-in entropy does not. The simulation reproduces both families to within 0.0006.

Remark 9 (Second-order accuracy varies strongly with N). At $N = 4$ — the size used for the simulations below — the second-order affine line deviates from the exact expectation by 0.02–0.07, i.e. by 3–11% of the attainable range of the index, and the exact chord slope over a representative interval is 0.84 against the predicted 1.082. The affine identity is qualitatively robust and quantitatively poor at small N ; all numerical claims about it should be made at $N \geq 16$ or computed exactly.

Rate heterogeneity

Lemma 10 (Γ transform and monotonicity). *For $r \sim \Gamma(\alpha, \text{scale} = 1/\alpha)$ and $\lambda < 0$,*

$$\mathbb{E}_r[e^{\lambda r t}] = \left(1 - \frac{\lambda t}{\alpha}\right)^{-\alpha}, \quad (15)$$

and $\alpha \mapsto \mathbb{E}_r[e^{\lambda r t}]$ is strictly decreasing for every $t > 0$, with limit $e^{\lambda t}$ as $\alpha \rightarrow \infty$. Consequently $\Psi_\alpha = 3\left(1 + \frac{8d}{3\alpha}\right)^{-\alpha}$ on a star under JC69, strictly decreasing in α , so $\mathbb{E}[I_{ss}]$ is strictly increasing in α .

Proof. The moment generating function of $\Gamma(\alpha, \alpha)$ is $(\alpha/(\alpha - s))^\alpha$ for $s < \alpha$; set $s = \lambda t < 0$. Writing $g(\alpha) = -\alpha \ln(1 + u)$ with $u = -\lambda t/\alpha > 0$, we get $g'(\alpha) = -\ln(1 + u) + u/(1 + u) < 0$ because $\ln(1 + u) > u/(1 + u)$ for all $u > 0$; and $-\alpha \ln(1 - \lambda t/\alpha) \rightarrow \lambda t$ as $\alpha \rightarrow \infty$. Since $\mathbb{E}[I_{ss}]$ is decreasing in Ψ by Eq. (9), the last claim follows. $\square$

The homogeneous case decays as $e^{-8d/3}$, exponentially; rate heterogeneity gives a *power law* $\Psi_\alpha \sim 3(3\alpha/8d)^\alpha$ as $d \rightarrow \infty$. A power-law tail is heavier than any exponential, so heterogeneity retards the approach of $\mathbb{E}[I_{ss}]$ to its ceiling at a rate set by α , not by any small perturbation.

Proposition 11 (Non-identifiability of the critical value). *Let $I_{ss,c} = I_{ss}(TL_c)$ with TL_c the tree length at which $\Pr[\tilde{\mathcal{T}} = \mathcal{T}]$ falls to 0.95. Then $I_{ss,c}$ is a strictly monotone function of α , and the tabulated values, computed under rate homogeneity, correspond to $\alpha = \infty$ alone.*

Proof. By Lemma 10 the map $\alpha \mapsto I_{ss}(TL; \alpha)$ is strictly increasing, with range

$$[I_{ss}(TL; 0^+), I_{ss}(TL; \infty)].$$

The level set $\{TL : I_{ss}(TL; \alpha) = \theta\}$ — and hence both TL_c and $I_{ss}(TL_c)$ — therefore depends on α . $\square$

The dependence is not, by itself, fatal, and Section 2.6 shows that it is weaker than it sounds: over the range of α that matters, TL_c moves by a factor of 15 while $I_{ss,c}$ moves by less than 0.01. The charge is therefore not that the threshold is unusable but that it is mis-calibrated in a way its user cannot detect: the critical tree length does depend strongly on α , the index does not carry that dependence, and a table with no α coordinate silently maps one onto the other.

2.4 When the index collapses onto p -distance

Definition 12 (Scaling family). Fix a topology with branch-length vector $\mathbf{d}^0$ and a rate distribution, and scale $\mathbf{d}(s) = s\mathbf{d}^0$. The family is *exponentially non-degenerate* if the functions $\{s \mapsto \xi_\mu(s)\}$, one per distinct non-zero eigenvalue, together with the constant function, are linearly independent on the family.

This holds generically. On a star under Γ -rates, $\xi_\mu(s) = (1 - 2\mu s d^0 / \alpha)^{-\alpha}$, and functions with distinct μ have distinct singularities in the complex plane.

Theorem 13 (Criterion). *For an exponentially non-degenerate scaling family the following are equivalent:*

1. Ψ is an affine function of $\wp$ along the family;
2. Δ is an affine function of Ψ along the family;
3. there is a constant B with $\tau_{[\mu]} + B w_{[\mu]} = 0$ for every distinct non-zero eigenvalue μ .

When (3) holds, $\Psi = B\wp - B(1 - \sum_j \pi_j^2)$.

Proof. (1) $\Rightarrow$ (3): suppose $\Psi = A + B\wp$. By Lemma 2, $\sum_\mu \tau_\mu \xi_\mu = A + B(1 - \sum_j \pi_j^2) - B \sum_\mu w_\mu \xi_\mu$, i.e. $\sum_\mu (\tau_\mu + B w_\mu) \xi_\mu$ is constant. Exponential non-degeneracy forces each coefficient to vanish, and summing within a distinct eigenvalue gives (3). (3) $\Rightarrow$ (1) is the same computation read backwards. (2) $\Leftrightarrow$ (3): Δ and Ψ are affinely related iff $w_\mu - \tau_\mu / K = c \tau_\mu$ for a constant c , which is (3) with $B = -(c + 1/K)^{-1}$. $\square$

The failure condition is therefore $\Delta \not\propto \Psi$, *not* $\Delta \neq 0$. A nonzero composition term is harmless as long as it is proportional to the phylogenetic one; this distinction is the difference between a model that satisfies the criterion and one that does not, as the F81 and HKY rows below show. To the accuracy of Eq. (9) the deviation of $\mathbb{E}[I_{ss}]$ from the p -distance prediction is $-\frac{N-1}{2Nh(\pi)}\Delta + O(3)$, so Δ measures the size of the failure and Theorem 13 measures its shape. Both statements are second-order; beyond it $\mathbb{E}[I_{ss}]$ depends on third and fourth moments of the site-pattern distribution, which pairwise identity does not determine, and Remark 18 below states what follows from that.

model	π	# distinct μ	λ_k	$\tau_{[\mu]}$	$w_{[\mu]}$	affine?
JC69	uniform	1	-1.3333×3	[3]	[0.75]	yes
F81	(.4, .3, .2, .1)	1	-1.4286×3	[3]	[0.70]	yes
HKY $\kappa=5$	uniform	2	$-0.5714, -1.7143 \times 2$	[1, 2]	[0.25, 0.50]	yes
HKY $\kappa=5$	(.4, .3, .2, .1)	3	$-0.6329, -1.6456, -2.1519$	[1, 1, 1]	[.2833, .1500, .2667]	no
GTR	uniform	3	$-0.6224, -1.2568, -2.1209$	[1, 1, 1]	[.25, .25, .25]	yes
GTR	(.4, .3, .2, .1)	3	$-0.5418, -1.1829, -2.1807$	[1, 1, 1]	[.1228, .3260, .2512]	no

Table 1: The criterion of Theorem 13 under six models. The ratios $\tau_{[\mu]}/w_{[\mu]}$ are constant exactly in the rows marked *yes*. Two rows carry the content: GTR with three distinct eigenvalues and uniform composition is affine, confirming Corollary 14 where the spectrum is fully non-degenerate; and F81 with strongly skewed composition is affine, confirming Corollary 15.

Corollary 14 (Uniform composition). *If π is uniform then $\Pi = I/K$, hence $w_\mu = \tau_\mu/K$ for every μ and $\tau_\mu/w_\mu = K$ identically — for every reversible Q , every topology, every branch-length vector and every rate distribution. The proof is one line and never mentions whether the spectrum is degenerate.*

Corollary 15 (Degenerate spectrum). *If Q has a single distinct non-zero eigenvalue μ — JC69, F81, and any model whose generator is a function of a single rank-one matrix — the identity holds for every composition, with $B = -\tau_{[\mu]}/w_{[\mu]}$.*

Proof. There is only one grouped pair to satisfy, so (3) is solved by any B . For F81, $P(t) = e^{-t}I + (1 - e^{-t})\mathbf{1}\pi^\top$, so $P_{jj}(t) = \pi_j + (1 - \pi_j)e^{-t}$, $\mu = -1$, $\tau_{[\mu]} = K - 1$, $w_{[\mu]} = \sum_j \pi_j(1 - \pi_j) = 1 - \sum_j \pi_j^2$, and $\wp = (1 - \sum_j \pi_j^2)(1 - \langle e^{-d} \rangle)$. $\square$

Corollary 16 (Generic failure). *If π is non-uniform and Q has at least two distinct non-zero eigenvalues with $\tau_{[\mu]}/w_{[\mu]}$ varying across μ , then $\mathbb{E}[I_{ss}]$ is a genuinely higher-dimensional function of ξ and cannot be recovered from $\wp$.*

Xia et al. [Xia et al., 2003] calibrate under F84 with $\kappa \in [1, 10]$ and base frequencies in $[0.1, 0.9]$ — with non-uniform composition and a non-degenerate spectrum, this is squarely inside the failure region. Table 1 evaluates the ratio for the models at issue.

The verdict is a property of the model alone, which is the first thing a reader will try to falsify and is worth settling before it is. $\tau_{[\mu]}$ and $w_{[\mu]}$ are read off π and Q , so the criterion cannot depend on the tree; the only place a topology enters is the vector ξ , and it enters Ψ and $\wp$ the same way. We checked the residual on six topologies — an equidistant star, quartets with internal fractions 0.10, 0.35 and 0.50, and balanced trees of 8 and 16 taxa — under three models. Where the criterion holds the relation is affine to 2×10^{-15} on all six; where it fails, the largest departure from a straight line is 2.3×10^{-2} to 3.7×10^{-2} in Ψ units, a spread of $\pm 20\%$ over a range of tree shapes from four taxa to sixteen. The failure is a property of the substitution model, and the tree changes its size by an amount that does not change what it means.

Lemma 17 (Shape of the failure). $\Delta(0) = 1/K - \sum_j \pi_j^2$, $\Delta(\infty) = 0$, and Δ need not be monotone.

Proof. At $\xi_\mu \equiv 1$, $\Delta = \sum_\mu (w_\mu - \tau_\mu/K) = (1 - \sum_j \pi_j^2) - (K - 1)/K = 1/K - \sum_j \pi_j^2$; at $\xi_\mu \rightarrow 0$, $\Delta \rightarrow 0$. Non-monotonicity: Δ is a signed mixture of exponentials with distinct rates, and the coefficients $w_\mu - \tau_\mu/K$ change sign across μ .

Numerically, for HKY with $\kappa = 5$ and $\pi = (.4, .3, .2, .1)$ the ratio Δ/Ψ sweeps from -0.0155 to $+0.0333$ as d runs over $[0.05, 8]$, so the identity fails; for F81 with the same composition it is




Figure 1: Theorem 13 evaluated at the two levels it distinguishes. (A) Ψ against the mean p -distance, computed exactly from the spectral decomposition on a star ($d \in [0.02, 1.5]$; dashed lines are the best fits). The two models whose criterion holds lie on their fits to 10^{-15} ; HKY with $\kappa = 5$ and $\pi = (.4, .3, .2, .1)$ does not. (B) The same departures multiplied by $(N - 1)/2Nh(\pi)$ at $N = 16$, i.e. expressed in the units of the index itself. The criterion-failing model is not distinguished at this level: the $O(3)$ remainder of Eq. (9) exceeds the affinity failure by a factor of 18 to 27 depending on the model, which is why the criterion must be checked on Ψ rather than on $\mathbb{E}[I_{ss}]$.

constant at $w/\tau - 1/K = 0.70/3 - 0.25 = -0.016667$ to all printed digits, so the identity holds even though $\Delta \neq 0$. The HKY extremum falls near $d \approx 1$, in the mid-divergence region where the critical values are calibrated. $\square$

Theorem 13 constrains Ψ and $\wp$, and the choice of those two scalars rather than $\mathbb{E}[I_{ss}]$ is not cosmetic. Ψ and $\wp$ are exact linear functionals of the spectral coordinates, whereas $\mathbb{E}[I_{ss}]$ is not: by Eq. (9) it carries the composition term Δ plus an $O(3)$ remainder, and a test performed on it measures those as much as it measures the criterion.¹ Figure 1 separates the two levels. Panel A evaluates Ψ against $\wp$ from the spectral decomposition alone, on a star of root-to-leaf length d scanned over $d \in [0.02, 1.5]$, for models exercising all three corollaries. Where the criterion holds, the relation is affine to machine precision: the largest departure from the best-fit line is 1.8×10^{-15} for HKY with uniform composition and 2.2×10^{-15} for F81 with $\pi = (.4, .3, .2, .1)$, which is double-precision residue rather than residual structure. Where it fails — HKY with $\kappa = 5$ and the same skewed composition, the model of the last row of Table 1 — the departure reaches 1.3×10^{-2} in Ψ units, or 4.9×10^{-3} in the units of I_{ss} at $N = 16$. Panel B then evaluates the same three models in $\mathbb{E}[I_{ss}]$ instead, and there the criterion-holding and criterion-failing models are indistinguishable: all three depart from a straight line by 10–15% of the range of $\mathbb{E}[I_{ss}]$, and the failing model sits in the middle of that pack rather than at its edge. The $O(3)$ channel moves the index by some twenty times what the criterion’s failure moves it. The practical consequence is that the failure cannot be diagnosed by plotting I_{ss} against $\bar{p}$, because at nucleotide skew it is buried under the higher-order terms; the diagnostic is the ratio $\tau_{[\mu]}/w_{[\mu]}$, which is read off the model and requires no data.

¹An earlier version of this check did exactly that, and returned its *largest* residual for the model whose criterion holds — the signature of a test aimed at the wrong quantity. It is recorded here because the mistake is easy to repeat.

Remark 18 (Scope of the collapse). The one-line argument fails exactly when the Π -average on some non-trivial eigenspace departs from the global mean $1/K$. Three qualifications must be stated, and the third is the strongest. (i) For nucleotides the departure is small but not zero: with non-uniform composition and a non-degenerate spectrum the deviation of I_{ss} from the p -distance prediction is 0.014 (least-squares) to 0.022 (chord), which is 4% to 44% of the separation between the symmetrical and asymmetrical critical values, depending on the number of taxa — the figure is *not* a single number. (ii) The criterion depends on the stationary composition of the *model*; a uniform-frequency model in any alphabet satisfies it. In practice nucleotide compositions are close to uniform while amino-acid compositions are not, so the collapse is a nucleotide-scale result. (iii) The decomposition above is *second order*, and it would be wrong to read it as a general statement that everything I_{ss} adds beyond p -distance is composition. At second order $\mathbb{E}[I_{ss}]$ depends on $\text{Var}(\hat{p}_j)$ alone, hence only on pairwise identity $\langle P_{jj} \rangle$; beyond second order it depends on third and fourth moments of the site-pattern distribution, which pairwise identity does *not* determine. The site patterns carry structure that the mean mismatch proportion cannot see — for $K = 4$ a 2:2 site is parsimony-informative while a 3:1 site is not, and the two have different entropies. The magnitude of this higher-order channel is measurable and is reported in Section 2.13; whether it is phylogenetically informative rather than merely compositional is a separate question, answered by testing conditional independence:

$$\Pr(Y \mid \bar{p}, I_{ss}) \quad \text{versus} \quad \Pr(Y \mid \bar{p}), \quad Y \in \{\text{parsimony-informative fraction, recovery}\}. \quad (16)$$

That test, and not the correlation between the two statistics, is what adjudicates redundancy.

2.5 The critical value depends on sequence length

Theorem 19 (Length invariance and its price). *With $\bar{H} = \frac{1}{L} \sum_i H_i$ and $B = h(\hat{\mathbf{p}})$,*

1. $\mathbb{E}[\bar{H}] = \mu(s)$ *exactly, for every L ;*

2. $\text{Var}(\bar{H}) = \sigma^2(s)/L$ *with $\sigma^2 = \text{Var}(H_1)$;*

3. $\text{Var}(I_{ss}) = \Theta(L^{-1})$;

4. $\mathbb{E}[I_{ss}] = I_{ss}^\infty + c/L + O(L^{-2})$, *with $I_{ss}^\infty = \mu/h(\boldsymbol{\pi})$ and $c = \frac{\mu \text{Var}(B)}{h(\boldsymbol{\pi})^3} - \frac{\text{Cov}(\bar{H}, B)}{h(\boldsymbol{\pi})^2}$.*

Proof. (1)–(2) Sites are i.i.d., so $\bar{H}$ is an i.i.d. mean. (3)–(4): expand B^{-1} about $h(\boldsymbol{\pi})$ to second order,

$$\mathbb{E}[\bar{H}B^{-1}] = \frac{\mu}{\mu_B} - \frac{\text{Cov}(\bar{H}, B)}{\mu_B^2} + \frac{\mu \text{Var}(B)}{\mu_B^3} + O(L^{-2}),$$

and

$$\text{Var}(I_{ss}) = \left(\frac{\mu}{\mu_B} \right)^2 \left[\frac{\text{Var}(\bar{H})}{\mu^2} - \frac{2 \text{Cov}(\bar{H}, B)}{\mu \mu_B} + \frac{\text{Var}(B)}{\mu_B^2} \right];$$

all three terms are $\Theta(L^{-1})$ because $\hat{\mathbf{p}}$ is pooled over NL observations while $\bar{H}$ fluctuates at $L^{-1/2}$. $\square$

Theorem 19 is the analytic content of an empirical fact the original paper already reports. Figure 3 of Xia et al. [2003] is titled *The critical index of substitution saturation depends on the sequence length and the number of OTUs*, and the surrounding text uses that dependence as advice — a user with an asymmetrical topology and $I_{ss} = 0.73$ at $N = 8$ is told to exceed 500 bp if the tree is symmetrical and 3000 bp if it is strongly asymmetrical. Their Figure 2 likewise reports

that $\Pr[\hat{\mathcal{T}} = \mathcal{T}]$ depends on sequence length. The dependence is therefore not overlooked, and the threshold being L -indexed is a deliberate design decision rather than an error.

What Theorem 19 supplies is why the dependence has the form it does. I_{ss} is an average over sites and cannot depend on L except through finite-sample noise; H_{FSS} is pooled over NL observations and is $\sqrt{N}$ better determined than $\bar{H}$; so $\text{Var}(I_{ss}) = \Theta(L^{-1})$ exactly, and the critical value must move with L at rate $L^{-1/2}$ to hold a decision threshold fixed. That is the mechanism behind Figure 3, and it fixes the functional form of the tabulated surface rather than merely exhibiting it.

Corollary 20 (A fixed margin does not stay fixed). *Since $\text{SE} = \Theta(L^{-1/2})$, the statistic $T = (I_{ss,c} - I_{ss})/\text{SE}$ satisfies $T = \Theta(\sqrt{L})$ for any fixed margin. A user who reads a tabulated $I_{ss,c}$ for the wrong L , or who compares a margin across data sets of different length, is reading a number whose significance changes with the site count.*

Xia et al.’s own worked example makes the scale concrete: their $I_{ss} = 0.5573$, $I_{ss,c} = 0.8093$, $\text{SE} = 0.01830$ on $\text{DF} = 731$ give $T = 13.77$, and the same margin at ten and a hundred times the length gives $T = 43.55$ and $T = 137.70$. This is a property of any consistent test whose statistic has an $L^{-1/2}$ standard error — it is not peculiar to I_{ss} — but it does mean the threshold is only meaningful jointly with the length it was tabulated for.

The dependence is visible in published regressions without any theory at all. Duchêne et al. [2022] fit the entropy t -statistic against six simulation variables, the variable-site count and all two-way interactions, over roughly 3,100 simulations per coefficient. Sequence length is the only main effect with a substantial positive coefficient (+38.7, $t = 89.5$); every other variable enters negatively (tree length -14.6 , variable sites -29.0 , stemminess -5.7 , proportion of invariable sites -6.1 , imbalance -4.1 , number of taxa -1.1). A statistic whose value is driven up by more sites and down by everything else is the $T = \Theta(\sqrt{L})$ of Corollary 20 read off an empirical fit. It is also the sharper form of the observation: the tabulated threshold is not merely indexed by L , it is dominated by it.

Remark 21. Theorem 19(4) also shows the *mean* carrying an $O(L^{-1})$ bias that the median does not, since the median is unaffected by the monotone per-site structure. The median still fluctuates at $O(L^{-1/2})$, so this is not a claim of noise-freeness; it is that the $1/L$ term, which has no phylogenetic content, is absent from the median and present in the mean. That makes the two a clean contrast for detecting it.

2.6 Resolution collapse in the near-saturated region

The index is bounded above, and the critical value sits close to that bound. The two facts together do more than either alone.

Corollary 22 (Where the critical value sits). $\mathbb{E}[I_{ss}] \leq I_{ss}^{\max}$ with $I_{ss}^{\max} = \mathbb{E}_{\text{Mult}(N, \boldsymbol{\pi})}[h]/h(\boldsymbol{\pi})$ under convention (B) and $I_{ss}^{\max} = 1$ under (A), and the bound is approached as $\Psi \rightarrow 0$. The sensitivity of the index to any parameter that acts through Ψ is

$$\left| \frac{d \ln I_{ss}}{d \ln \Psi} \right| = \frac{N-1}{2Nh(\boldsymbol{\pi})} \cdot \frac{\Psi}{I_{ss}}, \quad (17)$$

which vanishes as $\Psi \rightarrow 0$: the closer the critical point lies to the ceiling, the more inert the index is to everything that produced it.

Proof. Both parts follow from Eq. (9) by differentiation; $\partial I_{ss}/\partial \Psi = -(N-1)/(2Nh(\boldsymbol{\pi}))$ is a constant, so the logarithmic derivative is that constant times Ψ/I_{ss} . $\square$

The corollary has a quantitative consequence that is easy to overlook, and because it carries the section it is worth setting out exactly how each number is obtained. Take $N = 4$, $L = 1000$, and a quartet with 35% of its length in the internal branch. In a dedicated experiment (Figure 2) the critical tree lengths at which $\Pr[\hat{\mathcal{T}} = \mathcal{T}]$ falls to 0.95 under maximum likelihood are

α	∞	32	8	4	2	1	0.5
TL_c	7.19	7.32	7.78	10.04	13.02	25.47	107.72
$I_{ss,c}$	0.9807	0.9788	0.9735	0.9782	0.9718	0.9692	0.9715
SE	0.0096	0.0096	0.0097	0.0096	0.0098	0.0099	0.0098

(18)

so the tree length runs over a factor of 15 end to end while the index spans 0.0115. Recovery is by IQ-TREE rather than by our own code, at 150 replicates per point; the crossings are bracketed by the grid and located by interpolation in $\log TL$, and the brackets are reported in the archived output. The last row is the exact standard error of a single reading at $L = 1000$, from Lemma 23. It is flat across the row, so the whole of the α effect that the two rows above describe amounts to 1.2 standard errors at this length, and the row is there so that a reader can see that without recomputing it. The sweep is denser than the table shows, and the density is the point: $\alpha = 64, 32, 24, 16, 12$ are all run in addition, and from $\alpha = \infty$ down to $\alpha = 8$ the critical tree length does not leave 7.1–7.8, a band one grid step wide. Rate heterogeneity has to reach $\alpha \approx 4$ before the threshold notices it at all.

The sweep stops at $\alpha = 0.5$ because below it the discretisation, not the model, is what the curves would be plotting. Equal-probability four-category quadrature puts a quarter of the sites in its slowest class, and that class falls to 2×10^{-3} of the mean rate at $\alpha = 0.25$ and to 5×10^{-7} at $\alpha = 0.1$ — a pile of effectively invariant sites, not the tail of a continuum, and a quadrature error no longer small next to the effect being measured. An earlier version of this experiment ran the ladder down to 0.1 and found the recovery curve turning over and rising again with tree length; that turnover is what the inadequacy looks like, and which we no longer report. Nothing in this section depends on the suppressed range: the point is made from 0.5 upwards, where four categories are adequate and the curves are monotone. The scope of the exclusion is worth stating, because the index itself does not need it: what fails below 0.5 is the monotone crossing this section takes as its estimand, not the index, which stays a well-defined function of an alignment simulated under the same discretisation — and which Section 2.12 reads down to $\alpha = 0.25$ for exactly that reason.

The $I_{ss,c}$ row is exact, not simulated and not taken from the second-order form. At $N = 4$ the transition matrix has only 4^4 site patterns, so Lemma 6 can be evaluated by exhaustive enumeration over the count vectors of an internal branch; those are the values of Theorem 7, extended from the star to a quartet by that enumeration. The convention matters and is stated: these are convention (A) values, where the ceiling is 1. Under (B) at $N = 4$ the exact ceiling is 0.6620 (Corollary 8) and the same points read 0.64–0.65; the collapse is a statement about the width of the window, not about its location, and it appears under either convention with the span little changed.

Because the index is so flat there, the critical points are not even ordered by α — 0.9735 at $\alpha = 8$ sits below 0.9782 at $\alpha = 4$ — and the small non-monotonicity is the grid and replicate noise of the crossing, not a property of the index. That is what a quantity with nothing left to move looks like.

The collapse also has an estimator-level reading, which is the form in which a user would meet it and which is worth stating because it looks like an opportunity. On an equal-length star the index is monotone in the tree length, so a reading of I_{ss} can be inverted to a divergence — once the model is fixed. Fixing the model is not a formality. At $N = 8$ under convention (A), a reading of 0.75 implies a root-to-leaf length of 0.548 under JC69, 0.626 under GTR with uniform composition and 0.616 under GTR with $\pi = (.31, .24, .19, .26)$; the three composition-varying models agree to within




Figure 2: The critical point against rate heterogeneity. **A**, recovery of the true quartet under maximum likelihood against tree length, one curve per Γ shape, $N = 4$, $L = 1000$, 150 replicates per point; the dashed line is 0.95. **B**, the tree length at which each curve crosses it, on axes that put $\alpha = \infty$ one doubling beyond the largest finite shape. **C**, the exact expected index at those tree lengths, on the same 0–1 scale as the index itself. A fifteen-fold change in the first has no visible effect on the second: the threshold is nearly invariant to rate heterogeneity because the index is nearly at its ceiling, not because it is well calibrated. The ladder stops at $\alpha = 0.5$; see the text for why.

15%, which is the small size of the leakage term established in Remark 18(i). The same reading under the skewed GTR with Γ rates of shape 0.5 implies 2.87 — a factor of 5, in the same units, for the same reading of the same statistic. At a reading of 0.85 the spread runs from 0.75 to 6.97. The model has to be supplied before the inversion can be attempted, and it is supplied by a method the index cannot check; the resulting inverse is ill-conditioned exactly where the index is used, and the conditioning is set by the rate heterogeneity.

This is why the linearity that does exist does not help. By Eq. (9) the index is an exact affine functional of the spectral moments $\xi_\mu = \langle e^{\mu d} \rangle$, so it is a moment condition and can be inverted — but it is *one* moment against a parameter vector of eight model parameters plus a tree, and it is a linear functional of ξ rather than of the parameters, since ξ itself is nonlinear in both the branch lengths and the shape. Beyond second order the affinity is gone as well: Eq. (12) is a polynomial of degree N in pattern probabilities that are themselves sums of exponentials of $\lambda_k(\pi, t_{ij})$. The index therefore cannot support an estimation scheme, and it does not need to: the likelihood of these models is available directly by pruning, and estimating π , t_{ij} and α has never required I_{ss} .

The span depends on the tree shape, as Eq. (17) implies it must, but it does not approach the width of a tabulated row in either shape we tried. Equation (17) evaluates to 0.005 or less throughout. That equation is a second-order statement (Remark 9), and it is quoted only for the scale; the conclusion it supports does not need it, since the denominator of Eq. (17) is bounded below by $1 - I_{ss}^{\max}$ while the derivative $\partial \mathbb{E}[I_{ss}] / \partial \Psi$ is bounded above by a constant, so the sensitivity to any parameter acting through Ψ vanishes at the ceiling whether or not the second-order form is accurate there.

This is the resolution collapse, and it is worth separating into two claims of different strength, because only the first is a property of the statistic.

The first is intrinsic. By Eq. (17) the index’s elasticity with respect to Ψ vanishes as $\Psi \rightarrow 0$, whatever produced the divergence, so I_{ss} is a strictly decreasing function of Ψ whose slope goes to zero at the point where it is read. A fifteen-fold change in the tree length moves it by 0.0115, which is 1.2% of its scale. No amount of data recovers a quantity that the statistic has stopped

responding to; that is a statement about what I_{ss} is a function of, and it is the one the rest of this paper rests on.

The second is a matter of power, and it is weaker. The exact variance of Lemma 23, evaluated on this quartet, gives $\text{SE}(I_{ss})$ between 0.0096 and 0.0099 across the α ladder — so the 0.0115 span is 1.2 standard errors, which is why it is invisible at this length. Since the standard error falls as $L^{-1/2}$, the same span reaches three standard errors at $L \approx 6300$, and ten at $L \approx 7 \times 10^4$. The insensitivity is therefore not that the index cannot be resolved, but that resolving it costs an alignment long enough that the tree could have been estimated directly. Any claim that two partitions differ in α carries an implied L , and at the lengths the index is ordinarily applied to the implied L is too long.

Two qualifications on the scope. First, the sensitivity in Eq. (17) is proportional to $\Psi_c/I_{ss,c}$, and Ψ_c depends on the tree shape: it is small for a quartet with a long internal branch and larger for one with a short internal branch, so the collapse is sharpest for the configurations where saturation matters most and weakest where the internal branch is short. Second, the argument concerns the *level* of the threshold, not its *ordering*: the critical value still moves monotonically with α (Lemma 10), just not by enough to be usable as a coordinate. The rate-heterogeneity mis-calibration established by Proposition 11 stands; what this section adds is that a user who tries to detect it by comparing tabulated thresholds will not see it.

2.7 The index as a moment of the count distribution

Everything so far treats I_{ss} as a function of the model. It is worth writing it as a function of the *data*, because that is the form in which it could be used for inference, and that form also shows exactly how much of the alignment it uses.

Write $X_1, \dots, X_L$ for the columns, each in $\{A, C, G, T\}^N$, let $\hat{f}_x$ be the observed frequency of pattern x , and let $c(x) \in \mathbb{Z}_{\geq 0}^K$ be its count vector. Then

$$\bar{H} = \frac{1}{L} \sum_{i=1}^L h\left(\frac{c(X_i)}{N}\right) = \sum_x w(x) \hat{f}_x, \quad w(x) = h\left(\frac{c(x)}{N}\right), \quad (19)$$

identically — for every N , every alphabet size, every tree, and every model, reversible or not. Equation (19) is what “expanded in observable quantities” means for this statistic, and it has a feature that the reader may not expect: *the weights carry no parameter*. w is a function of an integer count vector, so it is a known constant of the analysis, fixed before any alignment is seen, and $\bar{H}$ is the sample mean of $W = w(X_1)$.

Lemma 23 (Exact mean and variance). $\mathbb{E}[\bar{H}] = \mu = \sum_x w(x)P(x)$ and $\text{Var}(\bar{H}) = \sigma^2/L$ with $\sigma^2 = \sum_x w(x)^2 P(x) - \mu^2$, exactly; σ^2 is the quantity left abstract in Theorem 19(2).

Proof. $\hat{f}$ is the empirical law of $X_1, \dots, X_L$ and the columns are independent and identically distributed under the model, so $\bar{H}$ is an i.i.d. mean of $w(X_i)$. $\square$

So the index *is* a moment, in the technical sense, and a well-behaved one: $0 \leq W \leq \ln K$ is bounded, so μ and σ^2 are finite and exact, the central limit theorem $\sqrt{L}(\bar{H} - \mu) \Rightarrow \mathcal{N}(0, \sigma^2)$ holds with no regularity beyond independence, and both moments are available from the same pattern enumeration as Theorem 7. Nothing prevents the index from being used in a method-of-moments, synthetic-likelihood or approximate-Bayesian scheme; it carries a computable standard error that nobody currently computes, and that is a fair criticism of current practice separate from the one this paper makes.

What it cannot do is identify the model, and the reason is visible in the same display. The map $\theta \mapsto \mu(\theta)$ is scalar. The space of linear functionals of $\hat{f}$ has dimension $\binom{N+K-1}{K-1}$, of which $\bar{H}$ is

exactly one; and on an equidistant star, where the leaves are exchangeable and $P(x)$ depends on x only through $c(x)$, the count histogram $\hat{g}_c = \sum_{x:c(x)=c} \hat{f}_x$ is *sufficient* with those $\binom{N+K-1}{K-1}$ entries. The index is one linear functional — the weight vector h — of a sufficient statistic with 84 entries at $N = 6$, 165 at $N = 8$ and 969 at $N = 16$.

That the discarded directions are not empty is checkable. At $N = 6$ on an equidistant star under GTR with Γ rates, the Jacobian of $\theta \mapsto \mathbf{g}(\theta)$ with respect to the ten parameters π, t_{ij}, d, α has full rank at both test points, and the elasticities $\theta_j \partial \log \mu / \partial \theta_j$ of the index’s own moment are

	π_A	π_C	π_G	t_{AC}	t_{AG}	t_{AT}	t_{CG}	t_{CT}	d	α
$\alpha = 70$	−0.039	+0.035	+0.013	+0.018	−0.061	+0.028	+0.013	−0.017	+0.458	+0.004
$\alpha = 1$	−0.037	+0.031	+0.018	+0.014	−0.048	+0.023	+0.010	−0.015	+0.418	+0.189

(20)

The composition and exchangeability coordinates move the index by at most a few per cent while the divergence moves it by more than 40%, and the shape moves it by 19% at $\alpha = 1$ but 0.4% at $\alpha = 70$. The information that separates those coordinates is present in the alignment — it is in the other 83 functionals — and what the index does with it is discard it.

This is not a defect peculiar to I_{ss} ; it is the price of being an *average of logs*. The sibling statistic on the other side of Jensen is the parilinear distance $\hat{d} = -\frac{1}{K} \ln \det \hat{F}$, where $\hat{F}$ is the observed state-pair matrix: a *log of an average*, hence nonlinear in the data and, by Lemma 1, exactly a sum of the eigenvalues, $\mathbb{E}[\hat{d}] = -\frac{1}{K} \ln \det \Pi - \frac{d}{K} \sum_k \lambda_k$ [Lockhart et al., 1994]. It is the same spectral decomposition read through a determinant instead of a trace, and the difference is that the log commutes with the sum on that side. Estimation lives on the log-of-average side; the index is on the other, and it is there by construction rather than by oversight.

2.8 Simulation: the design

Every claim below is checked against alignments simulated with AliSim [Ly-Trong et al., 2022] inside IQ-TREE 3.0.1 [Minh et al., 2020], under a custom Newick topology, explicit seeds, and base frequencies and rate categories set per arm. The one design decision worth stating is how the tree length is chosen. Sampling s uniformly wastes the budget in the saturated plateau, where the index is pinned; sampling it from

$$p_s(s) = \frac{\mu'(s)}{Z}, \quad Z = \mu(s_{\max}) - \mu(s_{\min}), \quad (21)$$

where $\mu(s) = \mathbb{E}[I_{ss}]$, makes I_{ss} uniform on $[\mu(s_{\min}), \mu(s_{\max})]$ exactly, since the induced density is $p_s/|\mu'| = 1/Z$. The normalising constant is the increment of μ itself and μ' is analytic, so no pilot run is needed. Panel C of Figure 3 is the check: the induced density is flat at $1/Z \approx 1.56$ for all four configurations. The same panel shows why the second-order form is not adequate for this purpose — its induced density drifts between 1.15 and 1.5 rather than staying flat.

2.9 Simulation: location and scale

Figure 4 maps the median and the standard deviation of I_{ss} over a grid of total tree lengths and sequence lengths, $N = 4$, 200 replicates per cell. The two halves of the picture behave differently, and that difference is the content of Theorem 19.

The median is a function of tree length alone. Reading across any row of panel A, it is flat in L : over a hundredfold range of sequence length the median moves by 0.5–2.0% of its value, at every tree length tested. This is the L -invariance of the location, and it is exact in expectation rather than approximate — Theorem 19(1) says $\mathbb{E}[\bar{H}]$ carries no L at all.




Figure 3: The uniform- I_{ss} sampling design of Eq. (21). **A**, the design curve $\mu(s)$ for four quartet configurations, exact (solid) against the second-order form (dotted); the dotted curves are shown past their range of validity, where they exceed the ceiling. **B**, the resulting sampling density on s . **C**, the density induced on I_{ss} , which is flat exactly when the sampler returns a uniform index; dotted is the second-order version.

The standard deviation is a function of sequence length alone, and it is where the L -dependence goes. Panel B shows the column-wise structure: SD depends on L and is nearly independent of tree length within a 2.5-decade range, falling as $L^{-1/2}$ — the ratio $\text{SD}(L)\sqrt{L}$ is constant to within 12% across the whole grid. An intensive statistic with an $L^{-1/2}$ standard error is exactly what makes the test statistic grow as $\sqrt{L}$, and this is that fact in simulation form. Figure 5 reads the same grid along its axes, so the flatness in L and the $L^{-1/2}$ decline can be seen directly rather than through a colour scale.

2.10 Simulation: the affine relation and the role of N

Figure 6 is the direct test of Theorem 13 and Corollary 8. Panels A and B use the same shape at three taxon counts — star trees at $N = 4, 8, 16$ — so that N is the only variable, and overlay the exact expectation of Theorem 7 on every cloud. The agreement is to 0.0006 at every one of the twenty-four cells, 200 replicates each; the closed form needs no fitting and no correction.

Against that baseline the two conventions separate. Under convention (A) the three exact curves lie almost on top of one another, and under (B) they fan apart, with $N = 16$ above $N = 4$ by 0.23 at $\bar{p} = 0.6$. The numbers are in the caption to Figure 6; the mechanism is Corollary 8. Xia’s normalisation is not a cosmetic rescaling: $H_{FSS}^{(A)}$ is itself a finite-sample quantity, and dividing by it removes the leading dependence on the number of taxa that the plug-in entropy retains.

Panel C shows the main arm, a quartet with a 35% internal branch at $N = 4$, where both conventions track their exact curves across four decades of tree length. The affine relation of Theorem 13 is a property of the process, not of the star.

2.11 Simulation: the Felsenstein–Farris transfer

The sharpest illustration of what the index does and does not measure is a one-parameter family that leaves the index fixed while moving the difficulty of the problem through its entire range. Take the unrooted quartet $((A, B), (C, D))$ with pendant lengths

$$(a, b, c, d) = (p - t, q, p, q + t), \quad t \in [0, p - q], \quad (22)$$




Figure 4: Location and scale of I_{ss} over (total tree length) $\times$ (sequence length), JC69, $N = 4$, 200 replicates per cell. **A1**, median with equal pendants; **A2**, median with pendants graded in a fixed ratio. **B1**, standard deviation over the same grid with equal pendants; **B2**, graded. The colour scale is separate for each row, since the median spans the whole range while the SD does not.

so that at $t = 0$ the two long branches are separated by the internal branch — the Felsenstein zone [Felsenstein, 1978] — and at $t = p - q$ they are adjacent, the Farris zone [Hendy and Penny, 1989]. The sum of the six pairwise distances is $6p + 6q + 4r$, independent of t ; and at $r = 0$ the endpoint distance multisets coincide exactly, so I_{ss} is identical at the two ends.

Figure 7 confirms both halves and shows what they cost. Along the path the mean index varies between 0.6506 and 0.6967, a range of 0.046; the exact $O(r)$ endpoint gap at the simulated $r = 0.02$ is 0.0024. Over the same path MPBoot goes from recovering the true topology 0 times in 200 to 200 times in 200. The index does not track this at all: for I_{ss} between 0.64 and 0.70 — a band 4% as wide as the index’s own range — MPBoot takes the value 0 at some points on the path and 1 at others, while maximum likelihood over the same band never exceeds 0.68 and never falls below 0.38. The three criteria span $[0, 1]$ between them across an index interval narrower than the precision with which the index is usually reported. This is Corollary 16 made concrete: an index that is affine in mean p -distance cannot distinguish configurations that the mean p -distance itself does not distinguish, and these configurations differ by the full range of attainable outcomes.




Figure 5: The same grid read along its axes. **A**, median I_{ss} against L at five tree lengths, with the interquartile band; the curves are flat. **B**, SD against L ; the dashed line is $L^{-1/2}$. **C**, the full distribution against tree length.




Figure 6: I_{ss} against mean pairwise p -distance. Lines are exact (Theorem 7); points are simulated. **A**, star trees under convention (A); the three exact curves are within 0.022 of one another over the whole range. **B**, the same under convention (B), where they separate. At $\bar{p} = 0.599$ the exact values are 0.7770, 0.7810 and 0.7983 under (A) for $N = 4, 8, 16$, against 0.5144, 0.6566 and 0.7403 under (B). **C**, the $N = 4$ quartet with a 35% internal branch, both conventions.

2.12 Simulation: one parameter at a time

Figure 8 moves one knob at a time from a common baseline and records I_{ss} , with no composite index and no power calculation — the aim is to show the shape of each dependence, not to rank tests.

Two knobs dominate. Total tree length moves the index across most of its range, from 0.280 to 0.929 — unsurprising, and the point of the index. Rate heterogeneity is next: $\alpha = 4$ gives 0.673 against the homogeneous 0.717, and $\alpha = 0.25$ gives 0.387, so the transition/transversion ratio and rate heterogeneity are comparable in effect at this scale. The transition/transversion ratio moves the index from 0.719 to 0.574 over $\kappa \in [1, 20]$.

The remaining knobs are weak. The internal fraction moves it between 0.678 and 0.723 over a sixfold range, base composition between 0.572 and 0.607 across uniform to strongly AT-rich, and sequence length between 0.7175 and 0.7229 over a hundredfold range — the last being the L -invariance again, from a second and independent design.

The ordering is the result. Of the five knobs a user might vary, two move the index substantially and three barely move it, and the three that barely move it are precisely the ones that govern




Figure 7: The transfer of Eq. (22), $N = 4$, $L = 1000$, 200 replicates per point, JC69. **A**, recovery of the true topology against the transfer parameter, for three criteria. **B**, the same path in I_{ss} coordinates: the exact expectation is a shallow arc rising to 0.696 at mid-path, with an exact endpoint gap of $0.0024 = O(r)$ at $r = 0.02$. **C**, recovery against I_{ss} : over a band of index values 4% as wide as the index’s own range, MPBoot takes both extreme values while maximum likelihood stays between 0.38 and 0.68.




Figure 8: One parameter at a time from a common baseline (JC69 quartet, $TL = 2$, $L = 1000$, 200 replicates), with no composite power index. **A**, rate heterogeneity α ; **B**, transition/transversion ratio; **C**, internal-branch fraction; **D**, total tree length; **E**, base composition. Error bars are ± 1 SD across replicates.

whether the topology can be recovered.

2.13 Conditional independence

Remark 18(iii) leaves open whether the higher-order channel — the part of the site-pattern distribution that the mean mismatch proportion cannot see — carries anything beyond composition. Reading it off a correlation between I_{ss} and $\bar{p}$ cannot answer that, because two statistics can be strongly correlated and still not be redundant. The question is whether I_{ss} predicts a phylogenetic outcome *after* $\bar{p}$ is accounted for,

$$\Pr(Y \mid \bar{p}, I_{ss}) \quad \text{versus} \quad \Pr(Y \mid \bar{p}), \quad (23)$$

and it is answered by comparing the two out-of-sample.

The design must be crossed for the question to be meaningful: if tree scale is the only knob then I_{ss} and $\bar{p}$ are deterministically linked to second order and there is nothing to condition on. We therefore vary shape (internal-branch fraction and pendant arrangement), composition, and model jointly, giving $720 \text{ cells} \times 150 \text{ replicates} = 108,000$ alignments in which I_{ss} and $\bar{p}$ are not in one-to-one correspondence.

Figure 9 shows the comparison. Three targets are reported. The fraction of parsimony-informative sites is exact and requires no inference. Recovery of the true quartet is measured twice, by MPBoot and by maximum likelihood in IQ-TREE — two programs we did not write. All fits use a cubic spline in $\bar{p}$ rather than a polynomial, for a reason worth stating: a quadratic in $\bar{p}$ is worse than chance in the transition band below, and comparing against it makes I_{ss} look harmful over the full range when it is merely redundant there.

The full range and the transition band answer different questions, and we report both. Over the full range $\bar{p}$ alone already separates almost perfectly (AUC 0.889 for MPBoot, 0.914 for ML) and there is nothing left to add: $\Delta\text{AUC} = +0.001$ and $+0.002$. The parsimony-informative fraction is different — it is a compositional statistic that $\bar{p}$ summarises poorly, and there I_{ss} raises cross-validated R^2 from 0.227 to 0.297.

Restrict to the band $\bar{p} \in [0.63, 0.70]$, where the recovery probability is neither 0 nor 1 and the decision is actually made, and the picture inverts. There $\bar{p}$ alone is nearly useless — AUC 0.555 (MPBoot) and 0.604 (ML) — and I_{ss} lifts it to 0.724 and 0.785, gains of $+0.17$ and $+0.18$. Used on its own, with no $\bar{p}$ term at all, I_{ss} reaches AUC 0.750 and 0.800 in the band, and a single threshold at $I_{ss} \approx 0.98$ classifies 80% and 83% of the alignments correctly.

So the index is not information-free, and the honest statement is stronger than "it is redundant". Two things bound what that information is.

First, the window in which it operates is a sliver at the top of the scale. Across the band, which spans a 0.07 range of $\bar{p}$ covering the whole transition from reliable to unreliable recovery, I_{ss} varies only between 0.79 and 1.03, with the interquartile range 0.87 to 0.97. The index discriminates recovery only within a few percent of its ceiling — which is precisely the regime Section 2.6 identifies as one where it cannot discriminate divergence. The information is real, but it is not saturation information: the quantity that moves I_{ss} within that window is not the amount of substitution.

Second, part of that information is tree shape. The internal-branch fraction — a property of the true tree, which no user observes — reaches AUC 0.759 in the band for MPBoot, better than I_{ss} 's 0.724; and over the full range adding it raises AUC from 0.889 to 0.923, after which I_{ss} adds nothing. For ML the ordering reverses (0.753 against 0.785), and in all three targets I_{ss} retains something after the shape is included, so the two are not substitutes. The summary is that I_{ss} carries genuine predictive information about recovery, that the information lives in a narrow band at the ceiling rather than across the range, and that it overlaps substantially with the shape of the tree that generated the data.

2.14 Empirical: where real data sit

Everything above is about the statistic. Whether any of it bites in practice depends on where real alignments place the index, and that is measurable. We use the data set Xia and Lemey [2009] introduce the index with: eight vertebrate mitochondrial COI sequences, 1512 nt, no gaps or ambiguity, with a reference topology shipped alongside. The alignment is analysed whole and by codon position, which is how the index is normally applied to a protein-coding gene.

The results are in Table 2 and Figure 10. Third codon positions carry an index of 0.715, first and second positions 0.137 and 0.038, and the whole alignment 0.270. The source reports this same partition, through DAMBE, as $I_{ss} = 0.7145$: the two conventions agree to 0.0005 here, and the agreement is not a coincidence but a consequence of the difference being confined to the treatment of invariable sites, which this partition does not have.

Read against the boundary the source itself reports for it — $I_{ss,c} = 0.7518$ for a symmetrical topology and 0.6401 for an asymmetrical one, both at $N = 8$ and $L \approx 500$ — the partition *passes* one and *fails* the other, by 0.037 and 0.075 respectively. The verdict is a property of which of the

two definitions of the critical tree length is read, and both are in the table.

Those two numbers are not free constants of the index, and it is worth saying what they are a property of, because the source says it. They are the recovery threshold of a particular inference engine applied to a particular generator: the calibration simulated under F84 and reconstructed with neighbour joining on F84 distances and with maximum likelihood under the same F84 model for 4, 8 and 16 taxa, and — quoting the methods — “only the neighbour-joining method is used for reconstructing topologies with NOTU > 16” [Xia et al., 2003]. The tabulated surface is therefore indexed by a coordinate that is set by the software rather than by the data, and never listed as a condition of use. The later manual states the consequence itself: “establishing the existence of phylogenetic information in the sequences does not mean that the true tree can be recovered by any tree-building algorithm” [Xia and Lemey, 2009]. A quantity defined as the point at which recovery fails inherits the engine that failed.

The size of that inheritance is measurable on our own quartet, and it is small. Rescoring the 150 replicates per tree length of Section 2.6 with neighbour joining on JC distances in place of maximum likelihood in IQ-TREE gives a critical tree length of 7.31 against 7.19 and an $I_{ss,c}$ of 0.9817 against 0.9807 — a difference of 0.001, against the 0.0115 the α sweep spans. Two engines that differ substantially in how they recover a quartet happen to agree here on *when* they stop, so the engine is a provenance fact in this comparison rather than a measured effect. It remains a coordinate the table does not carry, and for the larger taxon counts, where the source switches to neighbour joining alone, it is the same coordinate the user cannot vary.

That is the resolution collapse of Section 2.6 appearing in real data. The partition every practitioner would identify as the saturated one sits between the two readings of its own threshold, so that the answer to “is this saturated?” is decided by a convention rather than by the alignment. The tabulated surface is also computed from 100 replicates per point with no interval reported, so margins of a few percent are not distinguishable from zero.

One caveat on comparing our columns with DAMBE’s, since it is the reason the agreement above is worth checking rather than assuming. DAMBE’s index divides the mean entropy by $1 - P_{inv}$ when an invariant-site proportion is supplied, so the numerator becomes the mean entropy *per variable site* while the denominator is unchanged. For the first and second codon positions of this alignment, where $P_{inv} = 0.74$ was estimated and entered, that inflates the index from 0.087 to the 0.3335 the source reports — a factor of $1/(1 - P_{inv})$, exactly. Both are legitimate readings of the same construction and they answer different questions: ours asks how saturated the alignment is, theirs how saturated its variable sites are. Because P_{inv} is estimated, a threshold on their scale is implicitly a threshold on an estimated quantity, and is not comparable across data sets with different rates of invariance. Which of the two scales the published $I_{ss,c}$ surface was computed on is not stated in the source, and we cannot settle it from the published text; if it is the rescaled one, then a user who reads the index without supplying P_{inv} is comparing across scales, and that is a wider exposure than the missing rate coordinate, because P_{inv} is a parameter users routinely set.

The informative-site counts make the point sharper. Second codon positions hold 2 parsimony-informative sites out of 504 and 466 constant sites; maximum likelihood cannot recover the reference topology from them, and the index calls the partition emphatically unsaturated at 0.038. First positions, with 4 informative sites, do recover it. So the ordering the index produces — second positions safest, third positions worst — is the opposite of the ordering by signal, and the partition it clears most confidently is the one carrying least.

One qualification, because a sparse partition and a saturated one are different regimes that both end up with few informative sites. Second positions are sparse: 466 of 504 sites are constant, so there is little signal because there is little variation, and the index correctly reports that by returning a low value. A saturated locus is the opposite kind of case: many variable sites of

partition	sites	$\bar{p}$	$I_{ss}^{(A)}$	$I_{ss}^{(B)}$	informative sites	RF
second positions	504	0.0295	0.038	0.031	2	2
first + second	1008	0.0712	0.087	0.073	6	0
first positions	504	0.1129	0.137	0.115	4	0
all sites	1512	0.2181	0.270	0.226	42	0
third positions	504	0.5120	0.715	0.594	36	2

Table 2: Vertebrate mitochondrial COI, eight taxa, whole alignment and codon partitions. RF is the Robinson–Foulds distance between the maximum-likelihood tree under GTR+ Γ and the reference topology supplied with the data (zero means recovered). The tabulated boundaries at $N = 8$ and $L \approx 500$ are $I_{ss,c} = 0.7518$ and 0.6401 , symmetrical and asymmetrical; the third-position partition passes the first by 0.037 and fails the second by 0.075 . Note that second positions, which the index clears by the widest margin, hold two parsimony-informative sites. Our I_{ss} values use the per-site convention throughout; the source’s $I_{ss} = 0.7145$ for third positions agrees with ours to 0.0005 , and its larger values for the conserved partitions reflect the $1/(1 - P_{\text{inv}})$ normalisation described in the text.

which few are parsimony-informative, so the index rises while informativeness falls. That is the association reported by Duchêne et al. [2022], whose loci flagged as saturated tend to be those with few informative sites, and it is the same association measured in Section 2.13. Both regimes are unresolvable, the index separates them, and only one of them is saturation. Passing the test is evidence against saturation, as the original paper claims; on this data set it is not evidence of resolution.

2.15 Empirical: a genome-scale alignment

The COI alignment is eight taxa and one gene. The opposite end of the range is available: the largest source in the survey of Duchêne et al. [2022] is the avian data set of Jarvis et al. [2014], whose ExaML concatenation is distributed through GigaDB and covers 51 taxa over 322,150,876 aligned sites. We index that file directly, streaming the 4.6 GB archive rather than expanding it, by accumulating the per-site count vector over the four states — which is all the index needs, and which fits in memory where the alignment does not.

The alignment is sparse: 90.1% of sites carry a gap or an ambiguity code in at least one taxon. The 9.9% that do not — 31,962,462 sites — are complete columns with all 51 taxa present, and those are what the index is computed on. Over them $\bar{p} = 0.0786$, $\bar{H} = 0.1595$, and

$$I_{ss}^{(A)} = 0.119, \quad I_{ss}^{(B)} = 0.116. \quad (24)$$

Splitting the concatenation into 31,962 windows of 1000 complete sites gives a distribution from 0.006 to 0.230, and not one window exceeds 0.5, let alone the tabulated boundary.

The restriction that Duchêne et al. [2022] recommend — computing the entropy on parsimony-informative sites only, 37.0% of the complete columns — raises the whole distribution without moving it into the interesting region. Those 11,833 windows give $I_{ss}^{(A)} = 0.268$ for the subset as a whole and a window range of 0.221 to 0.393; still no window reaches 0.5. Panels C and D of Figure 11 show it. The modification is doing what it was designed to do, doubling the index by removing the constant and near-constant sites that dilute it, and the verdict on this alignment is unchanged either way: it is not saturated under either reading, by a wide margin.

That recommendation is not covered by the theory of Section 2.4, and since this paper repeats it, the gap is worth closing. Theorem 13 is a statement about columns drawn independently from the




Figure 9: Conditional independence: does I_{ss} predict a target beyond the mean p -distance? **A–C**, within each bin of $\bar{p}$ the alignments are split at the bin’s median I_{ss} and the outcome is plotted for the lower and upper halves; the targets are the parsimony-informative fraction, MPBoot recovery and maximum-likelihood recovery, and the mean separation between the two curves is 0.025, 0.058 and 0.055 respectively. A gap is information that $\bar{p}$ does not supply, and the gap opening in the transition region is what the cross-validated fits below measure. **D**, what I_{ss} adds out of sample, over the full range against the transition band $\bar{p} \in [0.63, 0.70]$. The full-range gains for recovery are indistinguishable from zero, which is what a naive comparison reports; the band gains are an order of magnitude larger and are confined to the only region where the recovery decision is in doubt.

model; conditioning on a column being parsimony-informative changes that law, so Lemma 4 does not apply and the criterion is not formally available. What it costs is measurable. On an $N = 6$ star the residual of $\mathbb{E}[\bar{H}]$ against the mean p -distance is 1.8% of its range for GTR with uniform composition and 1.7% for HKY with $\kappa = 5$ and skewed composition — one satisfying the criterion, one not — and 1.4% and 1.7% respectively under the restriction. The sharper comparison is between two models that both satisfy the criterion exactly: JC69 and GTR with uniform composition disagree by at most 1.0% of the range at matched p -distance on all columns and 0.9% on informative ones. The restriction therefore adds no departure beyond the $O(3)$ remainder that Section 2.7 already shows masks the criterion at the level of $\mathbb{E}[I_{ss}]$. The criterion is not a theorem there, and on this evidence it is also not violated.

The value is corroborated by the theory rather than merely plausible. A JC69 star with the same mean p -distance and $N = 51$ predicts $\mathbb{E}[I_{ss}^{(A)}] = 0.134$ exactly; the observed 0.119 is a little below that, which is the right sign and roughly the right size, since a real tree at a fixed mean p -distance is less saturated than the star that maximises the number of paths between leaves. Nothing about the number is anomalous; the alignment is simply not saturated by this measure.

Read against Section 2.14, the two data sets bracket the problem. The third codon positions of vertebrate COI sit at 0.715, between the two readings of their own threshold: the index is at its ceiling and the resolution collapse of Section 2.6 applies. A genome-scale avian alignment sits at 0.119, five to six times below the boundary, where the collapse does not apply and the index is behaving as intended. Both regimes are real, which is why the question of where a given matrix sits has to be answered rather than assumed — and it is answerable, from the alignment, before any tree is built.

Sites with gaps or ambiguity codes are excluded rather than imputed, which is what DAMBE does as well, so the treatment matches the practice the comparison is drawn against; 0.119 is the index of the complete columns, and those are the conserved fraction of the alignment.




Figure 10: The vertebrate COI data set. **A**, the index by partition, against both tabulated boundaries the source reports for it, 0.7518 symmetrical and 0.6401 asymmetrical, which it falls between; green partitions recover the reference topology, red ones do not. **B**, the index is not a function of how much parsimony-informative signal a partition carries. **C**, recovery against the index: the two failures lie at opposite ends of the scale, for opposite reasons — third positions are saturated, second positions are empty.

2.16 The other tree-free indices

Two other indices are computed from an alignment alone and are therefore available at exactly the point where I_{ss} is used: the C (convergence) value of Struck et al. [2008] and λ of Jermin and Misof [2020]. We compute both on the alignments above, and the comparison turns out to be analytic rather than merely numerical for one of them.

λ is defined per pair as

$$\lambda_{ij} = \frac{\delta_{\text{obs}}(ij)}{\delta_{\text{ran}}(ij)}, \quad \delta_{\text{ran}}(ij) = 1 - \sum_k f_{ik} f_{jk}, \quad (25)$$

where f_{ik} is the frequency of state k in sequence i . The denominator is a counterfactual: the divergence expected of two sequences that have evolved independently to equilibrium. It is not a reshuffling of the observed sites but an independent stationary process, and the endpoints are $\lambda = 0$ when the variable sites agree and $\lambda = 1$ when the pair is random with respect to itself.

The semantics decide how the quantity can be used, and they are not the semantics of a substitution count. $\lambda = 0$ does not mean no substitutions occurred, only that the variable sites are identical; $\lambda = 1$ does not mean every site has substituted repeatedly, only that the pair has become random. What λ measures is how much of the *historical* signal remains — information attributable to the order and timing of divergences — which is a different question from how many substitutions produced the data. We take the framing from the source, which is explicit that λ quantifies “loss of historical signal” and that $\lambda = 0$ and $\lambda = 1$ are definitional endpoints rather than decision thresholds; it supplies no critical value, no standard error and no test.

δ_{ran} also has a reading this paper already contains. Setting $\xi_\mu \rightarrow 0$ in Lemma 2 gives $\wp(\infty) = 1 - \sum_j \pi_j^2$, the value the mean p -distance approaches under saturation for any reversible model. That is the same quantity, reached from an independence argument on one side and from the spectral decomposition on the other. The familiar upper bound of 3/4 is therefore not a constant of nature but the value of $1 - \sum_k q_k^2$ at uniform composition: it is 0.700 at (.40, .30, .20, .10), 0.590 at (.45, .45, .05, .05) and 0.932 for amino-acid data. For the third codon positions of the COI alignment




Figure 11: The avian genome-scale alignment of Jarvis et al. [2014], 51 taxa over 322 million sites, indexed directly from the compressed archive. **A** and **B**: all complete columns, 31,962 windows of 1000 sites, against the whole-alignment value and the threshold the index is normally read against, then in the $(\bar{p}, I_{ss})$ plane with the exact expectation for a JC69 star at $N = 51$. Orange lines mark the vertebrate COI partitions of Section 2.14, one of which sits at the boundary. **C** and **D**: the same restricted to parsimony-informative sites, the modification Duchêne et al. [2022] recommend, which shifts the whole distribution from 0.119 to 0.268 and leaves the conclusion intact.

it is 0.684, so the rule of thumb $\bar{p}/0.75$ reports that partition as 32% short of saturation while λ reports three quarters of the historical signal gone. Both follow from the same arithmetic; the constant is answering the wrong question, by an amount that grows with compositional skew — and in the other direction, by +0.18, for protein data.

Now the part that bears on this paper. If the sequences share a composition — $f_{ik} = q_k$ for every i , the second-order regime the rest of this paper works in — then δ_{ran} is the *same* constant for every pair, so λ is the pairwise p -distance multiplied by a constant and

Proposition 24 (The mean λ is the mean p -distance rescaled). *If every sequence has composition $\mathbf{q}$, then $\lambda_{ij} = p_{ij}/(1 - \sum_k q_k^2)$ for every pair, and therefore $\bar{\lambda} = \bar{p}/(1 - \sum_k q_k^2)$.*

Proof. $\delta_{\text{obs}}(ij)$ is by construction the proportion of differing sites between i and j , i.e. p_{ij} ; $\delta_{\text{ran}}(ij)$ reduces to $1 - \sum_k q_k^2$ under the stated condition. Average over pairs. $\square$

This is not an approximation that happens to hold here. Table 3 evaluates it on seven alignments spanning I_{ss} from 0.038 to 0.749, and the right-hand side reproduces the observed $\bar{\lambda}$ to four decimal




Figure 12: λ on six taxa with one long pendant (1.6 expected substitutions against 0.05 elsewhere), gamma rates $\alpha = 1$, 4000 sites, averaged over three replicates. The global summaries are unremarkable — $I_{ss} = 0.348$, $\bar{\lambda} = 0.370$ — while the pairwise matrix isolates the five comparisons involving the long-branch taxon. This resolution among pairs is what λ is for, and it is what the mean, and any single global index, discards.

places on every one of them; across the seven $\bar{\lambda}$ correlates with $\bar{p}$ at $r = 0.999$ and with I_{ss} at $r = 0.998$. As a single scalar, then, $\bar{\lambda}$ carries what $\bar{p}$ carries and inherits everything Theorem 13 establishes about it.

The scope of that is narrower than it looks, and we state it because the opposite reading is available. It is a statement about the *mean over pairs*, and the mean is not what λ is for. λ is defined pairwise, the source’s own tooling plots the matrix as a heat map rather than a summary, and the resolution among pairs is the point: no global number can say *which* sequences are the problem. Figure 12 is a case in point. Six taxa, one of them on a long pendant, gamma rates, 4000 sites: the global summaries are unremarkable — $I_{ss} = 0.348$, $\bar{\lambda} = 0.370$, neither near anything — while the matrix separates the five pairs involving the long-branch taxon at 0.71–0.72 from the other ten at 0.12–0.22. Proposition 24 holds exactly on these data, the ratio being constant as it must be, and the structure that matters is invisible to it. The index of this paper shares that blindness and more so: being a ratio of two entropies, it has no pairwise form at all.

Two limitations belong alongside, both of which the source acknowledges in part. λ is built on the uncorrected p -distance and carries no standard error, so near the bound its decisions are close to coin flips: at the bound, $\Pr(\hat{p} \geq 1 - \sum_k q_k^2)$ is near one half at every length we tried, from 200 sites to 20000. The source notes that $\lambda > 1$ occurs occasionally and attributes it to finite alignment length or to sequencing and alignment error, which is the same admission without the magnitude. And the advantage claimed for it — using each sequence’s own composition makes λ “less likely to be overestimated” when lineages diverge under different conditions — has a cost that neither the source nor we had stated: once composition varies, δ_{ran} is no longer the model’s saturation bound but an empirical baseline, and the equivalence between $\lambda < 1$ and the existence of an interior maximum-likelihood estimate no longer follows. The index of this paper does not enter that regime, since it assumes a single stationary composition throughout.

The C value behaves differently. It is $\sigma(\text{ti/tv})/\sigma(p)$ over all pairwise comparisons, a ratio of two spreads rather than a ratio of two divergences, and it does not track I_{ss} : across the same seven alignments the correlation is -0.29 , and the ordering it produces is not the ordering by signal. On the COI partitions it ranks second positions ($C = 55.8$, $I_{ss} = 0.038$) as *more* saturated than first positions ($C = 115.3$, $I_{ss} = 0.137$), while third positions — the ones everyone agrees are saturated — come bottom at $C = 43.1$. A statistic built from the transition/transversion ratio is measuring the substitution process, not the amount of it, and the two come apart on conserved partitions where there is little of either. We report C because it is in use, and because it is the one of the three that

alignment	sites	$I_{ss}^{(A)}$	λ	C	$\lambda/\bar{p}$
vertebrate COI, 2nd positions	504	0.038	0.041	55.8	1.404
vertebrate COI, 1st + 2nd	1008	0.087	0.096	221.8	1.344
avian genome-scale, complete columns	31,962,462	0.119	0.108	7.1	1.345
vertebrate COI, 1st positions	504	0.137	0.151	115.3	1.339
avian genome-scale, informative	11,833,701	0.268	0.253	4.0	1.338
vertebrate COI, all sites	1512	0.270	0.294	101.8	1.350
vertebrate COI, 3rd positions	504	0.715	0.749	43.1	1.462

Table 3: The three tree-free indices on the same alignments. λ is the mean over all pairs, and it tracks I_{ss} closely, as Proposition 24 requires; the last column is its ratio to the mean p -distance, which is $1/(1 - \sum_k q_k^2)$ and therefore nearly constant. That column also gives the saturation bound for each alignment, which is 0.75 only where the composition is uniform. The mean is a compression of λ rather than its definition, and the compression is what makes the proposition apply; Figure 12 shows what it discards. C does not track I_{ss} , and in particular does not place the saturated third codon positions at the top of its scale.

is genuinely independent of the p -distance; we do not claim the criterion above applies to it.

3 Discussion

The two halves of the test come apart. The index is an affine image of the mean pairwise p -distance whenever the composition-weighted term in Eq. (7) is proportional to the phylogenetic term — exactly so for uniform composition or a degenerate spectrum, approximately so for nucleotides generally. Where that holds, the entropy construction adds nothing about the number of substitutions actually realised beyond the uncorrected proportion of differing sites. Where it does not, what the index adds is the composition of the model rather than the divergence, and the size of that addition is $\Delta = \sum_{\mu \neq 0} (w_\mu - \tau_\mu/K) \xi_\mu$, a function of π and the exchangeabilities and of nothing else; by Eq. (9) the deviation of $\mathbb{E}[I_{ss}]$ from its p -distance prediction is $-\frac{N-1}{2Nh(\pi)} \Delta$ to second order, so Δ is the size of the failure and the criterion is its shape. Saturation, in its standard sense, is the gap between the observed difference and the true divergence; a quantity affine in the former inherits the gap rather than seeing through it. The rate of loss follows: the index resolves true divergence at the rate the uncorrected distance does, exponentially under rate homogeneity and as a power law under Γ -distributed rates.

The simulation of Section 2.13 is the one place where this statement has to be read with care, and it is worth stating what it does and does not show. Over a crossed design spanning models, composition and tree shape — a design in which the affine relation is not maintained, because the user has not fixed what it is conditional on — the index does predict recovery in the transition band better than $\bar{p}$ alone, by +0.17 and +0.18 in AUC. That information is real. It is also not saturation information: it tracks the shape of the generating tree and the higher-order composition of the sites, and it is confined to a window a few percent wide beneath the ceiling, which is the window where the preceding paragraph says the index cannot see divergence. The two findings are consistent, and their conjunction is the sharper claim: the index carries predictive information whose source is not the quantity it is named after.

The critical value, meanwhile, is tabulated without a rate-heterogeneity coordinate, and in the region where the test is actually applied the index is close enough to its upper bound that the threshold cannot distinguish moderate from severe saturation. Its dependence on the sequence

length, by contrast, is deliberate and necessary: the threshold must move with L because the statistic's standard error shrinks as $L^{-1/2}$, and the original paper reports that dependence and reasons from it. The practical consequence is asymmetric. Passing the test is evidence against saturation — the index would be near its ceiling if the data were saturated. Failing it is weak evidence, and is not by itself grounds for discarding sites; the remedy for insufficient resolution is more data or more taxa, not fewer sites.

3.1 What the 2003 paper got right

The critique above should not obscure four things.

First, the problem was real. Before 2003, saturation was assessed by plotting transition/transversion ratios against distance, or by permutation and skewness tests. Both report signal in a data set that contains two closely related species and an arbitrary number of fully saturated ones. The paper states that defect explicitly and attributes it to the first three of the four approaches it reviews. The fourth is faulted on different grounds — clumsy beyond four taxa, parsimony-specific, undeveloped [Steel et al., 1993, Halanynch, 1995]. The 2003 paper's own point of entry is that the fourth approach's diagnostics had not been developed further, and that assessment has aged well.

Second, and most durably, the paper's central observation is correct: sequences fail to recover the true tree long before I_{ss} reaches 1, so the quantity of interest is a critical value, not the full-saturation limit. That observation is what motivated later work to rebuild the test around it [Duchêne et al., 2022]. It should also be said that the critical tree length is defined *two different ways* in the original: for the symmetrical topology it is the tree length at which $P_{\text{true}} = 0.95$, whereas for the asymmetrical topology, where the maximum recovery probability need not reach 1, it is the tree length at which P_{true} falls to 95% of its own maximum. The two are not comparable, and every departure below is measured on the first; the tabulated asymmetrical values therefore carry an additional, unquantified source of drift.

Third, the simulation grid — F84, transition/transversion ratio from 1 to 10, base frequencies from 0.1 to 0.9, 4 to 32 taxa, two extreme topologies, 100 replicates per point — is a serious design for its time, and the paper reports the grid it used, together with the finding that composition and transition/transversion ratio are negligible relative to topology, taxon number and sequence length.

Fourth, the discussion section does not ignore rate heterogeneity, but what it does with it is worth setting out precisely, because the usual reading — that the author was aware of the problem and the calibration simply predates it — is more generous than the text supports.

The section defines the rates-across-sites model as two halves: “A typical RAS model includes a proportion (P) of invariable sites ... and the rest of sites evolving at the Γ -distributed rates”. It then argues that the test is applicable to the $(1 - P)$ fraction, that DAMBE can estimate P , and that assuming $P = 0$ when $P > 0$ makes the test conservative. That argument is correct, and it is correct about the first half only: invariable sites depress $\bar{H}$, hence I_{ss} , hence the rate of rejection, which is the same channel as the $+I$ under-detection of Section 2.6. The second half — the sites that do evolve, at Γ -distributed rates — never enters the argument, and the calibration that produced the table is rate-homogeneous. Conservativeness with respect to P says nothing about it, and a reader who takes the paragraph as covering the RAS model has been told something the paragraph does not show.

What makes this more than an omission of scope is the sentence two lines earlier, in the same paragraph: the RAS model “has been studied recently in relation to phylogenetic reconstruction ... because most commonly used phylogenetic methods fail to recover the true tree when sequences evolved according to the RAS model”. The source of that claim is a simulation under unequal rates [Kuhner and Felsenstein, 1994], which is the regime the table was built in — and the table is

precisely a statement about when commonly used methods recover the true tree. The observation was available, in the paragraph, and was used to motivate the RAS discussion rather than to question the homogeneous grid. The criticism here is not that a 2003 paper failed to anticipate 2005 practice; it is that the applicability limit was stated as a property of P when the model it was stated about has two parameters.

The covarion discussion has the same shape in miniature. It concludes that the covarion process “does not seem to cause any complication in our test”, on the grounds that sequences evolving under covarion are much less susceptible to saturation than under RAS [Penny et al., 2001]. Susceptibility to saturation is a property of the data; what the test needs is a statement about the statistic, and less saturation is not the same as a preserved relation between I_{ss} and the tree length at which recovery fails. The conclusion may well be right; the premise does not reach it.

The failures are of a different kind: a normalisation whose entropy form implies more than the second-order expectation delivers, a calibration whose conditions were not stated as conditions of use, and a threshold whose behaviour is governed by sample size and by the proximity of the index to its own ceiling. None of these makes the reasoning wrong; together they make the verdict the index issues unsupported.

3.2 Relation to other approaches

The index has been questioned before, and the earlier objections are different in kind from the one made here. Verbruggen et al. [2007] and Letsch and Kjer [2011] report cases where the tabulated threshold gives the wrong verdict; both accept the construction and dispute the calibration. That is a disagreement about a number, and it leaves open the possibility that a better-calibrated threshold would work. The argument above is about the statistic: $\mathbb{E}[I_{ss}]$ is an affine function of the mean pairwise p -distance under conditions that hold for every reversible model at uniform composition and for every composition with a degenerate spectrum, so no threshold, however calibrated, can make it carry resolution the p -distance does not carry. The two critiques agree on the practical verdict and disagree on what would fix it.

The most substantial recent development touches the argument at an unexpected point. Duchêne et al. [2022] keep the entropy construction but restrict it to parsimony-informative sites, on the grounds that constant and slowly evolving sites distort the base-composition measurement on which the index depends. That is the same channel Section 2.13 isolates from the other direction: of the three targets we test, the parsimony-informative fraction is the one I_{ss} predicts beyond the mean p -distance, and their empirical finding that saturated loci are those with few informative sites is the same association read as a diagnostic. They also report that the tests are sensitive to sequence length and to short internal branches, which is Theorem 19 and the shape term of Section 2.6 respectively, arrived at by simulation.

What none of this changes is that the quantity being thresholded still reads only apparent divergence. Modifying which sites enter the entropy changes the statistic, but the modified statistic remains a function of the same site-pattern composition, and no choice of sites makes it a function of the number of substitutions that produced them.

Posterior branch-level tests such as Manuel et al. [2025] are a different construction and not subject to the criterion above. Given a tree and a model, saturation is identifiable per branch, and estimating it there is the right thing to do: the branch length and the number of substitutions implied by it are both in hand. What is not identifiable is saturation *before* a tree is committed to, which is the only regime in which a tree-free index is needed. The distinction is worth stating plainly because it delimits the critique: we are not arguing against measuring saturation, only against measuring it with this statistic and at that point in the workflow.

3.3 Where the collapse bites

Whether the rate-heterogeneity mis-calibration bites in practice depends on where real data place the critical point on the index curve, and the two alignments we measure answer it in opposite directions. The saturated partition of the COI alignment (Section 2.14) lands between the two readings of its own threshold — inside the collapse, where the threshold cannot discriminate. The genome-scale avian alignment (Section 2.15) sits at 0.119, five to six times below the boundary, where the index is behaving as designed and the collapse is irrelevant.

So the collapse is neither universal nor a limiting case that real data avoid. It is a property of where a particular matrix sits, and that is readable from the alignment before any tree is built, at the cost of one pass over the sequences. Two alignments chosen for their prominence in the literature disagree about which regime genomics is in, and the disagreement is the finding rather than a gap in it: the threshold is specific to the model, the topology and the recovery engine, all three of which belong to a particular analysis, so there is no regime that genomic data as such occupy. What a reader can be given is the value, which is a measurement, rather than the verdict, which is not.

4 Material and methods

4.1 Common notations and background

N sequences of length L related by a tree $\mathcal{T}$; leaf states generated by a reversible rate matrix Q stationary with respect to π , normalised so that $-\sum_j \pi_j Q_{jj} = 1$; K the number of states; $\Pi = \text{diag}(\pi)$; $S = \Pi^{1/2} Q \Pi^{-1/2}$ the symmetrised generator with eigenvalues $0 = \lambda_1 > \lambda_2 \geq \dots \geq \lambda_K$; P_μ the spectral projector onto the eigenspace of μ ; $\tau_\mu = \text{tr } P_\mu$ its multiplicity and $w_\mu = \text{tr}(\Pi P_\mu)$ its composition weight; d a branch or path length and $\langle \cdot \rangle$ an average over ordered leaf pairs and rate classes.

The two conventions for H_{FSS} are Eq. (2). Convention (A), $\mathbb{E}_{\text{Mult}(N, \hat{\mathbf{p}})}[h]$, is the one used by Xia et al. and by DAMBE; convention (B), $h(\hat{\mathbf{p}})$, is the plug-in composition entropy. They are related on any data set by Eq. (3).

4.2 Theoretical results

All proofs are given inline with their statements in Results. The two numerically delicate points are worth recording here. First, the spectral projectors must be evaluated in the orthogonal form of Eq. (5), $P_\mu = \sum_k u_k u_k^\top$, and not as the Lagrange interpolant $\prod_{\nu \neq \mu} (S - \nu I) / (\mu - \nu)$: at $K = 20$ the latter returns $\sum_\mu \tau_\mu \approx 34$ where the trace must be 19, with errors in w_μ of order 10^0 , because the interpolant is numerically unstable precisely when eigenvalues are close. The orthogonal form is exact, and we verified it against `expm` to 10^{-16} . Second, τ_μ is the algebraic multiplicity as a consequence of orthonormality rather than by a separate argument, so no care is needed in grouping eigenvalues. $K = 4$ is the last state space for which the non-zero eigenvalues have closed forms in radicals; the criterion itself does not depend on this, only the explicit expressions do.

The general reversible case is checked against an independent route. Equation (12) with $q_{\rho j}$ built from S in Eq. (13) is evaluated directly and compared with the pattern enumeration of Lemma 6 at $N = 8$: the two agree to 10^{-16} under both conventions, and the spectral resolution reproduces $\exp(Qt)$ to 10^{-16} at three values of t . The linear expansion of Eq. (19) is likewise checked against the per-site computation of $\bar{H}$, and the variance of Lemma 23 against the spread of 4000 synthetic alignments. Both checks, and the Jacobian table, are in the archived scripts. The Jacobian is a

central difference, and its rank and conditioning are unchanged when the relative step is moved from 10^{-4} to 10^{-6} , which is the only numerical care that table needs.

4.3 Simulation

Alignments were generated with AliSim [Ly-Trong et al., 2022] as implemented in IQ-TREE 3.0.1 [Minh et al., 2020]. Every cell stores its seed, so any alignment is reproducible bit for bit; trees are supplied as Newick so that branch lengths are controlled exactly; base frequencies and discrete Γ rate categories are set per arm. We use the AliSim mean-rate Γ discretisation [Yang, 1994], which is the convention the closed forms assume: the closed form and the simulator agree to 4×10^{-4} under it and disagree at 6×10^{-3} under the median discretisation.

Two designs are used. The first is the uniform- I_{ss} sampler of Eq. (21), applied per configuration with s obtained by inverting $\mu(s) = \mu_{\min} + UZ$ analytically. The second is the Felsenstein–Farris transfer of Eq. (22).

Tree shape is parameterised by the internal fraction f_i , the proportion of total tree length assigned to the internal branch, with pendants either equal or graded in a fixed ratio. Sequence lengths range over 100 to 10,000, tree lengths over roughly 0.2 to 100 expected substitutions per site, and replicates number 150 to 200 per cell.

The critical tree lengths of Section 2.6 come from a separate sweep: the $f_i = 0.35$ quartet at $N = 4$ and $L = 1000$, six Γ shapes from $\alpha = \infty$ to 0.5, and a geometric grid of fourteen tree lengths from 2 to 320 chosen to bracket the crossing at every shape, at 150 replicates per point. Each crossing is located by linear interpolation of P_{true} in $\log TL$ between the last grid point at or above 0.95 and the first below it; the bracketing pairs are in the archived output, so the interpolation can be checked or redone. The index values reported alongside them are exact and do not come from this simulation.

Recovery of the true topology is measured by three criteria. Maximum likelihood uses IQ-TREE with the generating model; maximum parsimony uses MPBoot [Hoang et al., 2018]. Both are established, independently published programs, and no searching criterion is computed from our own code: the search is the part of the pipeline where an unrecognised defect would be hardest to detect, and where a vetted program exists it is used. The third criterion is neighbour joining, which on a four-taxon quartet reduces to evaluating a deterministic function of the six pairwise distances — no search, no heuristic, and short enough to be checked by inspection. All three are used in the transfer experiment of Section 2.11; the conditional-independence experiment of Section 2.13 uses the two searching criteria only, since recovering a quartet from a p -distance criterion is not what the index’s critical value is calibrated against. For its JC-corrected distance we cap p at 0.74, since the correction diverges at $p = 3/4$ and returning ∞ would make the split sums arbitrary — a real limitation of distance methods in the saturated regime, reported as such.

4.4 Empirical data

The empirical alignment is `VertebrateMtCOI.FAS` as distributed with DAMBE and used by Xia and Lemey [2009] to introduce the index: eight vertebrate mitochondrial cytochrome oxidase I sequences, 1512 nt, in an RNA alphabet, with no gaps or ambiguity codes and no length variation. It ships with a reference topology, which we take as given rather than re-estimating. Codon partitions are taken by position modulo three from the reading frame supplied in the file. Sequence length is the same across codons, so no partition truncation is needed.

The index, the mean pairwise p -distance and the parsimony-informative count are computed on each partition exactly as in the simulations. Recovery uses IQ-TREE [Minh et al., 2020] under

GTR+ Γ , and is scored by the Robinson–Foulds distance between the inferred topology and the reference, over the five non-trivial splits of an eight-taxon tree (so the maximum is ten, and zero means the tree is recovered).

4.5 Evaluation criteria

The numerical comparison reported here is I_{ss} against mean pairwise p -distance, one parameter at a time. No composite power index is constructed and no power calculation is reported: the aim is to exhibit the shape of each dependence, not to rank tests by a summary statistic, and a composite index would require committing to a loss function that the application does not supply.

Two other tree-free *a priori* indices are in use and are the natural extension of the comparison: the C value of Struck et al. [Struck et al., 2008], defined as $\sigma(\text{ti/tv})/\sigma(p)$ over all pairwise comparisons, and λ of Jermin & Misof [Jermin and Misof, 2020], defined pairwise by Eq. (25). Both are computed here on the same alignments as I_{ss} , from their published definitions, and compared in Section 2.16 and Table 3.

For the conditional-independence test of Section 2.13 we compare out-of-sample performance of a model in $\bar{p}$ alone against one in $\bar{p}$ and I_{ss} jointly, with a cubic spline in $\bar{p}$ (16 knots) and five-fold cross-validation. The spline is not cosmetic: a quadratic in $\bar{p}$ is worse than chance in the transition region, so it understates what $\bar{p}$ alone achieves and makes I_{ss} look harmful there. For a binary target we report cross-validated AUC; for a continuous one, cross-validated R^2 . Correlation between I_{ss} and $\bar{p}$ is reported for context but is not evidence either way.

Two caveats on the conditional-independence numbers. The transition band $\bar{p} \in [0.63, 0.70]$ and the single-threshold value $I_{ss} \approx 0.98$ were both chosen by inspecting the full data, not nested inside the cross-validation, so the in-band gains are descriptive rather than unbiased estimates of out-of-sample performance; the full-range numbers are unaffected because they involve no such choice. And the two AUCs quoted for I_{ss} alone in the band are not comparable to the full-range AUC of $\bar{p}$: they are fitted and evaluated on a restricted sample where recovery is genuinely in doubt, than which nothing in the full-range fit is harder.

Use of AI-assisted tools

Claude Code (Anthropic), an AI coding assistant, was used in the preparation of this work. It assisted in drafting and editing the manuscript text, and in deriving and checking the mathematical results reported here. All content was reviewed by the author, who takes full responsibility for the work and for its conclusions; the tool is not an author and cannot take responsibility for it.

Data availability

All simulation code, the scripts that generate every figure and table, the per-cell parameter records and the stored seeds are archived together with the result tables this paper quotes. Each of the 108,000 alignments in the conditional-independence experiment, and every other alignment reported here, can be regenerated bit for bit from its stored seed and cell key. The empirical alignment and reference topology of Section 2.14 are redistributed with DAMBE and are included unmodified in the supplementary archive, with their provenance recorded. The genome-scale alignment of Section 2.15 is neither deposited nor redistributed: at 4.6 GB it is fetched, by a script in the code repository, from the address recorded in the script.

supplementary material: <https://doi.org/10.5281/zenodo.23293285>
code, results and figures: https://github.com/Gipsy-The-Shellier/Iss_simulations

Acknowledgments

We thank the developers of IQ-TREE and of MPBoot for making AliSim and the parsimony search available in a form that can be scripted, and the maintainers of TeX Live and of the Python scientific stack.

CRedit authorship contribution statement

Zhi-Jie Xu: Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing – original draft, Writing – review and editing, Visualization.

Declaration of competing interest

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- David A Duchêne, Niklas Mather, Cara Van Der Wal, and Simon Y W Ho. Excluding loci with substitution saturation improves inferences from phylogenomic data. *Systematic Biology*, 71(3): 676–689, 2022. doi: 10.1093/sysbio/syab075.
- J. Felsenstein. Cases in which parsimony or compatibility methods will be positively misleading. *Systematic Zoology*, 27(4):401–410, 1978. doi: 10.2307/2412923.
- Kenneth M. Halanych. The phylogenetic position of the pterobranch hemichordates based on 18s rDNA sequence data. *Molecular Phylogenetics and Evolution*, 4(1):72–76, 1995. doi: 10.1006/mpev.1995.1007.
- Michael D. Hendy and David Penny. A framework for the quantitative study of evolutionary trees. *Systematic Zoology*, 38(4):297–309, 1989. doi: 10.2307/2992396.
- Diep Thi Hoang, Le Sy Vinh, Tomáš Flouri, Alexandros Stamatakis, Arndt von Haeseler, and Bui Quang Minh. Mpboot: fast phylogenetic maximum parsimony tree inference and bootstrap approximation. *BMC Evolutionary Biology*, 18(1):11, 2018. doi: 10.1186/s12862-018-1131-3.
- Erich D. Jarvis, Siavash Mirarab, Andre J. Aberer, Bo Li, Peter Houde, Cai Li, Simon Y. W. Ho, Brant C. Faircloth, Benoit Nabholz, Jason T. Howard, et al. Whole-genome analyses resolve early branches in the tree of life of modern birds. *Science*, 346(6215):1320–1331, 2014. doi: 10.1126/science.1253451.
- Lars S Jermin and Bernhard Misof. Measuring historical and compositional signals in phylogenetic data, 2020.

1070 Mary K. Kuhner and Joseph Felsenstein. A simulation comparison of phylogeny algorithms under
1071 equal and unequal evolutionary rates. *Molecular Biology and Evolution*, 11(3):459–468, 1994. doi:
1072 10.1093/oxfordjournals.molbev.a040126.

1073 Harald O Letsch and Karl M Kjer. Potential pitfalls of modelling ribosomal rna data in phylogenetic
1074 tree reconstruction: Evidence from case studies in the metazoa. *BMC Evolutionary Biology*, 11
1075 (1):146, 2011. doi: 10.1186/1471-2148-11-146.

1076 Peter J. Lockhart, Michael A. Steel, Michael D. Hendy, and David Penny. Recovering evolutionary
1077 trees under a more realistic model of sequence evolution. *Molecular Biology and Evolution*, 11(4):
1078 605–612, 1994. doi: 10.1093/oxfordjournals.molbev.a040136.

1079 Nhan Ly-Trong, Suha Naser-Khdour, Robert Lanfear, and Bui Quang Minh. Alisim: A fast and
1080 versatile phylogenetic sequence simulator for the genomic era. *Molecular Biology and Evolution*,
1081 39(5):msac092, 2022. doi: 10.1093/molbev/msac092.

1082 Cassius Manuel, Enes Sakalli, Heiko A Schmidt, Carme Viñas, Arndt von Haeseler, and Christiane
1083 Elgert. When the past fades: Detecting phylogenetic signal with satute. *Molecular Biology and*
1084 *Evolution*, 42(5):msaf090, 2025. doi: 10.1093/molbev/msaf090.

1085 Bui Quang Minh, Heiko A Schmidt, Olga Chernomor, Dominik Schrempf, Michael D Woodhams,
1086 Arndt von Haeseler, and Robert Lanfear. Iq-tree 2: New models and efficient methods for
1087 phylogenetic inference in the genomic era. *Molecular Biology and Evolution*, 37(5):1530–1534,
1088 2020. doi: 10.1093/molbev/msaa015.

1089 David Penny, Bennet J. McComish, Michael A. Charleston, and Michael D. Hendy. Mathematical
1090 elegance with biochemical realism: the covarion model of molecular evolution. *Journal of Molecular*
1091 *Evolution*, 53(6):711–723, 2001. doi: 10.1007/s002390010258.

1092 M. A. Steel, P. J. Lockhart, and D. Penny. Confidence in evolutionary trees from biological sequence
1093 data. *Nature*, 364(6436):440–442, 1993. doi: 10.1038/364440a0.

1094 Torsten H. Struck, Maximilian P. Nesnidal, Günter Purschke, and Kenneth M. Halanych. Detecting
1095 possibly saturated positions in 18s and 28s sequences and their influence on phylogenetic recon-
1096 struction of annelida (lophotrochozoa). *Molecular Phylogenetics and Evolution*, 48(2):628–645,
1097 2008. doi: 10.1016/j.ympev.2008.05.015.

1098 Heroen Verbruggen, Frederik Leliaert, Christine A. Maggs, Satoshi Shimada, Tom Schils, Jim
1099 Provan, David Booth, Sue Murphy, Olivier De Clerck, Diane S. Littler, Mark M. Littler, and
1100 Eric Coppejans. Species boundaries and phylogenetic relationships within the green algal genus
1101 codium (bryopsidales) based on plastid dna sequences. *Molecular Phylogenetics and Evolution*, 44
1102 (1):240–254, 2007. doi: 10.1016/j.ympev.2007.01.009.

1103 Xuhua Xia and Philippe Lemey. Assessing substitution saturation with dambe. In *The Phylogenetic*
1104 *Handbook*, pages 615–630. Cambridge University Press, 2009. ISBN 9780521877107. doi: 10.1017/
1105 CBO9780511819049.022.

1106 Xuhua Xia, Zheng Xie, Marco Salemi, Lu Chen, and Yong Wang. An index of substitution
1107 saturation and its application. *Molecular Phylogenetics and Evolution*, 26(1):1–7, 2003. doi:
1108 10.1016/S1055-7903(02)00326-3.

1109 Ziheng Yang. Maximum likelihood phylogenetic estimation from dna sequences with variable rates
1110 over sites: Approximate methods. *Journal of Molecular Evolution*, 39(3):306–314, 1994. doi:
1111 10.1007/BF00160154.