hexagon:2611.01233v1

Stable Type Proportions under Partial Information in Branching Trees

Viktor Diakin
Viramelis Laboratory

27 November 2026

Abstract

We study the type proportions of finite ancestor-closed sets obtained by sequential exploration of multitype branching trees. When a vertex enters the frontier, the explorer sees its type and a local signal about its own offspring. Let G be the conditional mean offspring matrix on type–signal states and L the projection onto physical types. Under bounded, everywhere nonempty offspring and irreducible supercritical G, the almost surely attainable state limits along stopping-time exhaustions are exactly the probability rows x with $xG \geq x$. The physical limits stable under adjoining all children are exactly those xL for which $xGL = (xG\mathbf{1})xL$. In a binary model the stable interval has radius equal to the mean absolute deviation of the posterior offspring probability from a median. Every point of this interval is attained along all exploration steps by an exhaustive priority rule with a sparse oldest-first schedule. An appendix records the corresponding deterministic cutoff results.

1 Introduction

An infinite branching tree can be counted by taking larger and larger finite sets. The order of exploration can affect the resulting type proportions. We ask which limits are attainable and which remain unchanged asymptotically after one round of adding children. We distinguish the physical types being counted from the information available when a frontier vertex is selected.

The frontier recursion is also the service-completion recursion of a branching bandit without immigration [6, 1]. The attainable vectors in [1] are stationary mean queue lengths in a subcritical system. In our supercritical setting the frontier remains visible on the linear scale, and the attainable vectors are empirical proportions of processed vertices. For strict targets we invert the draw-frequency formula of Janson [3, Theorem 3.21] by choosing activities from the desired limit. The proof at the boundary uses adaptive stopping times. The binary optimization is the upper-half expected-shortfall calculation [4], with the additional requirement that the extremizers be implemented by exhaustive explorations. Signal degradation is used in the sense of the comparison of experiments [2]. These antecedents delimit the contribution. We make no claim of established priority for the cutoff classification.

2 Model and filtration

Let $\mathcal{I}$ be a finite set of physical types and fix the root type. A vertex of type i has an ordered list K of child types of length between 1 and $B < \infty$. Before it is processed, its visible mark is a finite-valued signal S drawn by a fixed kernel given (i, K). For each type i, write $p_i(s) = \mathbb{P}_i(S = s)$ and $\mu_{i,s}$ for

the conditional law of K given $S = s$ when $p_i(s) > 0$. We use the state set $\mathcal{Z} = \{(i, s) : p_i(s) > 0\}$. The root is initially on the frontier with its state known. A step processes one frontier vertex, exposes its offspring list, and adds the children to the frontier with their states. The one-step expansion F^+ of a finite ancestor-closed set F is the union of F and all children of its vertices.

For a concrete probability space, give each potential vertex address u and each physical type i an independent pair $(K_{u,i}, S_{u,i})$ of the specified joint law. Addresses are finite words over $\{1, \dots, B\}$, with the empty word for the root. A child address and its type are determined by the parent's list. At address u use the pair for its actual type i. All other pairs are unused. An independent sequence U_n of uniform random variables allows randomized choices. Let $\mathcal{F}_n$ be the sigma-field generated by the initial state and, through step n, the uniforms used, selected addresses, revealed lists and newly discovered states. Put $\mathcal{H}_{n+1} = \mathcal{F}_n \vee \sigma(U_{n+1})$. An admissible rule selects a frontier address V_{n+1} measurably with respect to $\mathcal{H}_{n+1}$. The history then grows from $\mathcal{H}_{n+1}$ to $\mathcal{F}_{n+1}$ by revealing its list and the children's states. Extra independent randomization can be generated from the uniforms.

Lemma 2.1 (Conditional frontier law)**.** *Given $\mathcal{F}_n$, the hidden lists at the frontier are conditionally independent, and a vertex of observed state (i, s) has list law $\mu_{i,s}$. The same holds given $\mathcal{H}_{n+1}$.*

Proof. The assertion means a product identity for bounded test functions on the lists of any given finite collection of addresses on the visible frontier. It suffices to check the identity on events specifying the frontier and its states, since the possible address and state configurations form a countable set. At time zero the identity is conditioning on the root signal. Assume it holds at time n. Adding the independent U_{n+1} preserves it. On each event specifying the selected address, that address is $\mathcal{H}_{n+1}$-measurable. Thus revealing its list leaves the conditional product law at every other old frontier address unchanged. The new child addresses and types are determined by this list. On each event fixing that list and the parent address, none of the random pairs at the new addresses has been inspected. Indeed all observations so far concern discovered addresses or the new children's parent and its ancestors. Independence of the address-indexed pairs gives the product of the joint laws at the new children, independent of the hidden lists at the retained frontier. Conditioning on the individual new signals updates each new factor to $\mu_{i,s}$, and leaves the old factors unchanged. This gives the product identity at time $n + 1$. □

For $z = (i, s)$, $w = (j, t)$, define

$$G_{zw} = \mathbb{E}_i[\#\{\text{children of type } j\} \mid S = s]\, p_j(t), \qquad L_{zj} = \mathbf{1}_{\{i=j\}}.$$

Throughout the stochastic results assume G is irreducible and $\rho(G) > 1$. Vectors of counts and proportions are rows. Write F_n for the first n processed vertices, c_n for their state counts and b_n for the current frontier counts. Let $x_n = c_n/n$ for $n \geq 1$. If Y_{n+1} is the row of child state counts and Z_{n+1} the selected state, the lemma gives

$$\mathbb{E}[Y_{n+1} \mid \mathcal{H}_{n+1}] = G_{Z_{n+1},*}.$$

In particular $Y_{n+1} - G_{Z_{n+1},*}$ is a bounded vector martingale difference for the filtration $(\mathcal{F}_n)$, by the tower property. Coordinatewise Azuma–Hoeffding followed by Borel–Cantelli yields $\sum_{k=1}^{n}(Y_k - G_{Z_k,*}) = o(n)$ almost surely for each fixed rule. Since F_n^+ consists of the root and all children of processed vertices, disjointly by parent,

$$b_n = e_{\text{root}} + c_n(G - I) + o(n), \qquad \frac{c(F_n^+)}{n} = x_n G + o(1). \tag{1}$$

Here $c(F_n^+)$ counts observed states. Its physical proportions are obtained by multiplying by L and normalizing by $|F_n^+|$.

3 Attainability and stability

A stopping-time exhaustion consists of F_{τ_k} for strictly increasing, almost surely finite stopping times of $(\mathcal{F}_n)$, with union the whole tree almost surely. Write $p(F)$ for its physical type proportions. Stability means $\|p(F^+_{\tau_k}) - p(F_{\tau_k})\|_1 \to 0$ almost surely.

Theorem 3.1. *The set of deterministic state targets attainable almost surely by a stopping-time exhaustion is*

$$\mathcal{P}_G = \{x \geq 0 : x\mathbf{1} = 1,\ xG \geq x\}.$$

The set of deterministic physical targets attainable almost surely by a stable stopping-time exhaustion is

$$\mathcal{S}_{G,L} = \{xL : x \in \mathcal{P}_G,\ xGL = (xG\mathbf{1})xL\}.$$

Each target may use its own admissible rule and stopping times. For each fixed rule the necessary constraints hold almost surely along every subsequence for which the stated limits exist.

Proof. The simplex is compact. From $b_n \geq 0$ and (1), any limit x of a subsequence of x_n is in $\mathcal{P}_G$. On that subsequence the physical proportions before and after expansion tend to xL and $xGL/(xG\mathbf{1})$; the denominator is at least one. For necessity for a physical limit, extract a convergent state subsequence along the given snapshots. On each path satisfying (1), the physical limit and stability force the displayed equality. Thus a deterministic physical target attained almost surely belongs to $\mathcal{S}_{G,L}$, even if the state subsequence limit used to obtain a witness is random.

For sufficiency first take $x > 0$ with $b = x(G - I) > 0$ coordinatewise. Assign each vertex of state z the activity $a_z = x_z/b_z$, and select a frontier vertex with probability proportional to its activity. Conditional on its state z the net replacement is $Y - e_z$. The conditional frontier law makes this an urn whose count transition depends only on the counts with these fixed activities. Its row-form mean generator is

$$H = \operatorname{diag}(a)(G - I), \qquad bH = b, \qquad \sum_z a_z b_z = 1.$$

In Janson's column convention the matrix is $A = H^{\mathsf{T}}$, with column z equal to a_z times the expected increment column. It is irreducible with nonnegative off-diagonal entries. A sufficiently large diagonal shift and Perron–Frobenius show that 1 is its simple leading eigenvalue, with the positive right eigenvector b^{T}. We verify the six hypotheses of [3, Section 2] individually.

(A1) A selected ball is removed and the children are added. The replacement entry plus the matching Kronecker delta is a nonnegative integer, so the urn is tenable.

(A2) Every replacement coordinate has finite second moment because the degree is bounded.

(A3) The leading real eigenvalue is $1 > 0$.

(A4) It is simple by irreducibility and the diagonal-shift argument.

(A5) Every state is dominating since the activity of every state is positive and the reachability graph is irreducible. The initial frontier is nonempty, so it contains a ball of a dominating state.

(A6) The unique communicating class is the dominating class and contains the leading eigenvalue.

Each step inserts at least one ball. The frontier never empties, and with one class essential nonextinction has probability one. The strong-law assertion of [3, Theorem 3.21] now gives $c_{n,z}/n \to a_z b_z = x_z$ almost surely. The normalization there is $a^{\mathsf{T}} b^{\mathsf{T}} = 1$ when a is viewed as an activity column. The same conclusion holds for future draws started after a finite stopping time, conditionally on its history, if the activities are thereafter held fixed. This follows from the frontier law by conditioning on each possible time and the current history, and applying the urn law to the known count vector. Including any already processed vertices changes the proportions by a vanishing amount.

For general $x \in \mathcal{P}_G$, choose the positive normalized left Perron row u of G, and put

$$x^{(k)} = (1 - 1/k)x + u/k, \qquad x^{(k)}(G - I) \geq (\rho(G) - 1)u/k > 0.$$

At stage k, first finish processing all vertices of depth at most k by always taking a shallowest available vertex up to that depth. This is a finite procedure and exhausts those levels by ancestor closure. Then apply the fixed activities for $x^{(k)}$ until the first subsequent step for which $\|x_n - x^{(k)}\|_1 < 1/k$; call this time τ_k. The test uses only the revealed history, so τ_k is a stopping time and is strictly later than τ_{k-1}. If the fixed-activity phase were continued forever, the conditional urn law would give $x_n \to x^{(k)}$. The first qualifying step is consequently finite almost surely given the history at the start of the phase. Induction and a countable intersection give finite stages almost surely. The forced levels ensure exhaustion and $\|x_{\tau_k} - x\|_1 \leq 1/k + \|u - x\|_1/k \to 0$. This argument also covers zero coordinates in x or in $x(G - I)$. If the defining equality of $\mathcal{S}_{G,L}$ holds, (1) gives stability of these snapshots. □

If $G\mathbf{1} = d\mathbf{1}$, then $xG\mathbf{1} = d$ and the set of physical limits is the linear image of a polytope. With a constant signal L is the identity and the unique stable limit is the normalized left Perron row. More generally, if a coarser signal is simulated from a finer one and the physical type by an independent, vertexwise kernel, the coarse rule can be run using the simulated observations [2]. Its physical tree and selected sets have the same law. Consequently its stable limit set is included in that of the finer signal.

4 Binary example and convergence along all steps

Each vertex independently has two children, either both of type A or both of type B, with probabilities $1/2$ each. This choice is independent of the parent's physical type. The same finite signal kernel is used at every vertex. For a signal s of positive mass set

$$\nu_s = \mathbb{P}(S = s), \qquad r_s = \mathbb{P}(AA \mid S = s).$$

Let $R = r_S$, so $\mathbb{E}[R] = 1/2$. For every pair of states there is a two-generation path of positive mean: a parent can produce at least one child of some type, that child can have any positive-mass signal, and averaging its offspring possibilities permits either type at the next generation. Thus G^2 is positive, and $G\mathbf{1} = 2\mathbf{1}$.

Theorem 4.1. *The stable fractions of type A form the closed interval*

$$[1/2 - D, 1/2 + D], \qquad D = \mathbb{E}|R - m|$$

for any median m of R. Each point is attained almost surely by an exhaustive rule for which both $p(F_n)$ and $p(F_n^+)$ converge along all steps. In particular, if the signal reveals the offspring type with probability α independently of that type and is otherwise an erasure, the interval is $[(1-\alpha)/2, (1+\alpha)/2]$.

Proof. For an observed-state probability row x let $q_s = \sum_i x_{i,s}$ and $a = \sum_s q_s r_s$. Then $xG = 2(a, 1-a) \otimes \nu$. Stability forces $xL = (a, 1-a)$, and $xG \geq x$ forces $q_s \leq 2\nu_s$. Conversely any probability row q with $q_s \leq 2\nu_s$ has a feasible lift $x = (a, 1-a) \otimes q$ with that same a, including the cases $a = 0, 1$. Thus the extrema of a are the extrema of a linear functional on the capped probability simplex.

Write $q_s = \nu_s(1+h_s)$, with $|h_s| \leq 1$ and $\sum_s \nu_s h_s = 0$. The deviation from $1/2$ is $\sum_s \nu_s h_s (r_s - m)$, which is bounded in absolute value by D. To attain the upper bound set $h_s = 1$ above m, $h_s = -1$ below m. Set $p_+ = \mathbb{P}(R > m)$, $p_- = \mathbb{P}(R < m)$, $p_0 = \mathbb{P}(R = m)$. The median inequalities $p_+, p_- \leq 1/2$ give $|p_+ - p_-| \leq p_0$. If $p_0 > 0$ use the constant $(p_- - p_+)/p_0$ on the median. If $p_0 = 0$, the two strict tails have mass $1/2$ and no adjustment is needed. Negating h attains the lower bound. Convexity gives all intermediate values among snapshots. This also shows that D is the same for all medians. For the erasure channel $1/2$ is a median, and $D = \alpha/2$.

We next construct an all-step exploration for each endpoint. Fix decreasing priority in r_s for the upper endpoint, with ties broken by a fixed signal order. At times $n = k^2$, process the oldest discovered frontier vertex. At all other times process the oldest vertex with the highest available priority. Fix a discovery order for siblings. Within each signal, processing is in discovery order under both clauses. Every discovered vertex has finitely many predecessors, and infinitely many square steps give it eventual service. Induction on depth shows exhaustion.

Let U be a prefix of the priority order, T_n the total number of new vertices with signals in U discovered through step n, and C_n the number of processed vertices with such signals. The conditional mean of each new count given the pre-choice history and choice is $t = 2\nu(U)$. It does not depend on the selected signal, since child signals have law ν for either physical type. Bounded martingale differences give $T_n = tn + o(n)$ and $\max_{j \leq n} |T_j - tj| = o(n)$ almost surely. The frontier balance gives $C_n \leq b_0(U) + T_n$ and $C_n \leq n$. If $j \leq n$ is the last nonsquare step selecting outside U, the frontier for U is empty before j, hence $C_{j-1} = b_0(U) + T_{j-1}$. After j, every nonsquare step selects in U, giving

$$C_n \geq b_0(U) + T_{j-1} + n - j - \lfloor \sqrt{n} \rfloor \geq \min(t, 1)n - o(n).$$

If no such j exists, $C_n \geq n - \lfloor \sqrt{n} \rfloor$. We conclude $C_n/n \to \min(1, 2\nu(U))$ for every prefix. Differences of consecutive prefix limits give the limiting processed signal frequencies q_s. These fill the caps $2\nu_s$ in descending priority order up to total mass one. They maximize $\sum_s q_s r_s$: moving any available mass from lower to higher r_s increases the objective, and ties have equal value.

Put $a = \sum_s q_s r_s$ for these limiting frequencies. For each fixed s, the conditional expected numbers of discovered vertices with that signal and of discovered A vertices with that signal, at a step selecting signal v, are $2\nu_s$ and $2\nu_s r_v$, respectively. The bounded martingale law gives discovered counts $2\nu_s n + o(n)$ and $2\nu_s a n + o(n)$ almost surely. The root can change either count by at most one. The fraction of A in the discovery list for signal s therefore tends to a. The processed vertices in that list form a prefix. If $q_s > 0$, the prefix tends to infinite length and has the same limiting fraction. If $q_s = 0$, its contribution divided by n tends to zero, even if there are infinitely many such draws. The full processed fraction therefore tends to a, and the fraction among all discovered vertices, which constitute F_n^+, tends to a as well. Reversing the signal priority attains the lower endpoint.

Finally retain each original signal with independent probability β and otherwise substitute a fresh erasure symbol. The posterior law is the mixture of the original law with mass $1 - \beta$ at $1/2$. For this law $D(\beta) = \min_{m \in [0,1]} \mathbb{E}_\beta |R - m|$. To see the minimization, the function $m \mapsto \mathbb{E}_\beta |R - m|$ is convex and its left and right slopes at a median straddle zero. It follows also that $|D(\beta) - D(\beta')| \leq |\beta - \beta'|$, by the uniform bound on the change of the objective. The radius is zero at $\beta = 0$ and

equals D at $\beta = 1$. For any desired deviation in $[0, D]$, the intermediate value theorem supplies a β; the endpoint rule with the chosen sign then attains the desired limit along all steps. Necessary bounds along all steps follow already from Theorem 3.1 and the capped simplex. □

5 Relation to stationary branching queues

The bookkeeping identity preceding (1) is the recursion in [1, Eq. (3)] without replenishment from the external class. In [1, Assumption A(a) and Lemma 3.1], the mean matrix N for ordinary customers satisfies $\rho(N) < 1$. Under stationarity the expected increment vanishes; the service rates solve $r = rN + r_0 u$, with u the immigration row. Their achievable-region theorems concern stationary mean queue lengths and use second-moment balances [1, Theorems 3.2, 4.2 and 5.4]. A stationary integrable frontier satisfies $b_n/n \to 0$ almost surely by tail summability and Borel–Cantelli. For our G, a positive right Perron column v gives instead

$$(b_n/n)v = (\rho(G) - 1)x_n v + o(1), \qquad \liminf_n (b_n/n)v \geq (\rho(G) - 1) \min_z v_z > 0.$$

A nonnegative nonzero linear recoding Q preserving the mean offspring via $NQ = QG$ also preserves the eigenvalue $\rho(G)$, since $Qv \neq 0$. Thus this class of recodings cannot yield a subcritical N. The divergence of stationary means as an auxiliary matrix approaches criticality in [1, Lemma 3.6 and Theorem 3.7] concerns a family of stationary systems. It supplies no within-tree stopping-time construction in a fixed supercritical system. This comparison concerns the published stationary results and mean-preserving recodings.

6 Computational checks and scope

The separate reproducibility script checks 4,646 deterministic quota constructions (256,355 individual additions), eight exact rational pole orders, the binary median formula for 504 finite channels and 101 erasure probabilities, and 200 inverse-activity identities. Seven seeded priority runs of 200,000 steps include both full-information endpoints, partial information, and a noisy channel. The checks test the stated formulae and implementations. The almost-sure and all-step assertions are established by the proofs above. A general boundary target in Theorem 3.1 is attained along stopping-time snapshots; the all-step conclusion here is for Theorem 4.1. The paper concerns type proportions in branching trees.

AI assistance. Viktor Diakin initiated and directed the project. OpenAI language models assisted with mathematical development, proof checking, literature comparison, verification code, and English exposition. Additional model-based critiques informed the revision. The checks reported here are mathematical arguments and computational tests, with no claim of human peer review.

A Deterministic offspring and depth balls

In this appendix every vertex of physical type i has exactly M_{ij} children of type j. Restrict the nonnegative integer matrix M to types reachable from the root r. Assume the tree is infinite; then the reachable graph has a directed cycle and $\rho(M) \geq 1$. Terminal vertices are allowed in this appendix. An exhaustion here consists of nested finite ancestor-closed sets with union the whole

tree, with unrestricted finite additions between snapshots. With $c(F)$ counting physical types, the exact identities are

$$c(F^+) = c(F)M + e_r, \qquad c(F^+ \setminus F) = c(F)(M - I) + e_r. \tag{2}$$

Theorem A.1. *The attainable limits are $P = \{x \geq 0 : x\mathbf{1} = 1,\ xM \geq x\}$. Every nonempty compact subset of P is the exact cluster set of some exhaustion. An exhaustion is stable if and only if all its cluster points lie in $E = \{x \in P : xM = (xM\mathbf{1})x\}$. In particular every nonempty compact subset of E occurs as the cluster set of a stable exhaustion.*

Proof. Necessity of the balance inequality follows from (2) and nonnegativity of the frontier. For sufficiency, fix $x \in P$ and set $J_+ = \{i : x_i > 0\}$. From any finite ancestor-closed starting set, add paths to include a representative of each type in J_+. For a sufficiently large integer T, impose quotas $\lceil Tx_i \rceil$ for $i \in J_+$, all at least the current counts. Keep the counts of other types fixed. Repeatedly add an available frontier vertex of a type below its quota.

If the process stops short, put $t = \min_{i \in J_+} c_i/x_i > 0$ and $J = \{i \in J_+ : c_i = tx_i\}$, with c the current counts. A short quota gives $t < T$, since an integer less than $\lceil Tx_i \rceil$ is less than Tx_i. Every type in J is below quota, so its frontier count is zero. For $i \in J$,

$$c_i = (cM)_i + (e_r)_i \geq t(xM)_i + (e_r)_i \geq tx_i = c_i.$$

The equalities imply $(e_r)_i = 0$ and no edge from any currently represented type outside J into i: for each such type k, $c_k - tx_k > 0$. Yet a path from the root to any represented vertex of J has a first edge entering J, a contradiction. Hence all quotas are filled.

At each snapshot first include the previous snapshot, the depth ball for the current index, and the representative paths. Then choose T so large that the starting counts are negligible relative to T, and fill the quotas. Rounding contributes at most the number of types, and the resulting proportions approximate x arbitrarily well. This proves attainment by an exhaustion even when some coordinates vanish. For a compact cluster set choose successive finite $1/k$-nets in that set, concatenate them, and approximate the successive target rows with errors tending to zero, making the same depth-ball inclusion each time. The resulting cluster set is exactly the prescribed set.

On any subsequence along which the proportions tend to $x \in P$, equation (2) gives expanded proportions tending to $xM/(xM\mathbf{1})$, with denominator at least one. Hence the stability residual tends to zero on that subsequence exactly when $x \in E$. Compactness gives the full-sequence equivalence. □

The following classification is the row-vector form of the Frobenius–Victory theorem [5, Theorem 2.1]. Write r_C for the spectral radius of the diagonal block of a strongly connected class C of the graph $i \to j$ when $M_{ij} > 0$. A class C generates a normalized nonnegative left eigenvector in E precisely when $r_C \geq 1$ and the radius of every distinct class reachable from C is strictly less than r_C. Its generator is supported on C and its descendants. For each eigenvalue $\lambda \geq 1$, the corresponding part of E is the convex hull of the generators with $r_C = \lambda$.

For completeness, in topological block order the equation for a block D is $x_D(\lambda I - M_{DD}) = b_D$, where $b_D = \sum_{C \neq D} x_C M_{CD} \geq 0$. Multiplying by a positive right Perron vector of M_{DD}, also available for a singleton zero block, shows that $b_D \neq 0$ forces $\lambda > r_D$. In that case the resolvent series gives a unique positive solution in the block. With zero input the block is either zero or a Perron row of eigenvalue $r_D = \lambda$. Each positive block with zero input therefore generates positive values downstream and must have radius strictly larger than each downstream block. Conversely one such class produces a solution by the block recursion, and superposition gives all solutions. Normalizing gives the claimed hulls. Distinct eligible classes of the same radius are mutually unreachable, and

different eigenvalues give distinct normalized rows. Therefore all stable exhaustions converge to the same limit exactly when there is a single eligible class.

Theorem A.2. *For the depth ball B_n, put $S(z) = e_r(I - zM)^{-1}/(1-z)$, $R = 1/\rho(M)$, and let m be the pole order of S at R. The order of a vector function is the maximum of its coordinate orders after cancellation. The balls are stable exactly when every other pole on $|z| = R$ has order strictly less than m. In that case their proportions converge to the normalization of the leading coefficient at R.*

Proof. Write $s_n = c(B_n) = \sum_{j=0}^{n} e_r M^j$, so $S(z) = \sum_{n\geq 0} s_n z^n$. A block of radius $\rho = \rho(M)$ is reached from the root, and its positive right Perron vector gives a lower bound $s_n\mathbf{1} \geq c\rho^n$ for all large n and some $c > 0$. A Jordan bound gives an upper bound by a polynomial times ρ^n. Thus the radius is R, and $S\mathbf{1}$ diverges as real $z \uparrow R$, so $m \geq 1$. For each coordinate S_j, $|S_j(t\zeta)| \leq S_j(tR)$ for $|\zeta| = R$, $0 < t < 1$, by nonnegative coefficients. No pole on the circle can have larger order than m.

By the block Perron–Frobenius theorem the ratios ζ/R for peripheral poles are roots of unity. Choose a common period h. Partial fractions yield, on residue class a modulo h,

$$\frac{s_n}{\rho^n \binom{n+m-1}{m-1}} \longrightarrow W_a = \sum_{\substack{|\zeta|=R \\ \operatorname{ord}_\zeta S = m}} v_\zeta (\zeta/R)^{-a}, \qquad v_\zeta = \lim_{z\to\zeta} (1 - z/\zeta)^m S(z).$$

The mean of the W_a is $v_R \neq 0$, and $s_n \geq s_{n-1}$ gives $W_a \geq W_{a-1}/\rho$. Consequently all $W_a\mathbf{1} > 0$. The scale displayed above tends to infinity; when $\rho = 1$ this follows from $s_n\mathbf{1} \geq n+1$ and boundedness of the scaled sequence. Since $B_n^+ = B_{n+1}$, stability is equivalent to equality of all normalized phase rows. If the common probability row is x, write $W_a = t_a x$ with $t_a > 0$. The identity $s_{n+1} = s_n M + e_r$ gives $\rho W_{a+1} = W_a M$. The ratios t_{a+1}/t_a are constant and have product one, so all W_a are equal. Uniqueness of the finite Fourier coefficients is then equivalent to vanishing of the other order-m terms. Conversely, absence of those terms gives the single normalized limit from v_R. The recursion shows that this limit is a left eigenvector for ρ. □

For $\rho = 1$, the nonnegative-coefficient argument applied first to $e_r(I - zM)^{-1}$ bounds all peripheral pole orders by the one at 1. The factor $1/(1-z)$ then makes the positive pole strictly dominant. For irreducible M and $\rho > 1$, the depth balls are stable exactly when M is primitive. Primitivity gives a single peripheral pole. If the period is greater than one, hypothetical stability and simplicity of the positive Perron root would give $s_n/\rho^n \to W > 0$, hence $e_r M^n/\rho^n \to (1 - 1/\rho)W > 0$. Each generation is confined to one cyclic class, contradicting the latter limit. As an example, $\left(\begin{smallmatrix} 0 & 2 \\ 2 & 0 \end{smallmatrix}\right)$ with root of the first type gives two depth-ball limits, $(2/3, 1/3)$ and $(1/3, 2/3)$. Its stable exhaustion limit is uniquely $(1/2, 1/2)$ by Theorem A.1.

Terminal types. In block form write $M = \left(\begin{smallmatrix} A & B \\ 0 & 0 \end{smallmatrix}\right)$, with terminal types in the second block, and $x = (u, v)$. The conditions for attainable limits are $u \geq 0,\ uA \geq u,\ 0 \leq v \leq uB$, and $u\mathbf{1} + v\mathbf{1} = 1$. Let $b_* = \max\{wB\mathbf{1} : w \geq 0,\ w\mathbf{1} = 1,\ wA \geq w\}$. This maximum exists since the feasible set is compact and $\rho(A) \geq 1$. With terminal mass $t = v\mathbf{1}$, the substitution $u = (1-t)w$ gives $t \leq (1-t)b_*$. The upper endpoint $b_*/(1 + b_*)$ follows by taking a maximizing w and $v = uB$; the lower endpoint zero follows by taking $v = 0$. Convexity and Theorem A.1 give the full interval. Stable limits satisfy $uA = \lambda u,\ v = uB/\lambda$, $\lambda \geq 1$. The block classification shows that a fixed terminal type has positive fraction in every stable limit exactly when it is reachable from every class generating a stable eigenvector. Finally, if $p(F_n) \to x \in E$, equation (2) yields $|F_n^+ \setminus F_n|/|F_n| \to xM\mathbf{1} - 1$. A vanishing relative frontier always suffices for stability because the residual is at most $2|F_n^+ \setminus F_n|/|F_n^+|$. If

every row in E has eigenvalue one, compactness and Theorem A.1 show that this condition is also necessary. For example, $M = \left(\begin{smallmatrix} 2 & 1 \\ 0 & 0 \end{smallmatrix}\right)$ gives terminal fractions $[0, 1/2]$ among all exhaustions and the unique stable terminal fraction $1/3$.

References

[1] D. Bertsimas, I. Ch. Paschalidis, and J. N. Tsitsiklis, Branching bandits and Klimov's problem: achievable region and side constraints, *IEEE Trans. Automat. Control* **40** (1995), 2063–2075. Author-hosted text.

[2] D. Blackwell, Equivalent comparisons of experiments, *Ann. Math. Statist.* **24** (1953), 265–272. doi:10.1214/aoms/1177729032.

[3] S. Janson, Functional limit theorems for multitype branching processes and generalized Pólya urns, *Stochastic Process. Appl.* **110** (2004), 177–245. doi:10.1016/j.spa.2003.12.002.

[4] R. T. Rockafellar and S. Uryasev, Optimization of conditional value-at-risk, *J. Risk* **2** (2000), 21–41. doi:10.21314/JOR.2000.038.

[5] B.-S. Tam and H. Schneider, On the core of a cone-preserving map, *Trans. Amer. Math. Soc.* **343** (1994), 479–524. doi:10.1090/S0002-9947-1994-1242787-6.

[6] G. Weiss, Branching bandit processes, *Probab. Engrg. Inform. Sci.* **2** (1988), 269–278. doi:10.1017/S0269964800000826.