Exact simulation of the sample paths of a diffusion with a finite entrance boundary

Paul A. Jenkins

Departments of Statistics
University of Warwick

Dynstoch, University of Warwick
11 Sep 2014

Joint work with Dario Spanò

Outline

Diffusion model

The time-evolution of a genetic variant, or allele, is well approximated by a diffusion process on the interval [0, 1].

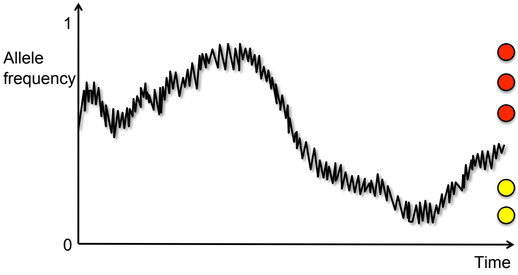

Diffusion model

The time-evolution of a genetic variant, or allele, is well approximated by a diffusion process on the interval [0, 1].

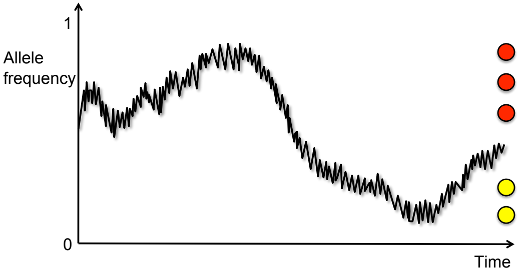

Wright-Fisher SDE

$$dX_t = \mu_\theta(X_t)dt + \sqrt{X_t(1 - X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$

The infinitesimal drift, $\mu_\theta(x)$, encapsulates directional forces such as natural selection, migration, mutation, ...

Population genetic Motivation I: Demographic inference

Given a sample of DNA sequences obtained in the present-day, what can we infer about the demographic history of the population?

Example (Gutenkunst *et al.*, 2009)

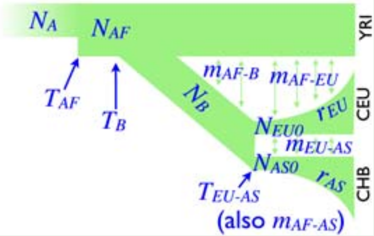

Expansion out-of-Africa

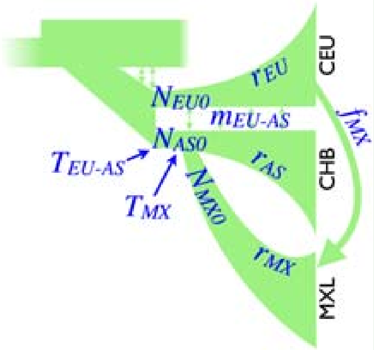

Settlement of the New World

Population genetic Motivation II: Time-series analysis of selection

Given a sample of genetic data obtained over several generations, what can we infer about the strength of natural selection?

Example (*Biston betulaeria*; Mathieson & McVean, 2013)

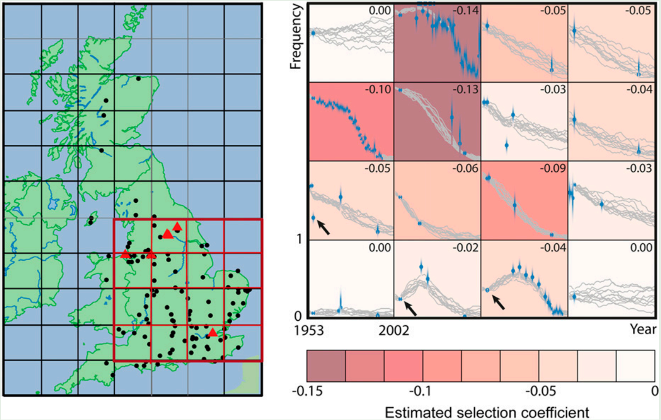

Inference from diffusion processes

- Like many interesting diffusions, the transition function of the Wright-Fisher diffusion is unknown.

Inference from diffusion processes

- Like many interesting diffusions, the transition function of the Wright-Fisher diffusion is unknown.
- Inference typically proceeds by
  1. Model-discretization such as an Euler approximation:

$$X_{t+dt} \mid (X_t = z) \sim \mathcal{N}(\mu_\theta(z)dt, \sigma^2(z)dt),$$

  2. ...followed by (sequential) Monte Carlo simulation (or numerical solution of Kolmogorov PDEs, or spectral expansions, ...)

Inference from diffusion processes

- Like many interesting diffusions, the transition function of the Wright-Fisher diffusion is unknown.
- Inference typically proceeds by
  1. Model-discretization such as an Euler approximation:

$$X_{t+dt} \mid (X_t = z) \sim \mathcal{N}(\mu_\theta(z)dt, \sigma^2(z)dt),$$

  2. . . . followed by (sequential) Monte Carlo simulation (or numerical solution of Kolmogorov PDEs, or spectral expansions, . . .)
- But—discretization introduces a bias we would like to remove.

Inference from diffusion processes

- Like many interesting diffusions, the transition function of the Wright-Fisher diffusion is unknown.
- Inference typically proceeds by
  1. Model-discretization such as an Euler approximation:

$$X_{t+dt} \mid (X_t = z) \sim \mathcal{N}(\mu_\theta(z)dt, \sigma^2(z)dt),$$

  2. ...followed by (sequential) Monte Carlo simulation (or numerical solution of Kolmogorov PDEs, or spectral expansions, ...)
- But—discretization introduces a bias we would like to remove.

Inference from diffusion processes

- Like many interesting diffusions, the transition function of the Wright-Fisher diffusion is unknown.
- Inference typically proceeds by
  1. Model-discretization such as an Euler approximation:

  $$X_{t+dt} \mid (X_t = z) \sim \mathcal{N}(\mu_\theta(z)dt, \sigma^2(z)dt),$$

  2. ...followed by (sequential) Monte Carlo simulation (or numerical solution of Kolmogorov PDEs, or spectral expansions, ...)
- But—discretization introduces a bias we would like to remove.

Three sentence summary

- There exist so-called exact algorithms for simulating diffusions without discretization error, even if the transition density is unknown.
- They can perform poorly when there are entrance boundaries.
- I will outline how to fix these problems.

Outline

Exact algorithm (EA)—one-dimensional bridge version

Goal: return exact bridge samples from the one-dimensional diffusion $X = (X_t : t \geq 0)$ on $\mathbb{R}$ satisfying

$$dX_t = \mu_\theta(X_t)dt + \sigma(X_t)dW_t, \qquad X_0 = x, \quad 0 \leq t \leq T.$$

Exact algorithm (EA)—one-dimensional bridge version

Goal: return exact bridge samples from the one-dimensional diffusion $X = (X_t : t \geq 0)$ on $\mathbb{R}$ satisfying

$$dX_t = \mu_\theta(X_t)dt + \sigma(X_t)dW_t, \qquad X_0 = x, \quad 0 \leq t \leq T.$$

1. Reduce the problem to unit diffusion coefficient via the Lamperti transform $X_t \mapsto Y_t$:

$$Y_t := \int^{X_t} \frac{1}{\sigma(u)} du,$$

so now we work with

$$\boxed{dY_t = \alpha_\theta(Y_t)dt + dW_t, \qquad Y_0 = y, \quad 0 \leq t \leq T.}$$

$$\boxed{dY_t = \alpha_\theta(Y_t)dt + dB_t, \qquad Y_0 = y, \quad 0 \le t \le T.}$$

Exact algorithm (EA)

2. Now we can consider a rejection algorithm using Brownian bridge paths as candidates.

   If $\mathbb{Q}_y$ is the target law (of Y) and $\mathbb{W}_y$ is the law of a Brownian motion then we need

   $$\frac{d\mathbb{Q}_y}{d\mathbb{W}_y}(Y)$$

   to provide the rejection probability

$$\boxed{dY_t = \alpha_\theta(Y_t)dt + dB_t, \qquad Y_0 = y, \quad 0 \le t \le T.}$$

Exact algorithm (EA)

2. Now we can consider a rejection algorithm using Brownian bridge paths as candidates.

   If $\mathbb{Q}_y$ is the target law (of Y) and $\mathbb{W}_y$ is the law of a Brownian motion then we need

$$\frac{d\mathbb{Q}_y}{d\mathbb{W}_y}(Y) = \exp\left\{\int_0^T \alpha_\theta(Y_t)dY_t - \frac{1}{2}\int_0^T \alpha_\theta^2(Y_t)dt\right\}$$

   to provide the rejection probability, by the Girsanov theorem.

$$\boxed{dY_t = \alpha_\theta(Y_t)dt + dB_t, \qquad Y_0 = y, \quad 0 \le t \le T.}$$

Exact algorithm (EA)

2. Now we can consider a rejection algorithm using Brownian bridge paths as candidates.

   If $\mathbb{Q}_y$ is the target law (of Y) and $\mathbb{W}_y$ is the law of a Brownian motion then we need

$$\frac{d\mathbb{Q}_y}{d\mathbb{W}_y}(Y) = \exp\left\{ \int_0^T \alpha_\theta(Y_t)dY_t - \frac{1}{2}\int_0^T \alpha_\theta^2(Y_t)dt \right\}$$

   to provide the rejection probability, by the Girsanov theorem.

- Such a rejection algorithm is impossible: it requires simulation of complete (infinite-dimensional) Brownian sample paths!

$$\boxed{dY_t = \alpha_\theta(Y_t)dt + dB_t, \qquad Y_0 = y, \quad 0 \le t \le T.}$$

Exact algorithm (EA)

3. Key observation: The Radon-Nikodým derivative can be put in the form

$$\frac{d\mathbb{Q}_y}{d\mathbb{W}_y}(Y) \propto \exp\left\{-\int_0^T \phi(Y_s)ds\right\} \le 1,$$

where $\phi(\cdot) := \frac{1}{2}[\alpha_\theta^2(\cdot) + \alpha_\theta'(\cdot)] + C.$

$$\boxed{dY_t = \alpha_\theta(Y_t)dt + dB_t, \qquad Y_0 = y, \quad 0 \le t \le T.}$$

Exact algorithm (EA)

3. Key observation: The Radon-Nikodým derivative can be put in the form

$$\frac{d\mathbb{Q}_y}{d\mathbb{W}_y}(Y) \propto \exp\left\{-\int_0^T \phi(Y_s)ds\right\} \le 1,$$

where $\phi(\cdot) := \frac{1}{2}[\alpha_\theta^2(\cdot) + \alpha_\theta'(\cdot)] + C$.

Assume we can arrange for $\phi \ge 0$. Then the right-hand side is the probability that a Poisson point process of unit rate on $[0, T] \times [0, \infty)$ has no points under the graph of $t \mapsto \phi(Y_s)$.

$$\boxed{\frac{d\mathbb{Q}_y}{d\mathbb{W}_y}(Y) \propto \exp\left\{-\int_0^T \phi(Y_s)ds\right\} \leq 1.}$$

Exact algorithm (EA)

4. A proposed Brownian path should be rejected if a simulated Poisson point process has any points under its graph.

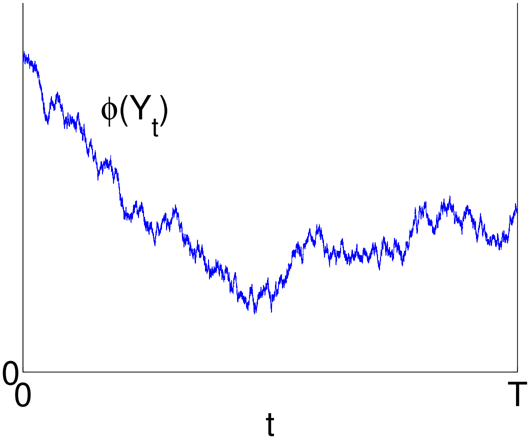

$$\boxed{\frac{d\mathbb{Q}_y}{d\mathbb{W}_y}(Y) \propto \exp\left\{-\int_0^T \phi(Y_s)ds\right\} \leq 1.}$$

Exact algorithm (EA)

4. A proposed Brownian path should be rejected if a simulated Poisson point process has any points under its graph.

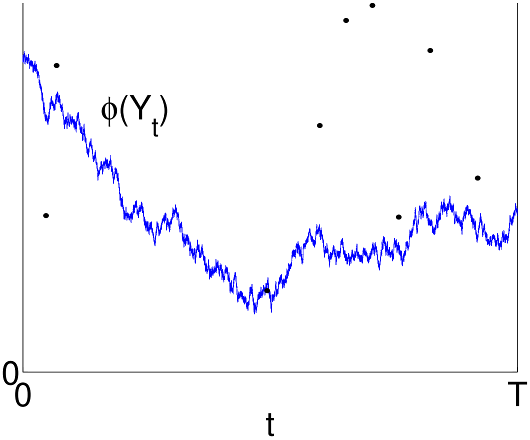

Retrospective sampling

Exact algorithm (EA) for simulating a bridge from Y_0 to Y_T

1. Simulate a Brownian bridge $(Y_t)_{0\leq t\leq T}$ from Y_0 to Y_T.

Retrospective sampling

Exact algorithm (EA) for simulating a bridge from Y_0 to Y_T

1. Simulate a Brownian bridge $(Y_t)_{0\leq t\leq T}$ from Y_0 to Y_T.
2. Simulate a Poisson point process of unit rate on $[0, T] \times [0, \infty)$.

Retrospective sampling

Exact algorithm (EA) for simulating a bridge from Y_0 to Y_T

1. Simulate a Brownian bridge $(Y_t)_{0\leq t\leq T}$ from Y_0 to Y_T.
2. Simulate a Poisson point process of unit rate on $[0, T] \times [0, \infty)$.
3. Accept if all points are in the epigraph of $t \mapsto \phi(Y_t)$, otherwise return to 1.

Retrospective sampling

Exact algorithm (EA) for simulating a bridge from Y_0 to Y_T

1. Simulate a Brownian bridge $(Y_t)_{0 \leq t \leq T}$ from Y_0 to Y_T.
2. Simulate a Poisson point process of unit rate on $[0, T] \times [0, \infty)$.
3. Accept if all points are in the epigraph of $t \mapsto \phi(Y_t)$, otherwise return to 1.

Problems

1. We still need an infinite-dimensional Brownian path.
2. The Poisson point process has unbounded intensity.

Retrospective sampling

Exact algorithm (EA) for simulating a bridge from Y_0 to Y_T

1. Simulate a Brownian bridge $(Y_t)_{0 \leq t \leq T}$ from Y_0 to Y_T.
2. Simulate a Poisson point process of unit rate on $[0, T] \times [0, \infty)$.
3. Accept if all points are in the epigraph of $t \mapsto \phi(Y_t)$, otherwise return to 1.

Problems

1. We still need an infinite-dimensional Brownian path.
2. The Poisson point process has unbounded intensity.

Solutions

1. Exploit retrospective sampling; switch the order of simulation!

Retrospective sampling

Exact algorithm (EA) for simulating a bridge from Y_0 to Y_T

1. Simulate a Brownian bridge $(Y_t)_{0\leq t\leq T}$ from Y_0 to Y_T.
2. Simulate a Poisson point process of unit rate on $[0,T]\times[0,\infty)$.
3. Accept if all points are in the epigraph of $t\mapsto\phi(Y_t)$, otherwise return to 1.

Problems

1. We still need an infinite-dimensional Brownian path.
2. The Poisson point process has unbounded intensity.

Solutions

1. Exploit retrospective sampling; switch the order of simulation!

Retrospective sampling

Exact algorithm (EA) for simulating a bridge from Y_0 to Y_T

1. Simulate a Poisson point process of unit rate on $[0, T] \times [0, \infty)$.
2. Simulate the Brownian bridge at the times of the Poisson points.
3. Accept if all points are in the epigraph of $t \mapsto \phi(Y_t)$, otherwise return to 1.

Problems

1. We still need an infinite-dimensional Brownian path.
2. The Poisson point process has unbounded intensity.

Solutions

1. Exploit retrospective sampling; switch the order of simulation!

Retrospective sampling

Exact algorithm (EA) for simulating a bridge from Y_0 to Y_T

1. Simulate a Poisson point process of unit rate on $[0, T] \times [0, \infty)$.
2. Simulate the Brownian bridge at the times of the Poisson points.
3. Accept if all points are in the epigraph of $t \mapsto \phi(Y_t)$, otherwise return to 1.

Problems

1. We still need an infinite-dimensional Brownian path.
2. The Poisson point process has unbounded intensity.

Solutions

1. Exploit retrospective sampling; switch the order of simulation!
2. Assume ϕ is bounded, $\phi \leq K$ (for now), and use Poisson thinning ("EA1").

Exact algorithm (EA1); Beskos & Roberts (2005)

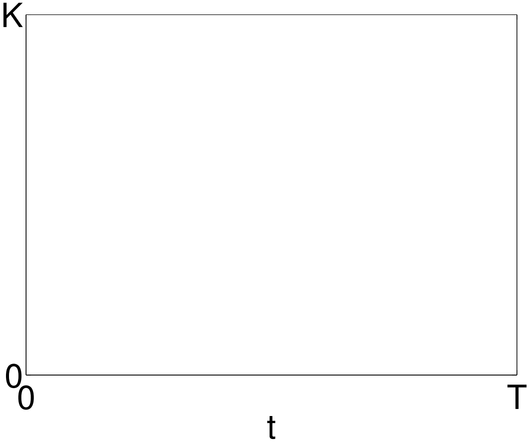

Exact algorithm (EA1); Beskos & Roberts (2005)

1. Simulate a Poisson point process on $[0, T] \times [0, K]$.

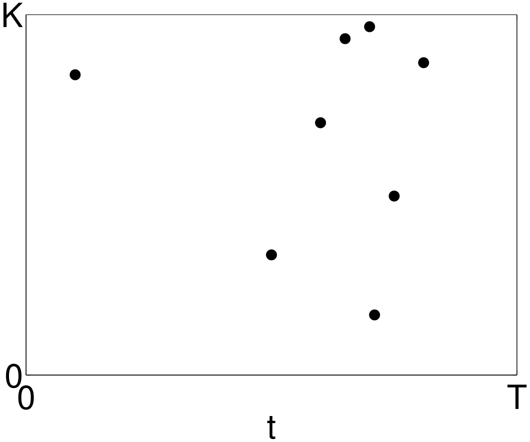

Exact algorithm (EA1); Beskos & Roberts (2005)

1. Simulate a Poisson point process on $[0, T] \times [0, K]$.
2. Simulate the Brownian bridge at the times of the Poisson points.

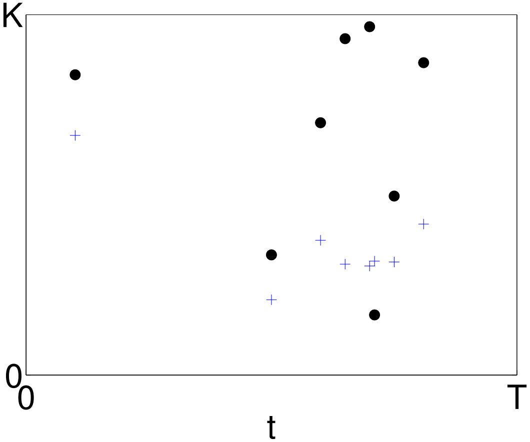

Exact algorithm (EA1); Beskos & Roberts (2005)

1. Simulate a Poisson point process on $[0, T] \times [0, K]$.
2. Simulate the Brownian bridge at the times of the Poisson points.
3. If any of the former are beneath any of the latter, return to 1.

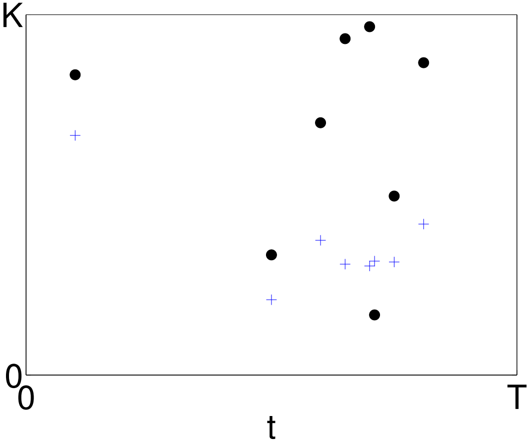

Exact algorithm (EA)

- Output of the algorithm is a set of skeleton points of the bridge.
- Any further points can be filled in by further draws from the Brownian bridge—no further reference to the target law, $\mathbb{Q}_y$, is necessary!

Exact algorithm (EA)

- Output of the algorithm is a set of skeleton points of the bridge.
- Any further points can be filled in by further draws from the Brownian bridge—no further reference to the target law, $\mathbb{Q}_y$, is necessary!

$$\boxed{\phi(\cdot) = \frac{1}{2}[\alpha_\theta^2(\cdot) + \alpha_\theta'(\cdot)] + C.}$$

- This function ϕ is important.
- The assumption $\phi \leq K$ is restrictive, but it can in fact be relaxed ("EA2", Beskos *et al.*, 2006).

Exact algorithm (EA)

- Output of the algorithm is a set of skeleton points of the bridge.
- Any further points can be filled in by further draws from the Brownian bridge—no further reference to the target law, $\mathbb{Q}_y$, is necessary!

$$\boxed{\phi(\cdot) = \frac{1}{2}[\alpha_\theta^2(\cdot) + \alpha_\theta'(\cdot)] + C.}$$

- This function ϕ is important.
- The assumption $\phi \leq K$ is restrictive, but it can in fact be relaxed ("EA2", Beskos *et al.*, 2006).

Exact algorithm (EA)

- Output of the algorithm is a set of skeleton points of the bridge.
- Any further points can be filled in by further draws from the Brownian bridge—no further reference to the target law, $\mathbb{Q}_y$, is necessary!

$$\boxed{\phi(\cdot) = \frac{1}{2}[\alpha_\theta^2(\cdot) + \alpha_\theta'(\cdot)] + C.}$$

- This function ϕ is important.
- The assumption $\phi \leq K$ is restrictive, but it can in fact be relaxed ("EA2", Beskos *et al.*, 2006).
- There have been many further refinements to this algorithm (multidimensions, jumps, killing, reflection, ...):
  Beskos *et al.* (2006, 2008, 2012), Casella & Roberts (2008, 2011), Chen & Huang (2013), Étoré & Martinez (2013), Giesecke & Smelov (2013), Gonçalves & Roberts (2013), Mousavi & Glynn (2013), Blanchet & Murthy (2014), Pollock *et al.* (2014).

Efficiency

- The exact algorithm will be less efficient wherever $\phi(X_t)$ is very large—unavoidable when the diffusion travels through a region where the drift (or its derivative) is very large.

Efficiency

- The exact algorithm will be less efficient wherever $\phi(X_t)$ is very large—unavoidable when the diffusion travels through a region where the drift (or its derivative) is very large.

Example: Entrance boundary at 0

- "A diffusion at x will almost surely not hit 0 before hitting any $b > x$. A diffusion started at 0 will enter $(0, \infty)$ in finite time."
- If $\sigma^2(x) = 1$, then ϕ explodes at the boundary.

Efficiency

- The exact algorithm will be less efficient wherever $\phi(X_t)$ is very large—unavoidable when the diffusion travels through a region where the drift (or its derivative) is very large.

Example: Entrance boundary at 0

- "A diffusion at x will almost surely not hit 0 before hitting any $b > x$. A diffusion started at 0 will enter $(0, \infty)$ in finite time."
- If $\sigma^2(x) = 1$, then ϕ explodes at the boundary.

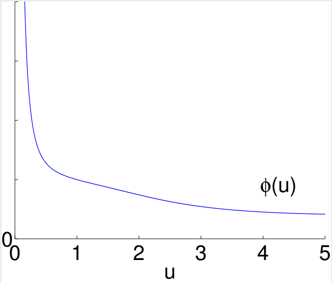

Efficiency

- The exact algorithm will be less efficient wherever $\phi(X_t)$ is very large—unavoidable when the diffusion travels through a region where the drift (or its derivative) is very large.

Example: Entrance boundary at 0

- "A diffusion at x will almost surely not hit 0 before hitting any $b > x$. A diffusion started at 0 will enter $(0, \infty)$ in finite time."
- If $\sigma^2(x) = 1$, then ϕ explodes at the boundary.

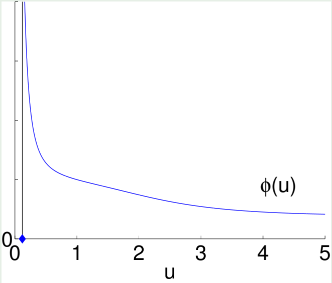

- Large ϕ is a symptom of a poor likelihood ratio, i.e. Brownian motion is a poor mimic of the target diffusion.

- Large ϕ is a symptom of a poor likelihood ratio, i.e. Brownian motion is a poor mimic of the target diffusion.
- Idea: Replace Brownian motion with a different candidate process—one with an entrance boundary.

- Large ϕ is a symptom of a poor likelihood ratio, i.e. Brownian motion is a poor mimic of the target diffusion.
- Idea: Replace Brownian motion with a different candidate process—one with an entrance boundary.
- But: the exact algorithms rely heavily on our knowledge about Brownian bridges:
  - The distribution of bridge coordinates.
  - The distribution of the minimum, m_T, and its time, t_m.
  - The distribution of bridge coordinates conditioned on (m_T, t_m).
  - The ability to sample from these distributions *exactly*.

- Large ϕ is a symptom of a poor likelihood ratio, i.e. Brownian motion is a poor mimic of the target diffusion.
- Idea: Replace Brownian motion with a different candidate process—one with an entrance boundary.
- But: the exact algorithms rely heavily on our knowledge about Brownian bridges:
  - The distribution of bridge coordinates.
  - The distribution of the minimum, m_T, and its time, t_m.
  - The distribution of bridge coordinates conditioned on (m_T, t_m).
  - The ability to sample from these distributions *exactly*.

- Large ϕ is a symptom of a poor likelihood ratio, i.e. Brownian motion is a poor mimic of the target diffusion.
- Idea: Replace Brownian motion with a different candidate process—one with an entrance boundary.
- But: the exact algorithms rely heavily on our knowledge about Brownian bridges:
  - The distribution of bridge coordinates.
  - The distribution of the minimum, m_T, and its time, t_m.
  - The distribution of bridge coordinates conditioned on (m_T, t_m).
  - The ability to sample from these distributions *exactly*.

Question. Does there exist a diffusion:

- with infinitesimal variance equal to 1,
- with an entrance boundary, and such that
- the finite-dimensional distributions of its bridges are known, and
- which can be simulated exactly, and
- (bonus) whose extrema are well characterized?

Outline

Bessel process

- Infinitesimal variance 1?
- Entrance boundary?
- Finite-dimensional distributions?
- Exact simulation?
- Distributions of extrema?

✓ Drift $\beta(y) = (\delta - 1)/(2y)$, variance $\sigma^2(y) = 1$.

Bessel process

- Infinitesimal variance 1?
- Entrance boundary?
- Finite-dimensional distributions?
- Exact simulation?
- Distributions of extrema?

✓ Drift $\beta(y) = (\delta - 1)/(2y)$, variance $\sigma^2(y) = 1$.

✓ Zero is an entrance boundary when $\delta \geq 2$.

Bessel process

- Infinitesimal variance 1?
- Entrance boundary?
- Finite-dimensional distributions?
- Exact simulation?
- Distributions of extrema?

✓ Drift $\beta(y) = (\delta - 1)/(2y)$, variance $\sigma^2(y) = 1$.

✓ Zero is an entrance boundary when $\delta \geq 2$.

✓ $p_{(y,0)\to(z,T)}(x;t) =$

$$\frac{T}{2t(T-t)} e^{-\left(\frac{z(T-t)}{2tT} + \frac{xT}{2t(T-t)} + \frac{yt}{2T(T-t)}\right)} \frac{I_\nu\left(\frac{\sqrt{xz}}{t}\right) I_\nu\left(\frac{\sqrt{xy}}{(T-t)^2}\right)}{I_\nu\left(\frac{\sqrt{yz}}{T^2}\right)},$$

where $\nu = 2(\delta + 1)$, is the transition density of the (squared) Bessel bridge.

Bessel process

- Infinitesimal variance 1?
- Entrance boundary?
- Finite-dimensional distributions?
- Exact simulation?
- Distributions of extrema?

✓ Drift $\beta(y) = (\delta - 1)/(2y)$, variance $\sigma^2(y) = 1$.

✓ Zero is an entrance boundary when $\delta \geq 2$.

✓ $p_{(y,0)\to(z,T)}(x;t) =$

$$\frac{T}{2t(T-t)} e^{-\left(\frac{z(T-t)}{2tT} + \frac{xT}{2t(T-t)} + \frac{yt}{2T(T-t)}\right)} \frac{I_\nu(\frac{\sqrt{xz}}{t}) I_\nu(\frac{\sqrt{xy}}{(T-t)^2})}{I_\nu(\frac{\sqrt{yz}}{T^2})},$$

where $\nu = 2(\delta + 1)$, is the transition density of the (squared) Bessel bridge.

✓ $\delta \in \mathbb{Z}_{\geq 0}$: radial part of a δ-dimensional Brownian motion.
$\delta \in \mathbb{R}_{\geq 0}$: See Makarov & Glew (2010).

Bessel process

- Infinitesimal variance 1?
- Entrance boundary?
- Finite-dimensional distributions?
- Exact simulation?
- Distributions of extrema?

✓ Drift $\beta(y) = (\delta - 1)/(2y)$, variance $\sigma^2(y) = 1$.

✓ Zero is an entrance boundary when $\delta \geq 2$.

✓ $p_{(y,0)\to(z,T)}(x;t) =$

$$\frac{T}{2t(T-t)} e^{-\left(\frac{z(T-t)}{2tT} + \frac{xT}{2t(T-t)} + \frac{yt}{2T(T-t)}\right)} \frac{I_\nu\left(\frac{\sqrt{xz}}{t}\right) I_\nu\left(\frac{\sqrt{xy}}{(T-t)^2}\right)}{I_\nu\left(\frac{\sqrt{yz}}{T^2}\right)},$$

where $\nu = 2(\delta + 1)$, is the transition density of the (squared) Bessel bridge.

✓ $\delta \in \mathbb{Z}_{\geq 0}$: radial part of a δ-dimensional Brownian motion.
$\delta \in \mathbb{R}_{\geq 0}$: See Makarov & Glew (2010).

(✓) Partly.

Bessel-EA

- Exact simulation from a diffusion with law $\mathbb{Q}_y$ using the Bessel process (law $\mathbb{B}^\delta_y \gg \mathbb{Q}_y$) is possible by the following:

Theorem.

Under regularity conditions (similar to EA), $\mathbb{Q}_y$ is the marginal distribution of Y when

$$(Y, \Phi) \sim (\mathbb{B}^\delta_y \otimes \mathbb{L})\Big| \left\{\Phi \subseteq \text{epigraph}\left[\widetilde{\phi}(Y)\right]\right\},$$

where $\mathbb{L}$ is the law of a Poisson point process Φ of unit rate on $[0, T] \times [0, \infty)$,and

$$\widetilde{\phi}(u) := \frac{1}{2}[\alpha^2_\theta(u) - \beta^2(u) + \alpha'_\theta(u) - \beta'(u)] + C.$$

Outline of proof.

Similar to the Brownian case: regularity conditions permit a Girsanov transformation and rearrangement so that

$$\frac{d\mathbb{Q}_y}{d\mathbb{B}^\delta_y}(Y) \propto \exp\left\{-\int_0^T \widetilde{\phi}(Y_t)dt\right\} \leq 1,$$

provides the rejection probability for sampling from the conditional law

$$(\mathbb{B}^\delta_y \otimes \mathbb{L})\Big| \left\{\Phi \subseteq \text{epigraph}\left[\widetilde{\phi}(Y)\right]\right\}. \quad \square$$

Outline of proof.

Similar to the Brownian case: regularity conditions permit a Girsanov transformation and rearrangement so that

$$\frac{d\mathbb{Q}_y}{d\mathbb{B}_y^\delta}(Y) \propto \exp\left\{-\int_0^T \widetilde{\phi}(Y_t)dt\right\} \leq 1,$$

provides the rejection probability for sampling from the conditional law

$$(\mathbb{B}_y^\delta \otimes \mathbb{L})\Big| \left\{\Phi \subseteq \text{epigraph}\left[\widetilde{\phi}(Y)\right]\right\}. \quad \square$$

So what?

- We have just replaced one candidate process for another, the only substantial difference the appearance of

$$\widetilde{\phi}(u) := \frac{1}{2}[\alpha_\theta^2(u) - \beta^2(u) + \alpha_\theta'(u) - \beta'(u)] + C.$$

instead of

$$\phi(u) := \frac{1}{2}[\alpha_\theta^2(u) + \alpha_\theta'(u)] + C.$$

Example: A population growth model.

- A diffusion $(X_t)_{0\leq t\leq T}$ with drift and diffusion coefficients

$$\mu(x) = \kappa x, \qquad \sigma^2(x) = x + \omega x^2,$$

commenced from $X_0 = x_0$ and grown to $X_T = x_T$.

Example: A population growth model.

- A diffusion $(X_t)_{0\leq t\leq T}$ with drift and diffusion coefficients

$$\mu(x) = \kappa x, \qquad \sigma^2(x) = x + \omega x^2,$$

commenced from $X_0 = x_0$ and grown to $X_T = x_T$.

- The population has not died out, so we can condition the process on non-absorption at 0.

Example: A population growth model.

- A diffusion $(X_t)_{0 \leq t \leq T}$ with drift and diffusion coefficients

$$\mu(x) = \kappa x, \qquad \sigma^2(x) = x + \omega x^2,$$

commenced from $X_0 = x_0$ and grown to $X_T = x_T$.

- The population has not died out, so we can condition the process on non-absorption at 0.
- Conditioning and Lamperti transforming leads to new drift

$$\alpha(y) = \frac{\kappa}{\sqrt{\omega}} \tanh\left[\frac{\sqrt{\omega} y}{2}\right] - \frac{\sqrt{\omega}}{2} \coth\left[\sqrt{\omega} y\right] + \frac{\omega - 2\kappa}{\sqrt{w}} \frac{\tanh\left[\frac{\sqrt{\omega} y}{2}\right]}{1 - \cosh^{\frac{4\kappa}{\omega} - 2}\left[\frac{\sqrt{\omega} y}{2}\right]},$$

with an entrance boundary at 0.

Example: A population growth model.

- What does the drift look like at the boundary?

$$\alpha(y) = \frac{3}{2y} + O(y) \qquad \text{as } y \to 0.$$

Example: A population growth model.

- What does the drift look like at the boundary?

$$\alpha(y) = \frac{3}{2y} + O(y) \qquad \text{as } y \to 0.$$

- Compare with the Bessel process: $$\beta(y) = \frac{\delta - 1}{2y}.$$

  So we should choose $\delta = 4$ for our candidate process.

Example: A population growth model.

- What does the drift look like at the boundary?

$$\alpha(y) = \frac{3}{2y} + O(y) \qquad \text{as } y \to 0.$$

- Compare with the Bessel process: $\beta(y) = \frac{\delta - 1}{2y}$.

  So we should choose $\delta = 4$ for our candidate process.

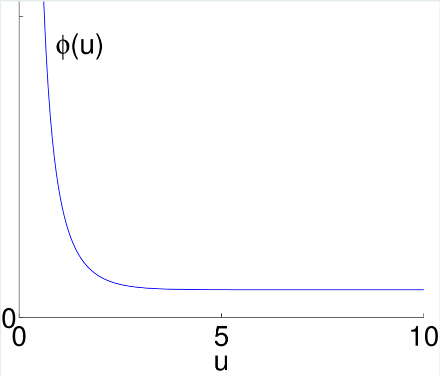

Example: A population growth model.

- What does the drift look like at the boundary?

$$\alpha(y) = \frac{3}{2y} + O(y) \qquad \text{as } y \to 0.$$

- Compare with the Bessel process: $\beta(y) = \dfrac{\delta - 1}{2y}$.

  So we should choose $\delta = 4$ for our candidate process.

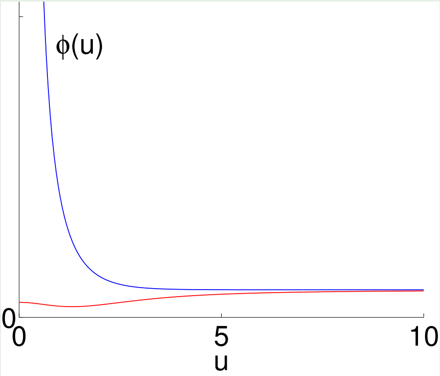

Example: A population growth model.

- $\widetilde{\phi}$ is (tightly) bounded (by K say), while ϕ is unbounded as $y \to 0$.

Example: A population growth model.

- $\widetilde{\phi}$ is (tightly) bounded (by K say), while ϕ is unbounded as $y \to 0$.
- Hence we can use the following Bessel-EA to return skeleton bridges:
  1. Simulate a Poisson point process on $[0, T] \times [0, K]$.
  2. Simulate a Bessel bridge of dimension $\delta = 4$ at the times of the Poisson points.
  3. If any of the former are beneath any of the latter, return to 1.

Results

Bessel-EA1 $Y_0 = y$ to $Y_{0.15} = 1$, $\omega = 3$.

κ	y	Attempts	Poisson points	Skeleton points	Random variables	Total Time (s)
1.0	10.0	1.1	0.2	0.2	1.9	0
1.0	1.0	1.0	0.2	0.2	1.9	0
1.0	0.25	1.0	0.2	0.2	2.0	0
1.0	0.15	1.0	0.2	0.2	2.0	1
1.0	0.1	1.1	0.2	0.2	2.0	1
1.0	0.025	1.0	0.2	0.2	2.0	0

Brownian-EA (“EA2”)

κ	y	Attempts	Poisson points	Skeleton points	Random variables	Total Time (s)
1.0	10.0	1.0	0.1	0.1	7.3	0
1.0	1.0	1.1	0.1	0.1	7.4	0
1.0	0.25	1.2	1288.6	420.6	3846.1	6
1.0	0.15	1.4	7531.1	617.4	16921.4	16
1.0	0.1	DNF	DNF	DNF	DNF	DNF
1.0	0.025	DNF	DNF	DNF	DNF	DNF

Results

Bessel-EA1 $\qquad Y_0 = y$ to $Y_{0.15} = 1$, $\omega = 3$.

κ	y	Attempts	Poisson points	Skeleton points	Random variables	Total Time (s)
10.0	10.0	5.2	14.1	6.8	56.4	1
10.0	1.0	3.0	7.9	4.9	36.4	1
10.0	0.25	2.3	6.1	4.4	30.8	1
10.0	0.15	2.2	6.0	4.3	30.3	0
10.0	0.1	2.2	5.9	4.4	30.4	0
10.0	0.025	2.1	5.8	4.3	29.6	1

Brownian-EA (“EA2”)

κ	y	Attempts	Poisson points	Skeleton points	Random variables	Total Time (s)
10.0	10.0	5.0	9.8	4.8	40.9	0
10.0	1.0	2.9	5.9	3.6	29.8	0
10.0	0.25	2.6	81.4	10.7	201.9	0
10.0	0.15	2.9	23052.1	1981.9	52056.9	52
10.0	0.1	DNF	DNF	DNF	DNF	DNF
10.0	0.025	DNF	DNF	DNF	DNF	DNF

When is the singularity in the drift at an entrance boundary matched by a Bessel process?

When is the singularity in the drift at an entrance boundary matched by a Bessel process?
Here's a partial answer.

Theorem.

Suppose we have a diffusion Y satisfying the requirements of EA1. Then the diffusion Y^* obtained by conditioning this process on $\{T_b < T_0\}$, can be simulated via Bessel-EA1 with $\delta = 3$.

When is the singularity in the drift at an entrance boundary matched by a Bessel process?
Here's a partial answer.

Theorem.

Suppose we have a diffusion Y satisfying the requirements of EA1. Then the diffusion Y^* obtained by conditioning this process on $\{T_b < T_0\}$, can be simulated via Bessel-EA1 with $\delta = 3$.

Outline of proof.

- Deduce regularity requirements for Bessel-EA1 from the assumptions of EA1.
- Compute the conditioned drift $\alpha^*(y)$ by bare hands, using a Doob h-transform.
- We find $\widetilde{\phi}^*(u)$ is bounded iff $\delta = 3$ (among all possible $\delta \geq 2$). $\square$

Outline

Example: The Wright-Fisher diffusion with natural selection

- The frequency $X_t \in [0, 1]$ of a gene in a large population evolves according to
$$dX_t = \gamma X_t(1 - X_t)dt + \sqrt{X_t(1 - X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$

Example: The Wright-Fisher diffusion with natural selection

- The frequency $X_t \in [0,1]$ of a gene in a large population evolves according to
$$dX_t = \gamma X_t(1 - X_t)dt + \sqrt{X_t(1 - X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$
- γ parametrizes the selective advantage of this gene.

Example: The Wright-Fisher diffusion with natural selection

- The frequency $X_t \in [0, 1]$ of a gene in a large population evolves according to
$$dX_t = \gamma X_t(1 - X_t)dt + \sqrt{X_t(1 - X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$
- γ parametrizes the selective advantage of this gene.
- A gene is observed to be at frequency $X_0 = x_0$ and $X_T = x_T$. What do its sample paths look like?

Example: The Wright-Fisher diffusion with natural selection

- The frequency $X_t \in [0,1]$ of a gene in a large population evolves according to
$$dX_t = \gamma X_t(1-X_t)dt + \sqrt{X_t(1-X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$
- γ parametrizes the selective advantage of this gene.
- A gene is observed to be at frequency $X_0 = x_0$ and $X_T = x_T$. What do its sample paths look like?
- 0 is absorbing so we should condition the diffusion not to have hit 0 ($T_0 > T$).

Example: The Wright-Fisher diffusion with natural selection

- The frequency $X_t \in [0,1]$ of a gene in a large population evolves according to
$$dX_t = \gamma X_t(1-X_t)dt + \sqrt{X_t(1-X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$
- γ parametrizes the selective advantage of this gene.
- A gene is observed to be at frequency $X_0 = x_0$ and $X_T = x_T$. What do its sample paths look like?
- 0 is absorbing so we should condition the diffusion not to have hit 0 ($T_0 > T$).
- Conditioning and Lamperti transforming leads to new drift
$$\alpha(y) = \frac{1}{2}\gamma \sin(y) \coth\left[\gamma \sin^2\left(\frac{y}{2}\right)\right] - \cot(y).$$

(Schraiber *et al.*, 2013).

Example: The Wright-Fisher diffusion.

- What does the drift look like at the boundary?

$$\alpha(y) = \frac{3}{2y} + O(y) \qquad \text{as } y \to 0.$$

Example: The Wright-Fisher diffusion.

- What does the drift look like at the boundary?

$$\alpha(y) = \frac{3}{2y} + O(y) \qquad \text{as } y \to 0.$$

- So we should again choose $\delta = 4$ for our candidate process.

Example: The Wright-Fisher diffusion.

- What does the drift look like at the boundary?

$$\alpha(y) = \frac{3}{2y} + O(y) \qquad \text{as } y \to 0.$$

- So we should again choose $\delta = 4$ for our candidate process.

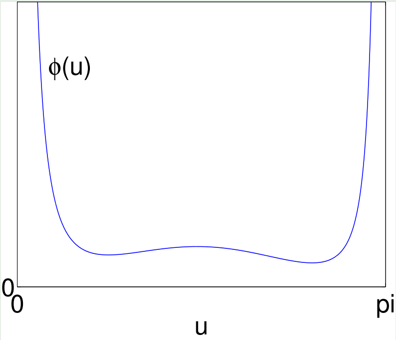

Example: The Wright-Fisher diffusion.

- What does the drift look like at the boundary?

$$\alpha(y) = \frac{3}{2y} + O(y) \qquad \text{as } y \to 0.$$

- So we should again choose $\delta = 4$ for our candidate process.

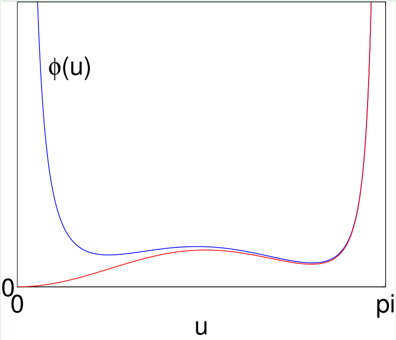

Example: The Wright-Fisher diffusion.

- What does the drift look like at the boundary?

$$\alpha(y) = \frac{3}{2y} + O(y) \qquad \text{as } y \to 0.$$

- So we should again choose $\delta = 4$ for our candidate process.

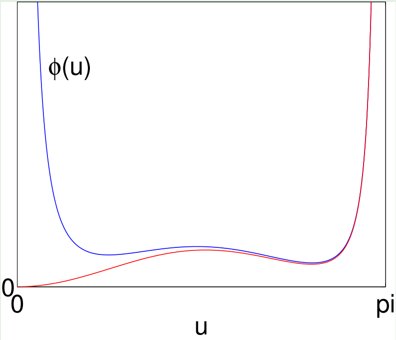

- Problem: $\widetilde{\phi}(u)$ is still unbounded as $u \to \pi$. This leaves us with only an approximately exact algorithm.

Ongoing work to fix this issue

In the spirit of (Brownian)-EA2, conditioning on the maximum of a Bessel bridge would solve the problem:

Ongoing work to fix this issue

In the spirit of (Brownian)-EA2, conditioning on the maximum of a Bessel bridge would solve the problem:

1. Simulate the maximum M_T (and the time, t_M, it is attained) of a Bessel bridge from Y_0 to Y_T.

Ongoing work to fix this issue

In the spirit of (Brownian)-EA2, conditioning on the maximum of a Bessel bridge would solve the problem:

1. Simulate the maximum M_T (and the time, t_M, it is attained) of a Bessel bridge from Y_0 to Y_T.
2. Use another version of the Wright-Fisher diffusion as our candidate process.

The Wright-Fisher diffusion with mutation but no selection

$$dX_t = [\theta_1(1 - X_t) - \theta_2 X_t]dt + \sqrt{X_t(1 - X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$

The transition density has eigenfunction expansion

$$f(x, y; t) = \sum_{m=0}^{\infty} q_m(t) \sum_{l=0}^{m} \underbrace{\mathcal{B}_{m,x}(l)}_{\text{Binomial PMF}} \cdot \underbrace{\mathcal{D}_{\theta_1+l,\theta_2+m-l}(y)}_{\text{Beta density}},$$

where $q_m(t)$ is the transition function of a certain pure death process on $\mathbb{N}$ (related to Kingman's coalescent).

The Wright-Fisher diffusion with mutation but no selection

$$dX_t = [\theta_1(1 - X_t) - \theta_2 X_t]dt + \sqrt{X_t(1 - X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$

The transition density has eigenfunction expansion

$$f(x, y; t) = \sum_{m=0}^{\infty} q_m(t) \sum_{l=0}^{m} \underbrace{\mathcal{B}_{m,x}(l)}_{\text{Binomial PMF}} \cdot \underbrace{\mathcal{D}_{\theta_1+l,\theta_2+m-l}(y)}_{\text{Beta density}},$$

where $q_m(t)$ is the transition function of a certain pure death process on $\mathbb{N}$ (related to Kingman's coalescent).

- This is a known infinite mixture of beta random variables.

The Wright-Fisher diffusion with mutation but no selection

$$dX_t = [\theta_1(1 - X_t) - \theta_2 X_t]dt + \sqrt{X_t(1 - X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$

The transition density has eigenfunction expansion

$$f(x, y; t) = \sum_{m=0}^{\infty} q_m(t) \sum_{l=0}^{m} \underbrace{\mathcal{B}_{m,x}(l)}_{\text{Binomial PMF}} \cdot \underbrace{\mathcal{D}_{\theta_1+l,\theta_2+m-l}(y)}_{\text{Beta density}},$$

where $q_m(t)$ is the transition function of a certain pure death process on $\mathbb{N}$ (related to Kingman's coalescent).

- This is a known infinite mixture of beta random variables.

Convenient for simulation!

1. Simulate $M \sim \{q_m(t) : m = 0, 1, \ldots\}$.
   (a realization of Kingman's coalescent with mutation, time t).

The Wright-Fisher diffusion with mutation but no selection

$$dX_t = [\theta_1(1 - X_t) - \theta_2 X_t]dt + \sqrt{X_t(1 - X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$

The transition density has eigenfunction expansion

$$f(x, y; t) = \sum_{m=0}^{\infty} q_m(t) \sum_{l=0}^{m} \underbrace{\mathcal{B}_{m,x}(l)}_{\text{Binomial PMF}} \cdot \underbrace{\mathcal{D}_{\theta_1+l,\theta_2+m-l}(y)}_{\text{Beta density}},$$

where $q_m(t)$ is the transition function of a certain pure death process on $\mathbb{N}$ (related to Kingman's coalescent).

- This is a known infinite mixture of beta random variables.

Convenient for simulation!

1. Simulate $M \sim \{q_m(t) : m = 0, 1, \ldots\}$.
   (a realization of Kingman's coalescent with mutation, time t).
2. Simulate $L \sim \text{Binomial}(M, x)$.

The Wright-Fisher diffusion with mutation but no selection

$$dX_t = [\theta_1(1 - X_t) - \theta_2 X_t]dt + \sqrt{X_t(1 - X_t)}dW_t, \qquad X_0 = x, \quad t \geq 0.$$

The transition density has eigenfunction expansion

$$f(x, y; t) = \sum_{m=0}^{\infty} q_m(t) \sum_{l=0}^{m} \underbrace{\mathcal{B}_{m,x}(l)}_{\text{Binomial PMF}} \cdot \underbrace{\mathcal{D}_{\theta_1+l,\theta_2+m-l}(y)}_{\text{Beta density}},$$

where $q_m(t)$ is the transition function of a certain pure death process on $\mathbb{N}$ (related to Kingman's coalescent).

- This is a known infinite mixture of beta random variables.

Convenient for simulation!

1. Simulate $M \sim \{q_m(t) : m = 0, 1, \ldots\}$.
   (a realization of Kingman's coalescent with mutation, time t).
2. Simulate $L \sim \text{Binomial}(M, x)$.
3. Return $Y \sim \text{Beta}(\theta_1 + L, \theta_2 + M - L)$.

Issues

Example

Mixture weights are known only as an infinite series:

$$q_m(t) = \sum_{k=m}^{\infty} (-1)^{k-m} \frac{(\theta + 2k - 1)\Gamma(\theta + m + k - 1)}{m!(k-m)!\Gamma(\theta + m)} e^{-k(k+\theta-1)t/2}.$$

Proposition (Jenkins & Spanò, in preparation).

The coefficients of the ancestral process of Kingman's coalescent,

$$\{q_m(t) : m = 0, 1, \ldots\},$$

can be rearranged in such a way that this distribution can be simulated exactly.

Outline

Summary

- It is possible to simulate efficiently from several diffusions with a finite entrance boundary, without discretization error.

Summary

- It is possible to simulate efficiently from several diffusions with a finite entrance boundary, without discretization error.
- Candidate diffusions other than Brownian motion:
  - Bessel process
  - Wright-Fisher diffusion

  suggest the potential for further generalizing the exact algorithms.

Plug

Jenkins, P. A. "Exact simulation of the sample paths of a diffusion with a finite entrance boundary." arXiv:1311.5777.

Acknowledgements

Many helpful conversations:

- Steve Evans
- Murray Pollock
- Gareth Roberts
- Joshua Schraiber

Plug

Jenkins, P. A. “Exact simulation of the sample paths of a diffusion with a finite entrance boundary.” arXiv:1311.5777.

Acknowledgements

Many helpful conversations:

- Steve Evans
- Murray Pollock
- Gareth Roberts
- Joshua Schraiber

Thank you for listening!