Quasistationary Distributions: Existence, Uniqueness and Characterization

Iddo Ben-Ari (University of Connecticut), joint with Ningwei Jiang

University of Colorado, Colorado Springs, 2023-03-16

Outline

Motivating Example: Birth & Death Chain

Consider the discrete-time Birth & Death chain $(X_t : t \in \mathbb{Z}_+)$ on the states $\mathbb{Z}_+$ with $q \in (\frac{1}{2}, 1)$.

- A unique stationary distribution π, a distribution invariant under the dynamics of the chain. Moreover,
- For any initial distribution μ,

$$(*)\ P_\mu(X_t \in \cdot\ ) \underset{t\to\infty}{\to} \pi,\ \pi \sim \text{Geom}(1 - \frac{1-q}{q}) - 1.$$

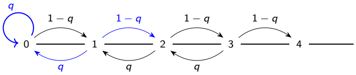

Figure: Birth & Death

Motivating Example: Birth & Death Chain

Consider the discrete-time Birth & Death chain $(X_t : t \in \mathbb{Z}_+)$ on the states $\mathbb{Z}_+$ with $q \in (\frac{1}{2}, 1)$.

- A unique stationary distribution π, a distribution invariant under the dynamics of the chain. Moreover,
- For any initial distribution μ,

$$(*)\ P_\mu(X_t \in \cdot\ ) \underset{t\to\infty}{\to} \pi,\ \pi \sim \text{Geom}(1 - \frac{1-q}{q}) - 1.$$

Now absorb ("kill") the process at 0, setting $p(0,0) = 1$.

- $(*)$ still holds, but with a trivial stationary distribution $\pi = \delta_0$.
- How would the process behave, conditioned on not being absorbed? Equivalently, is there an conditional version of $(*)$,

$$P_\mu(X_t \in \cdot \mid \mathbf{X} \text{ has not hit } 0 \text{ by time } t) \underset{t\to\infty}{\to} ?$$

- Quasistationary distributions (QSDs) are probability distributions appearing as such limits.
- "What would a biological system that has survived for a long time would look like?"

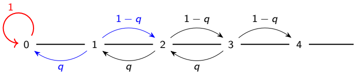

Figure: Birth & Death killed at 0

Definitions

Assumption 1

Let $\mathbf{X} = (X_t : t \in \mathbb{Z}_+)$ *be MC on state space* $\{\mathbf{0}\} \cup S$ *where* $S = \{1, \ldots, N\}$ *or* $S = \mathbb{N}$*, with transition function p satisfying*

1. *The state* $\mathbf{0}$ *is a unique absorbing state:* $p(0, 0) = 1$.
2. *The restriction of p to nonabsorbing states(*$= S$*)is irreducible.*

Let

$$\zeta = \inf\{t \in \mathbb{N} : X_t = 0\},$$

the absorption time.

3. $P_x(\zeta < \infty) = 1$ *for some (equivalently all)* $x \in S$.
4. $E_x[\beta^\zeta] < \infty$ *for some (equivalently all)* $x \in S$ *and* $\beta > 1$.

Definition 1 (QSD)

A probability measure ν *on S is a Quasistationary Distribution if*

$$P_\nu(X_t \in \cdot \mid \zeta > t) = \nu(\cdot)$$

for all $t \in \mathbb{Z}_+$.

First Observations

Proposition 1

- *(Necessary condition) If ν is a QSD, then under P_ν, ζ has a geometric distribution with parameter $1-\lambda \in (0,1)$:*

$$P_\nu(\zeta > t) = \lambda^t, \ t \in \mathbb{Z}_+. \tag{1}$$

  - *λ is called the survival probability for ν.*
- *(Eigenvector) A probability measure ν on S is a QSD with survival probability λ if and only if*

$$\nu p = \lambda \nu. \tag{2}$$

  Equivalently, ν satisfies the non-linear eigenvalue equation:

$$\nu p = ((\nu p)\mathbf{1})\nu.$$

Proposition 2 (Quasi-limiting $\Rightarrow$ QSD)

If ν is a probability measure on S satisfying

$$\lim_{t\to\infty} P_\mu(X_t \in \cdot \mid \zeta > t) = \nu \text{ for some } \mu,$$

then ν is a QSD .

Example: RW on an Interval

Example 1 (RW on an Interval)

Let $N \geq 2$ be an integer, and consider the following transition function:

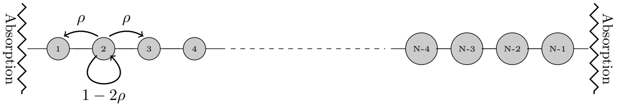

Figure: *RW absorbed outside an interval*

Solving (2) *yields a unique QSD ν_N, with a survival probability λ_N:*

$$\begin{cases} \nu_N(x) = C_N \sin(\frac{x}{N}\pi) & (C_N = \tan\frac{\pi}{2N}); \\ \lambda_N(\rho) = \frac{\rho}{2}\cos\frac{\pi}{N} + (1-2\rho) \end{cases} \tag{3}$$

Example 2 (Voter on a Cycle)

For $N \geq 2$, consider the N-cycle $\mathbb{Z}_N$. Assign each vertex an opinion "yes" or "no". At each unit of time, uniformly sample a vertex and a random neighbor (CW or CCW), and assign the neighbor's opinion to the chosen vertex.

- *Absorbing states are consensus states: all "yes" and all "no".*
- *Consensus is eventually reached.*

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.

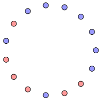

Figure: Initial opinion assignment

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of "yes" and "no".

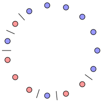

Figure: Interfaces

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of “yes” and “no”.
3. In terms of interfaces, each step either:

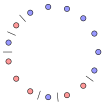

Figure: Interfaces

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of "yes" and "no".
3. In terms of interfaces, each step either:
   - Vertex & Neighbor in same cluster ⇒ no movement of interface.

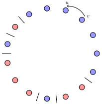

Figure: None of interface move

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of “yes” and “no”.
3. In terms of interfaces, each step either:
   - Vertex & Neighbor in same cluster $\Rightarrow$ no movement of interface.
   - Vertex & Neighbor on two sides of an interface $\Rightarrow$ an interface moves in one direction, with equal probability to each direction.

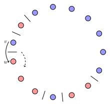

Figure: Interface moves

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of "yes" and "no".
3. In terms of interfaces, each step either:
   - Vertex & Neighbor in same cluster ⇒ no movement of interface.
   - Vertex & Neighbor on two sides of an interface ⇒ an interface moves in one direction, with equal probability to each direction.

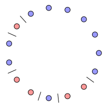

Figure: Interface moves, completed

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of "yes" and "no".
3. In terms of interfaces, each step either:
   - Vertex & Neighbor in same cluster ⇒ no movement of interface.
   - Vertex & Neighbor on two sides of an interface ⇒ an interface moves in one direction, with equal probability to each direction.
   - If interfaces meet, they are both eliminated.

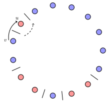

Figure: Interfaces meet and eliminated

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of “yes” and “no”.
3. In terms of interfaces, each step either:
   - Vertex & Neighbor in same cluster ⇒ no movement of interface.
   - Vertex & Neighbor on two sides of an interface ⇒ an interface moves in one direction, with equal probability to each direction.
   - If interfaces meet, they are both eliminated.

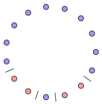

Figure: Interfaces cancel each other, completed

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of "yes" and "no".
3. In terms of interfaces, each step either:
   - Vertex & Neighbor in same cluster ⇒ no movement of interface.
   - Vertex & Neighbor on two sides of an interface ⇒ an interface moves in one direction, with equal probability to each direction.
   - If interfaces meet, they are both eliminated.
4. Eventually, the system has two interfaces ⇒ Looking for a QSD supported on states with two clusters.

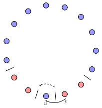

Figure: Down to two interfaces

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of "yes" and "no".
3. In terms of interfaces, each step either:
   - Vertex & Neighbor in same cluster ⇒ no movement of interface.
   - Vertex & Neighbor on two sides of an interface ⇒ an interface moves in one direction, with equal probability to each direction.
   - If interfaces meet, they are both eliminated.
4. Eventually, the system has two interfaces ⇒ Looking for a QSD supported on states with two clusters.

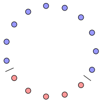

Figure: Down to two interfaces, completed

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of "yes" and "no".
3. In terms of interfaces, each step either:
   - Vertex & Neighbor in same cluster $\Rightarrow$ no movement of interface.
   - Vertex & Neighbor on two sides of an interface $\Rightarrow$ an interface moves in one direction, with equal probability to each direction.
   - If interfaces meet, they are both eliminated.
4. Eventually, the system has two interfaces $\Rightarrow$ Looking for a QSD supported on states with two clusters.

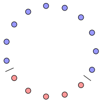

Figure: Down to two interfaces, completed

5. Per Step 3, the size of the remaining "yes" cluster performs a symmetric RW from Example 1, with $\rho = \frac{1}{N}$.

Example: Voter on the Cycle

Evolution

1. Some non-consensus initial opinion assignment.
2. Each non-absorbing state has an even number of interfaces between clusters of "yes" and "no".
3. In terms of interfaces, each step either:
   - Vertex & Neighbor in same cluster $\Rightarrow$ no movement of interface.
   - Vertex & Neighbor on two sides of an interface $\Rightarrow$ an interface moves in one direction, with equal probability to each direction.
   - If interfaces meet, they are both eliminated.
4. Eventually, the system has two interfaces $\Rightarrow$ Looking for a QSD supported on states with two clusters.

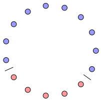

Figure: Down to two interfaces, completed

5. Per Step 3, the size of the remaining "yes" cluster performs a symmetric RW from Example 1, with $\rho = \frac{1}{N}$.
6. Comeback: QSD problem has been reduced to that of the RW from Example 1.

Example: Voter on a Cycle, Summary

Recall that for the RW on an Interval from Example 1 we had, (3):

$$\begin{cases} \nu_N(x) = \tan(\frac{\pi}{2N}) \sin(\frac{x}{N}\pi) \\ \lambda_N(\rho) = \frac{\rho}{2} \cos \frac{\pi}{N} + (1-2\rho). \end{cases}$$

Proposition 3

The unique QSD for the Voter Model on $\mathbb{Z}_N$ is a rotationally invariant distribution on configurations with exactly one cluster of each opinion, satisfying the following properties:

- *The size of each cluster is distributed according to ν_N.*
- *The survival probability is $\lambda_N(\frac{1}{N})$.*

Minimal Survival Probability

Necessary Condition: Geometric Tails

In light of Proposition 1 and the irreducibility, if ν is a QSD with survival probability λ,

$$E_x[\beta^\zeta] < \infty,\ x \in S,\ 1 < \beta < \lambda^{-1}.$$

This explains Assumption 1 part 4, leading to

Definition 2 (Minimal Survival Probability)

- *Define*

$$\lambda_0 = \inf\{\lambda < 1 : E_x[\lambda^{-\zeta}] < \infty \text{ for some } x \in S\}. \quad (4)$$

  That is λ_0 is the geometric tail of ζ under P_x for some (any) $x \in S$.
- *A QSD with survival probability λ_0 is called a minimal QSD.*

Corollary 4

1. $0 < \lambda_0 < 1$.
2. *For a QSD, the survival probability λ satisfies $\lambda_0 \le \lambda < 1$.*

- Why "minimal" QSD? For a QSD with survival probability λ,

$$E_\nu[\zeta] = \frac{1}{1-\lambda} \ge \frac{1}{1-\lambda_0},$$

QSD Regimes

Regimes Identified

Study of QSDs for a given survival probability λ is according to the following:

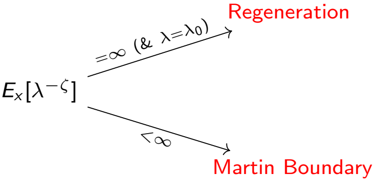

General features

- Regeneration
  - Reminiscent to positive recurrent MCs.
  - In this regime, if a QSD exists, it is unique.
- Martin Boundary
  - Reminiscent to Poisson Boundary for transient Markov chains.
  - Applicable to λ_0 in some cases.
  - Easy to construct examples where uniqueness does not hold.

Regeneration Regime

Definition 3 (Hitting times)

For $x \in S$ let

$$\tau_x = \inf\{t \geq \mathbb{N} : X_t = x\}.$$

Theorem 5 (Regeneration)

Suppose $E_x[\lambda_0^{-\zeta}] = \infty$. Then p possesses a QSD with survival probability λ_0 if and only if

$$E_x[\lambda_0^{-\zeta}, \zeta < \tau_x] < \infty \text{ for some } x \in S. \tag{5}$$

In this case the QSD with survival probability λ_0 is unique, given by

$$\nu(x) = \frac{\lambda_0^{-1} - 1}{E_x[\lambda_0^{-\zeta}, \zeta < \tau_x]} \tag{6}$$

Proposition 6 (Perron-Frobenius)

If S is finite, there exists a unique QSD. The QSD has survival probability λ_0 and is given by (6).

Martin Boundary

Preface

- Recall that this regime corresponds to existence and characterization of QSDs for survival probabilities λ satisfying $E_x[\lambda^{-\zeta}] < \infty$.
- As finite S was settled in Proposition 6: In what follows, we assume $S = \mathbb{N}$.
- Two main and newly obtained results, Theorem 7 and Theorem 10. The latter provides complete description of QSDs.

Theorem 7 (Asymptotics of GFs)

Suppose $\alpha > 1$ satisfy $E_x[\alpha^\zeta] < \infty$. Then

1. *If $\lim_{x\to\infty} E_x[\beta^\zeta] = \infty$ for some $\beta < \alpha$, then there exists a QSD corresponding to the survival probability α^{-1}.*
2. *If $\limsup_{x\to\infty} E_x[\alpha^\zeta] < \infty$, then there does not exist a QSD corresponding to the survival probability α^{-1}.*

Corollary 8 (Continuum of QSDs)

If $\lim_{x\to\infty} E_x[\beta^\zeta] = \infty$ for some $\beta < \frac{1}{\lambda_0}$, then for every $\lambda \in [\lambda_0, \beta^{-1})$ there exists a QSD corresponding to the survival probability λ.

Corollary 8: Two examples

Example 3 (Birth & Death)

Consider any Birth & Death process on $\{0\} \cup \mathbb{N}$ *satisfying the conditions of Assumption 1.*

- *Trivially, under* P_x, $\zeta \geq x$. *Therefore the condition in Corollary 8 holds for all* $\beta \in (1, \lambda_0^{-1})$.
- $\overset{\text{Corollary 8}}{\Longrightarrow}$ *existence of a QSD for each survival probability in* $[\lambda_0, 1)$

Example 4 (Subcritical Branching)

Consider a branching process with nondegenrate offspring distribution X, *satisfying* $E[X] < 1$. *Then*

- *A calculation with the generating function gives:*

$$\lambda_0 = E[X],$$
$$\lim_{x\to\infty} E_x[\beta^\zeta] = \infty \text{ for all } \beta \in (1, \lambda_0^{-1}).$$

- $\overset{\text{Corollary 8}}{\Longrightarrow}$ *existence of a QSD for each survival probability in* $[E[X], 1)$.
- *There exists a unique minimal QSD, obtained through Theorem 5.*

Martin Boundary

Overview

- Classically, Martin Boundary theory provides a compactification of the state space of a transient Markov Chain through a set of positive harmonic functions. These functions describe the tail of the chain: under the new topology the chain converges almost surely, with the limit viewed as where the process "exits" the state space.
- We borrow the ideas and obtain a similar compatification of the state space. In our work, the time arrow is reversed: we describe the behavior of the process according to how it is "coming from infinity".
- The result is a representation of all QSDs as a convex combination of the QSDs obtained as limits of Green's functions.

Preliminaries

- Fix $\alpha > 1$ satisfying $E_x[\alpha^\zeta] < \infty$.
- Define

$$K^\alpha(x,y) = \underbrace{\frac{\alpha - 1}{E_x[\alpha^\zeta - 1]}}_{\text{normalizing}} E_x[\sum_{s<\zeta} \alpha^s \delta_y(X_s)],$$

ℓ^1-normalized (in the second variable) Green's function for αp.

Martin Compactification: Construction

Definition 4 (Martin Compactification)

- *A sequence* $\mathbf{x} = (x_n : n \in \mathbb{N})$ *in* $\mathbb{N}$ *satisfying* $\lim_{n\to\infty} x_n = \infty$ *is convergent if* $\lim_{n\to\infty} K^\alpha(x_n, y)$ *exists for all* $y \in \mathbb{N}$.
- *Two convergent sequences* $\mathbf{x}$ *and* $\bar{\mathbf{x}}$ *are equivalent if*

$$\lim_{n\to\infty} K^\alpha(x_n, y) = \lim_{n\to\infty} K^\alpha(\bar{x}_n, y) \text{ for all } y \in \mathbb{N}.$$

- *Write* $[\mathbf{x}]$ *for the equivalency class of the convergent sequence* $\mathbf{x}$.
- *Martin Boundary: Let*

$$\begin{aligned} K^\alpha([\mathbf{x}], \cdot) &= \lim_{n\to\infty} K^\alpha(x_n, \cdot) && \leftarrow \text{ boundary points} \\ \partial^\alpha M &= \{[\mathbf{x}] : K^\alpha([\mathbf{x}], \cdot)\} && \leftarrow \text{ Martin Boundary} \\ M^\alpha &= \mathbb{N} \cup \partial^\alpha M && \leftarrow \text{ Martin Space} \end{aligned}$$

- *Metric: For* $a, b \in M^\alpha$*, let*

$$\rho^\alpha(a, b) = \sum_{n=1}^{\infty} \frac{1}{2^n} \left(|\delta_{a,n} - \delta_{b,n}| + d(K^\alpha(a, n), K^\alpha(b, n)) \right),$$

where $d(i, j) = \frac{|i-j|}{1+|i-j|}$.

Martin Compactification: Properties

Proposition 9 (Properties of the metric space)

- *(M^α, ρ^α) is a compact metric space and $\partial^\alpha M$ is closed.*
- *A sequence $\mathbf{a} = (a_n : n \in \mathbb{N})$ of elements of M^α is ρ^α convergent if and only if either*
  1. *There exists $a \in \mathbb{N}$ and $n_0 \in \mathbb{N}$ such that $a_n = a$ for all $n \geq n_0$:*
     - $\lim_{n\to\infty} a_n = a$*; or*
  2. *Condition 1 does not hold and there exists $[\mathbf{a}] \in \partial^\alpha M$ such that* $\lim_{n\to\infty} K^\alpha(a_n, \cdot) = K([\mathbf{a}], \cdot)$
     - $\lim_{n\to\infty} a_n = [\mathbf{a}]$.

Explanation

Roughly speaking (avoiding technical caveats):

- Each element of $x \in \mathbb{N}$ is identfied with the probability measure $K^\alpha(x, \cdot)$.
- M^α is obtained by closing this set with respect to pointwise limits, with set of "new" elements being $\partial^\alpha M$ (these limits may be sub-probability measures).
- The metric ρ^α corresponds to pointwise convergence.

Martin Boundary: Result

Let

$$\mathcal{K}^\alpha = \{[\mathbf{x}] \in \partial^\alpha M : K^\alpha([\mathbf{x}], \cdot) \text{ is a QSD}\}.$$

Theorem 10 (Martin/Choquet Representation)

Let $\alpha > 1$ satisfy $E_x[\alpha^\zeta] < \infty$. If ν is a QSD w/survival probability α^{-1} then there exists a Borel probability measure $\bar{F}_\nu$ on $\partial^\alpha M$ satisfying $\bar{F}_\nu(\mathcal{K}^\alpha) = 1$ and

$$\nu(y) = \int_{\partial^\alpha M} K^\alpha([\mathbf{x}], y) d\bar{F}_\nu([\mathbf{x}]).$$

Bottom line

- Every QSD is a convex combination of elements of $\mathcal{K}^\alpha$.

Theorem 10: Immediate Application

We revisit a previously introduced example:

Example 3: Birth & Death

- Corollary 8 $\Rightarrow$ a QSD for every survival probability in $[\lambda_0, 1)$.
- Theorem 10 $\Rightarrow$ a unique QSD for every survival probability in $[\lambda_0, 1)$.

Example: QSDs on a Tree

Example 5 (Example: QSDs on a Tree)

Consider the
d-regular tree with root ρ. Evolution:

- *From each state other than ρ move towards ρ with probability $q > \frac{1}{2}$.*
- *From each state move to a one of the neighbors away from the root with probability $1-q$, uniformly over the neighbors.*
- *From the root: move to the absorbing state Δ with probability $\delta \in (0, q)$, and stay put with probability $1-(1-q)-\delta = q-\delta$.*

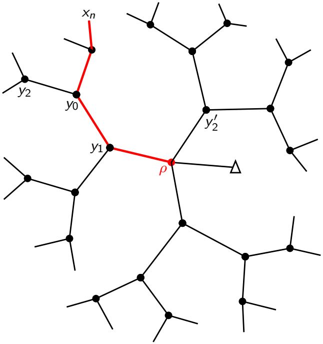

Example: QSDs on a Tree, completed

Proposition 11 (Minimal Survival Probability)

Consider the Tree from Example 5. Let

- $\lambda_\rho = 2\sqrt{q(1-q)}$
- *Let* $\delta_{cr} = \sqrt{q}(\sqrt{q} - \sqrt{1-q})$.

Then

$$\lambda_0 = \lambda_0(\delta) = \begin{cases} q - \delta + \frac{q(1-q)}{q-\delta} & \delta \in (0, \delta_{cr}) \\ \lambda_{cr} & \delta \in [\delta_{cr}, q] \end{cases}.$$

$\delta \in$	$(0, \delta_{cr})$	$\{\delta_{cr}\}$	$(\delta_{cr}, q]$
λ_0	$< \lambda_\rho$	$= \lambda_\rho$	
$E_\rho[\lambda_0^{-\zeta}]$	$= \infty$		$< \infty$

Proposition 12 (QSDs from K^α)

1. *For* $\lambda \leq \lambda_0$ *satisfying* $E_\rho[\lambda^{-\zeta}] < \infty$ *and every branch,* $\lim_{n\to\infty} K^{\lambda^{-1}}(x_n, \cdot)$ *exists along any sequence tending to infinity along the branch and is a QSD.*
2. *The QSDs obtained along each of the branches are distinct.*
3. *If* $E_\rho[\lambda_0^{-\zeta}] = \infty$, *there exists a unique QSD with survival probability* λ_0, *obtained through Theorem 5.*

Theorem 10 + Proposition 12 $\Rightarrow$ All QSDs for the model.

Thank you!