Noise-Tolerant Bayesian Bisection

Peter I. Frazier, German Gutierrez, Shane G. Henderson

Operations Research & Information Engineering, Cornell University

November 10, 2010
INFORMS Annual Meeting,
Austin, TX

Stochastic Root-Finding

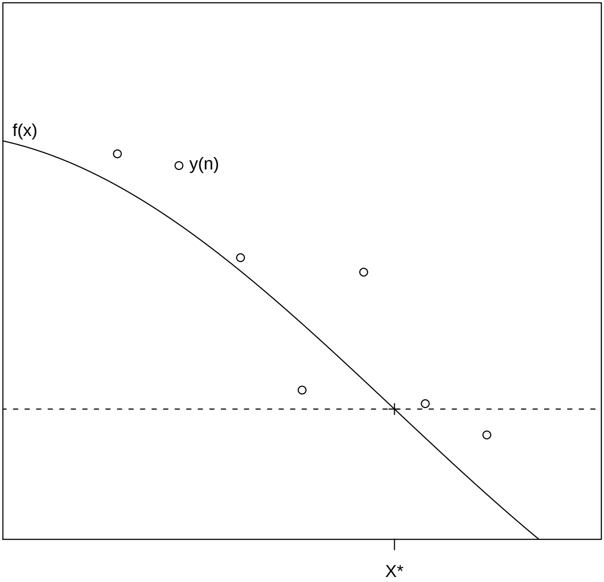

Stochastic Root-Finding

- $f : \mathbb{R} \mapsto \mathbb{R}$ is a decreasing function.
- The only way to evaluate f is via stochastic simulation.
- We observe $y_n = f(x_n) + \varepsilon_n$, where ε_n is independent noise.
- Our goal is to find a root x_*, i.e., a point x_* such that $f(x_*) = 0$.
- Central Question: Given a budget of N measurements, $x_1, \ldots, x_N$, how should we place them to find x_* as accurately as possible?

Overview

- Stochastic root-finding is a classic problem in simulation and sequential statistics, dating to the work of Robbins & Monro in 1951.
- Contribution of this talk:
  1. We formulate an idealized version of this problem as a dynamic program.
  2. We solve this dynamic program analytically and find the Bayes-optimal policy.
  3. We analyze this policy in the non-Bayesian non-idealized setting.

Stochastic Root-Finding: Model

- Assume we observe only whether y_n is larger or smaller than 0.

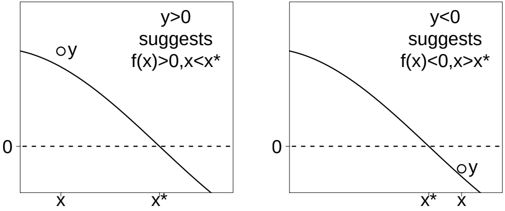

- Assume nature gives the incorrect sign with fixed known probability q:

$$\mathrm{sgn}(y_n) = \begin{cases} -\mathrm{sgn}(f(x_n)), & \text{with probability } q, \\ \mathrm{sgn}(f(x_n)), & \text{with probability } 1-q \end{cases}$$

Posterior Distributions

- Place a prior density p_0 on the root x_*, e.g., uniform on $[0,1]$.
- Each measurement x_n produces a new posterior density p_n on x_*:

$$p_n(x) = \mathbb{P}\{X_* \in dx \mid x_{1:n-1}, y_{1:n-1}\}$$

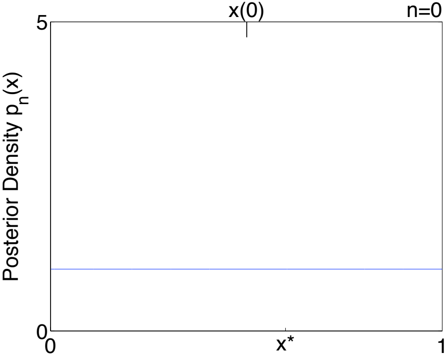

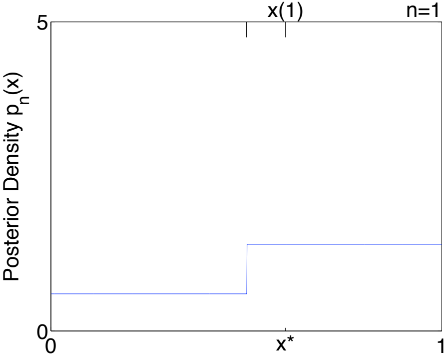

Posterior Distributions

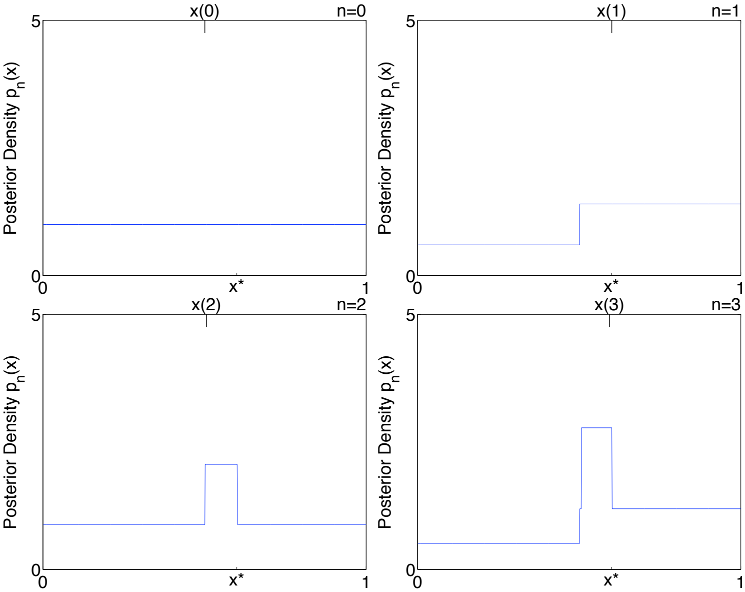

Bayes Optimal Sequential Policy

- At the final time, we take the entropy as our measure of uncertainty.

$$H(p_N) = -\int p_N(x) \log p_N(x) \, dx.$$

- A sequential policy π is an adaptive rule for choosing each x_n.
- $\mathbb{E}^{\pi}[H(p_N)]$ is the average final entropy when choosing $x_1, \ldots, x_N$ according to policy π,
- A **Bayes-Optimal Policy** is any solution to

$$\inf_{\pi \in \Pi} \mathbb{E}^{\pi}[H(p_N)],$$

where Π is the space of all policies.

Dynamic Programming

- To find the Bayes-optimal policy, we must solve:

$$\inf_{\pi} \mathbb{E}^{\pi}\left[H(p_N)\right],$$

- This is a **Markov decision process** (MDP), where the Markov process being controlled is the sequence of posteriors p_n.
- As an MDP, its solution is characterized by the **dynamic programming equations**,

$$\begin{aligned} V_N(p_N) &= H(p_N), \\ V_n(p_n) &= \inf_{x_n \in [0,1]} \mathbb{E}\left[V_{n+1}(p_{n+1}) \mid p_n\right]. \end{aligned}$$

Dynamic Programming

- Our dynamic program:

$$V_N(p_N) = H(p_N),$$
$$V_n(p_n) = \inf_{x_n \in [0,1]} \mathbb{E}\left[V_{n+1}(p_{n+1}) \mid p_n\right].$$

- Solving this dynamic program numerically via brute force is impossible because it would require computing and storing $V_n(p_n)$ for every possible density p_n....
- .... But in this problem, we can solve the dynamic program analytically!

Main Result: Bayes Optimality

Theorem

The value function can be written explicitly as

$$V(p_n) = H(p_n) - (N-n)\left[-q\log_2(q) - (1-q)\log_2(1-q)\right],$$

and **the policy that chooses x_n at the median of p_n is Bayes-optimal**.

Example

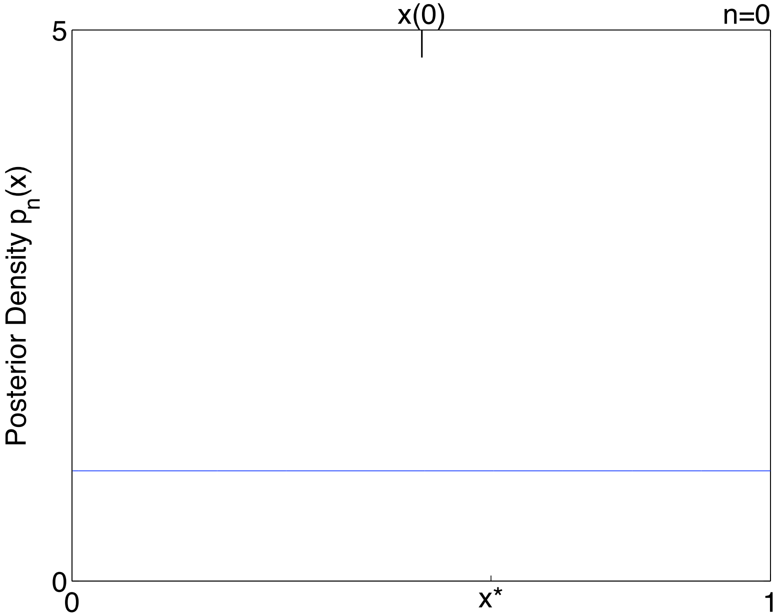

Example

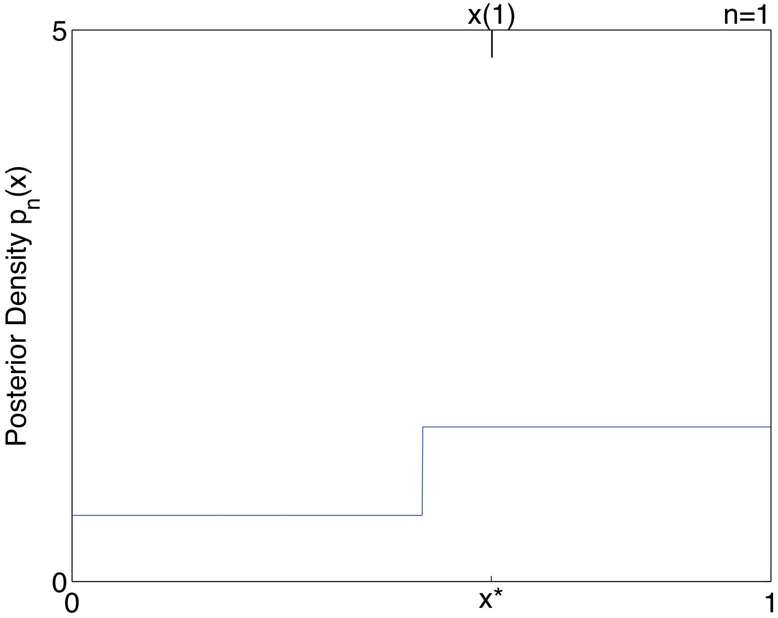

Example

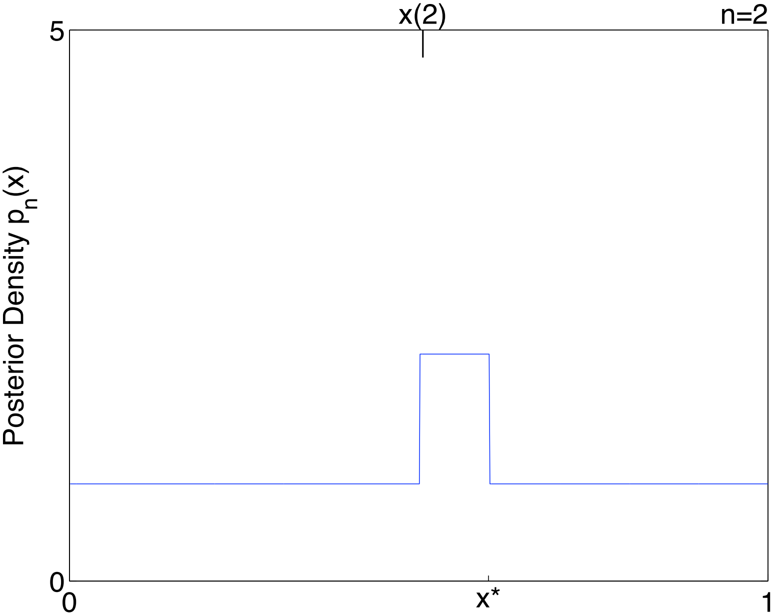

Example

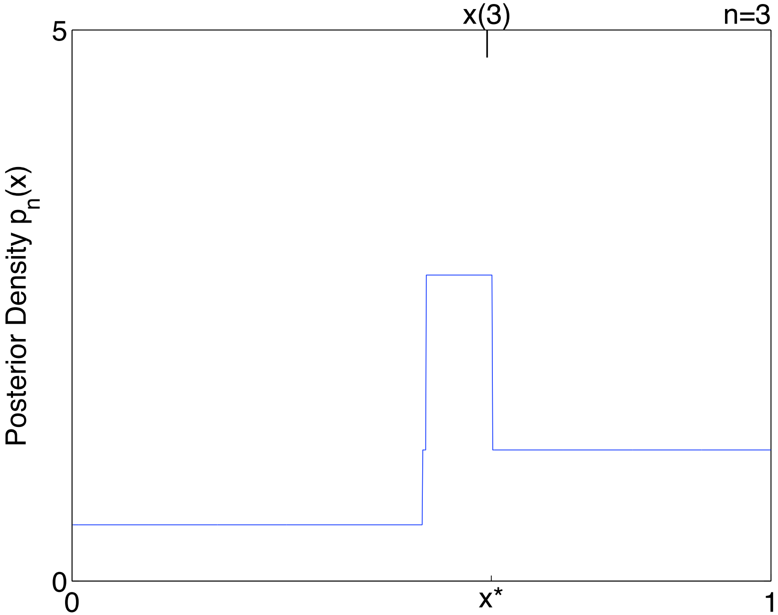

Example

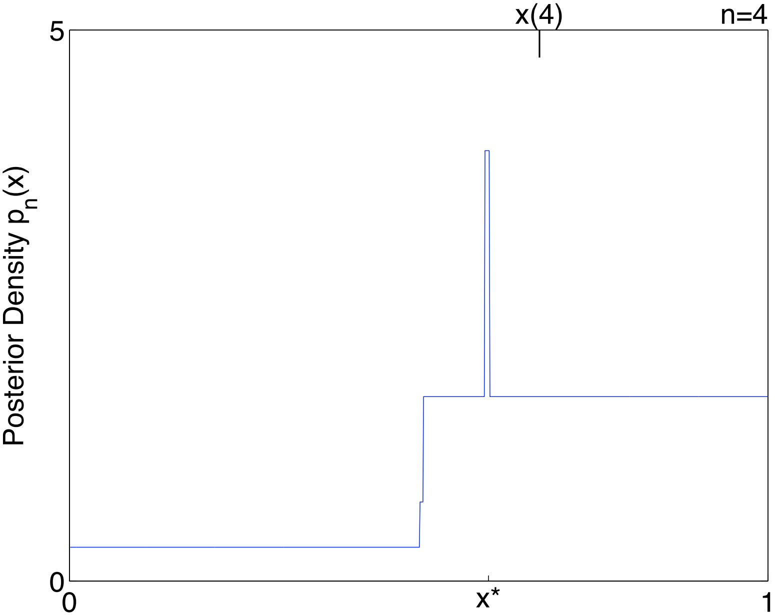

Example

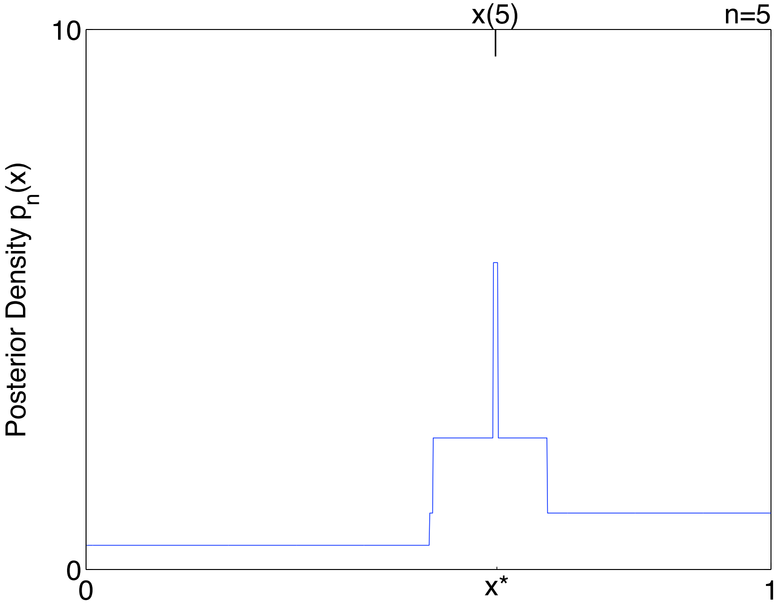

Example

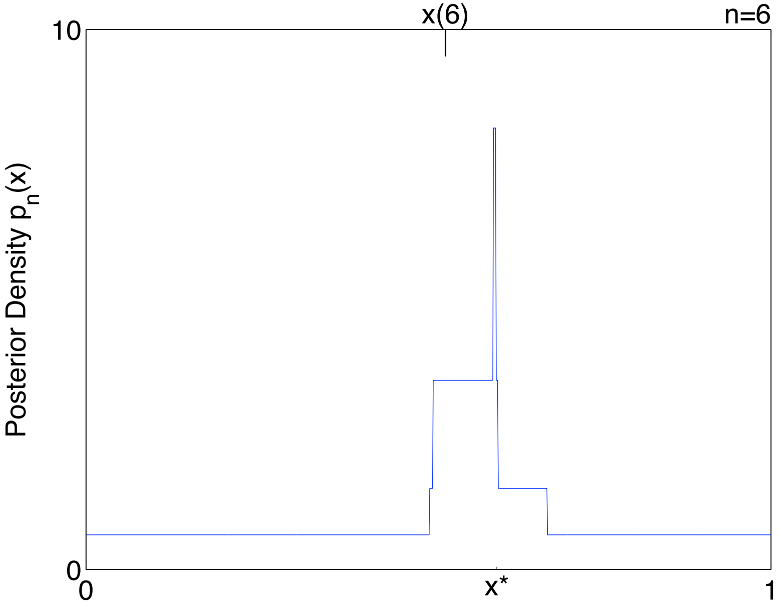

Example

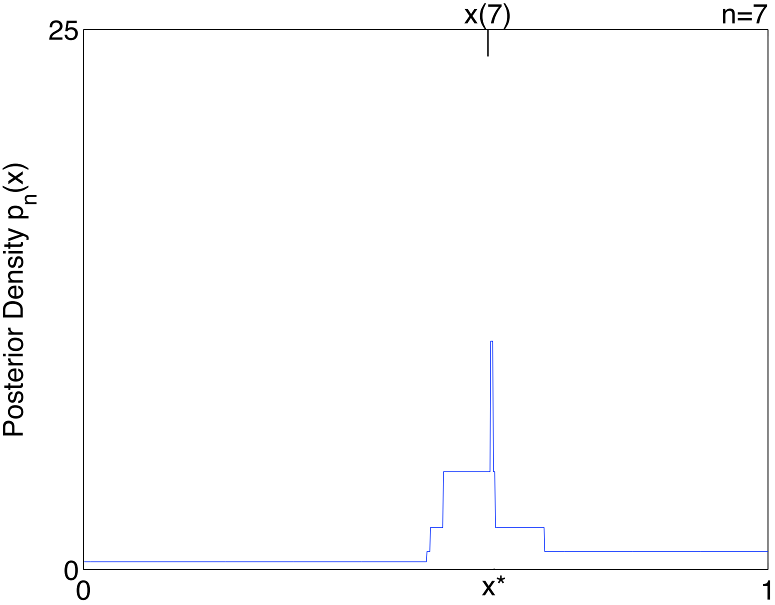

Example

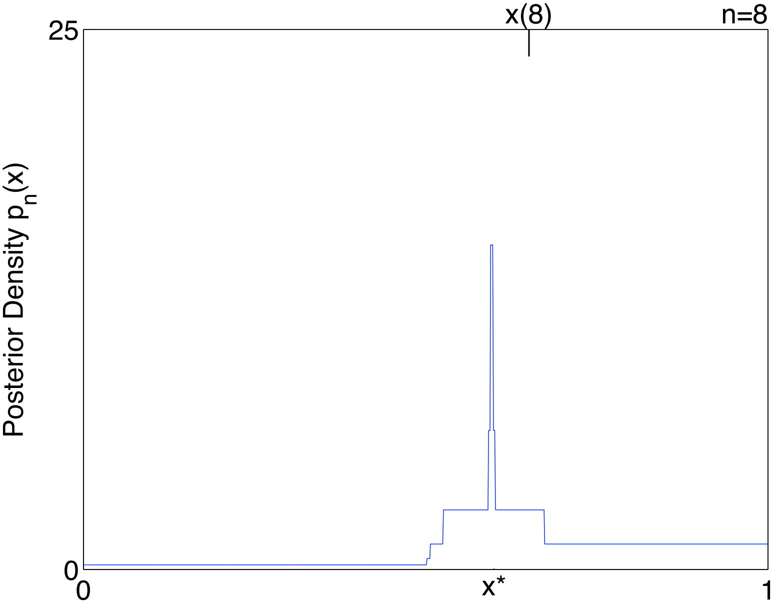

Example

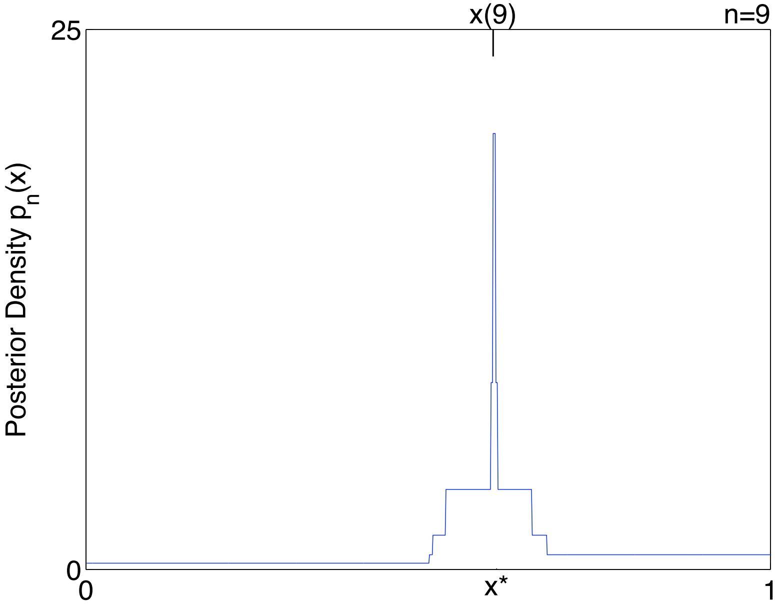

Example

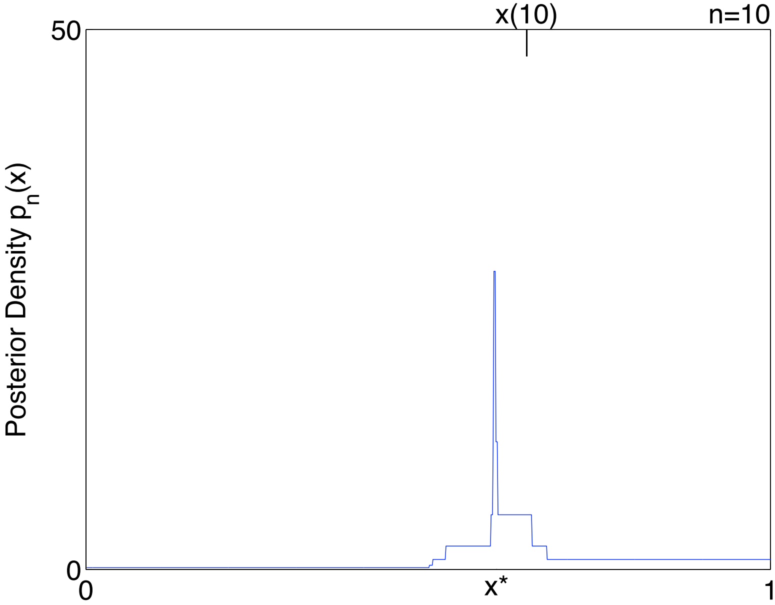

Example

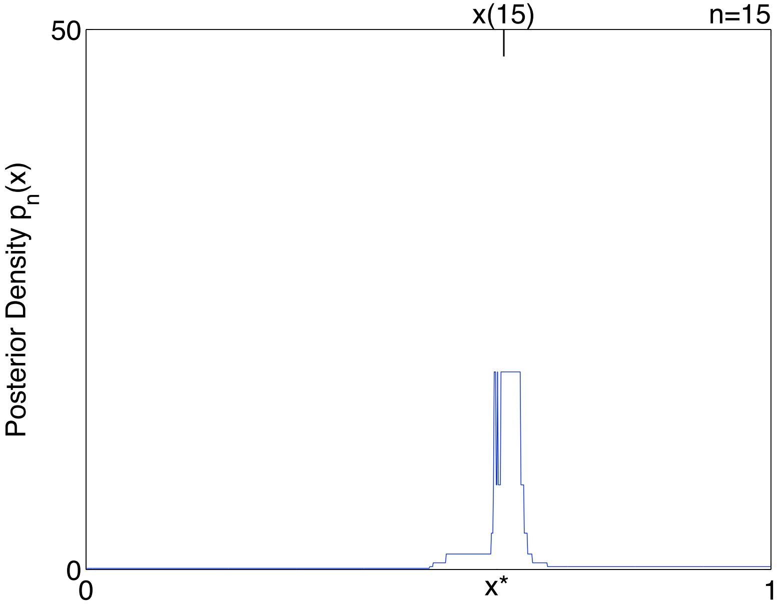

Example

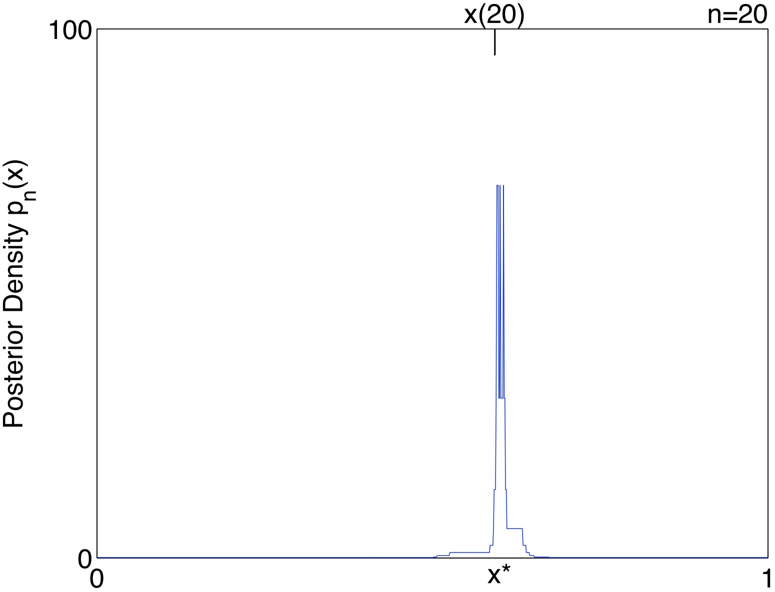

Example

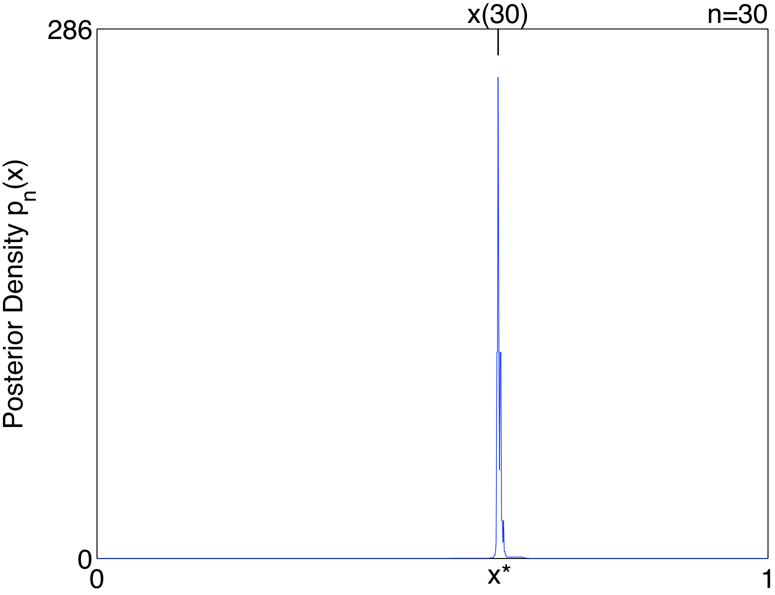

Consistency

- Allow $q(x)$ to vary with x.
- Continue to assume that updates are done with the correct $q(x)$.
- Then we have convergence to the root.

Theorem

$x_n \longrightarrow x_*$ almost surely as $n \longrightarrow \infty$.

Geometric Convergence

- q is fixed.
- Updates are done using q.
- Then x_n converges geometrically to the root x_*.

Theorem

Fix $\varepsilon > 0$ *and let* $A_\varepsilon = [x_* - \varepsilon, x_* + \varepsilon]$.
Let $\tau = \inf\{n : x_n \in A_\varepsilon\}$.
Let $c(q) = 1 + q\log_2(q) + (1-q)\log_2(1-q)$. *Then*

$$\mathbb{E}[\tau] \leq -\log(2\varepsilon)/c(q)$$

Unknown $q(x)$

- In practice, the probability of error $q(x)$ varies with x and is unknown.
- We can estimate $q(x_n)$ by performing multiple measurements at the same x_n, and then use this estimate in our update.
- Alternatively, we can sample sequentially to achieve an error probability $q(x_n)$ **bounded above by a constant**, call it $\widetilde{q}$.

Sampling with $q(x)$ bounded above by $\widetilde{q}$

How to sample sequentially to achieve bounded error probability $q(x_n) \leq \widetilde{q}$.

- We extend results from Siegmund 1985 for hypothesis testing of the mean of a normal random variable using curved boundaries.
- Fix x_n and sample a random walk that goes up when $y > 0$ and down when $y < 0$.

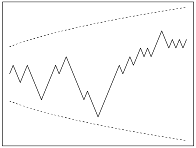

- If the random walk hits the top boundary first, this indicates $x_* > x_n$. If it hits the bottom, $x_* < x_n$. $q(x_n)$ is the error probability of this bit.
- Given any $0 < \widetilde{q} < \frac{1}{2}$, a test may be constructed whose $q(x_n) \leq \widetilde{q}$ whenever $P(y_n > 0) \neq \frac{1}{2}$.

Convergence Rate

- Allow $q(x)$ to vary with x.
- Updates are done with $\widetilde{q}$, which is an upper bound on $q(x_n)$.

Theorem

Fix $\varepsilon > 0$ *and let* $A_\varepsilon = [x_* - \varepsilon, x_* + \varepsilon]$.
Let $\tau = \inf\{n : x_n \in A_\varepsilon\}$.
Let $c(\widetilde{q}) = 1 + \widetilde{q}\log_2(\widetilde{q}) + (1-\widetilde{q})\log_2(1-\widetilde{q})$. *Then*

$$\mathbb{E}[\tau] \leq -\log(2\varepsilon)/c(\widetilde{q})$$

We still have geometric convergence in the number of different x_n sampled, but it may take more time to sample one x_n when x_n is closer to x_*.

Experimental Results (Simulation Problem)

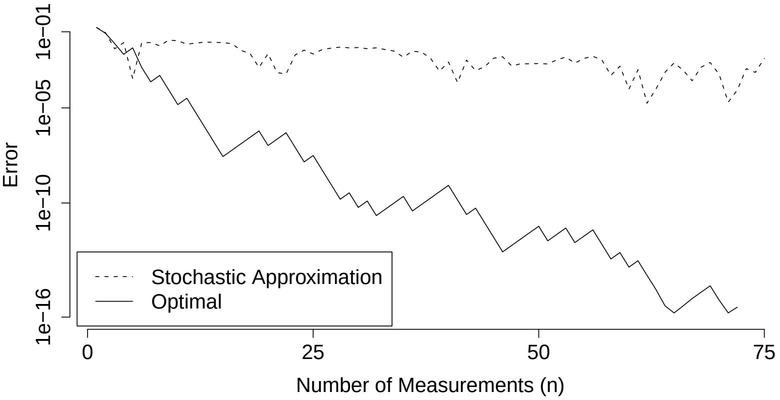

Performance on a problem with $f(x) = e^x - e^{1/3}$ and domain $[0,1]$. Error is $|x_n - x_*|$ on one sample path. Stochastic approximation used stepsize $1/n$.

Experimental Results (Laboratory Problem)

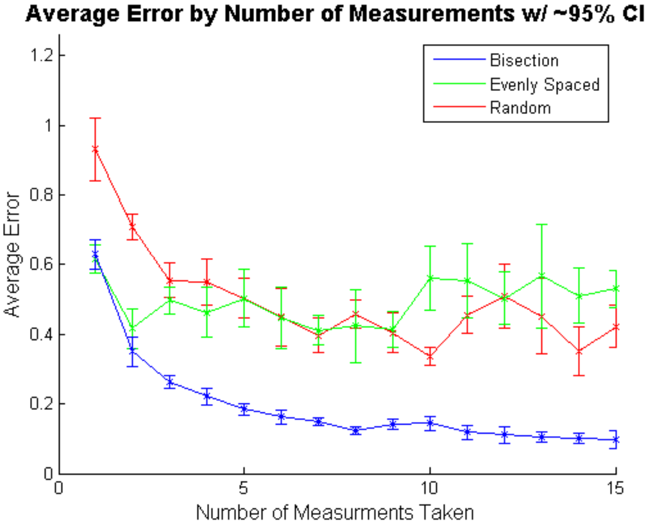

Joint work with Zachary Owen, Thorsten Joachims, and Rodrigo Bicalho.

Conclusion

- **We solved the dynamic program** for an idealized version of the stochastic root-finding problem.
- This provides a **Bayes-optimal** policy for this idealized problem.
- The resulting algorithm has **geometric convergence** to the root when the error probability is known.

Thank You