Simulation-based Approaches for Computing an Optimal Sequence of Adaptive Management Decisions

Robert M. Dorazio

U.S. Geological Survey

Florida Integrated Science Center

Gainesville, FL

Adaptive Management Conference Series
Laurel, MD
24–25 April 2003

Bayesian Inference, Prediction, and Decision Theory

Benefits:

- provides a coherent framework for sequentially updating beliefs (models and parameters) as new information is acquired through monitoring
- predicts consequences of future management actions while properly accounting for uncertainty in updated beliefs
- evaluates the expected consequences of future management actions in a manner that fully accounts for posterior uncertainty in beliefs and predictions
- (almost) places no limits on the complexity of models or utility functions, given modern computing capabilities

source: *Dorazio and Johnson (2003)*

Bayesian Inference, Prediction, and Decision Theory (cont'd)

Difficulties:

- Computing an optimal sequence of proposed management actions is computationally intensive
  - discrete decision spaces $\Rightarrow$ comparison of expected utilities of an enormous number of possible sequences of management actions.
  - For example, let
    $q =$ no. of possible actions taken at each decision time
    $\tau =$ no. of decisions to be made for each unit of observation
    $n =$ no. of units of observation

    Then, in general, $(q^\tau)^n =$ super-exponential no. of possible sequences of management actions

Algorithms for Computing Optimal Sequences of Management Actions

- Backward-induction/dynamic-programming (exploits recursion) *(Carlin et al. 1998, Hardwick and Stout 1999, 2002, Müller 1999, Müller et al. 2000)*
- Stochastic search algorithms
  - simulated annealing
  - genetic algorithms
- Reinforcement learning
- Others?

A Simple, yet Computationally Challenging, Example

Suppose we observe $n = 200$ units of observation initially (at $t = 0$), and half of them are in a favorable condition ($y_0 = 1$) while the other half are in an unfavorable condition ($y_0 = 0$).

At each decision time, one of 2 actions ("do nothing" or "manage") may be selected for each unit.

A cautious ecologist, being unsure about the potential benefits of management, randomly selects half of the units in favorable condition for management and leaves the other half unmanaged. The same allocation of actions is applied to the units in unfavorable condition.

Action	$y_0 = 1$	$y_0 = 0$
"do nothing"	50	50
"manage"	50	50

Suppose this initial allocation management actions is maintained for $T = 5$ years, and the condition of each unit is monitored in each of these years.

Data

Denote the management actions and conditions of units observed after $T = 5$ years as follows:

$$(\boldsymbol{X}_0, \boldsymbol{X}_1, \ldots, \boldsymbol{X}_{T-1}) = \ (n \times T) \text{ matrix } \boldsymbol{X}$$
$$(\boldsymbol{y}_0, \boldsymbol{y}_1, \ldots, \boldsymbol{y}_{T-1}, \boldsymbol{y}_T) = \ (n \times (T+1)) \text{ matrix } \boldsymbol{Y}$$

Utility

Denote the utility of a proposed sequence of τ management actions for $m \leq n$ units of observation as follows:

$$U(\tilde{\boldsymbol{X}}_T, \ldots, \tilde{\boldsymbol{X}}_{T+\tau-1}, \tilde{\boldsymbol{y}}_{T+1}, \ldots, \tilde{\boldsymbol{y}}_{T+\tau}) = \sum_{i=1}^{m} \sum_{t=T+1}^{T+\tau} U(\tilde{\boldsymbol{x}}_{i,t-1}, \tilde{y}_{i,t})$$

where $U(\tilde{\boldsymbol{x}}_{i,t-1}, \tilde{y}_{i,t})$ denotes the following state- and action-dependent utilities:

$\tilde{\boldsymbol{x}}_{t-1}$	$\tilde{y}_t = 1$	$\tilde{y}_t = 0$
"do nothing"	1.0	0.0
"manage"	0.8	-0.2

Objective

Find the sequence of management actions, $(\tilde{\boldsymbol{X}}_T, \ldots, \tilde{\boldsymbol{X}}_{T+\tau-1})$, that maximizes the expected utility

$$\begin{aligned}
&\overline{U}(\tilde{\boldsymbol{X}}_T, \ldots, \tilde{\boldsymbol{X}}_{T+\tau-1} \mid \boldsymbol{Y}, \boldsymbol{X}) \\
&\quad = \mathrm{E}_{(\tilde{\boldsymbol{y}}_{T+1},\ldots,\tilde{\boldsymbol{y}}_{T+\tau} \mid \tilde{\boldsymbol{X}}_T,\ldots,\tilde{\boldsymbol{X}}_{T+\tau-1},\, \boldsymbol{Y}, \boldsymbol{X})} \left[U(\tilde{\boldsymbol{X}}_T, \ldots, \tilde{\boldsymbol{X}}_{T+\tau-1},\, \tilde{\boldsymbol{y}}_{T+1}, \ldots, \tilde{\boldsymbol{y}}_{T+\tau}) \right]
\end{aligned}$$

For a sequence of $\tau = 5$ decisions and m units of observation, there are $(2^5)^m$ possibilities to consider!

$m =$	1	2	3	4	$\cdots$	200
	32	1,024	32,768	1,048,576	$\cdots$	$\approx 1 \times 10^{301}$

Recall: we have $T = 5$ years of observations for each of $n = 200$ units that may be used to predict the consequences of future management actions.

$$\boldsymbol{Y} = \begin{pmatrix} 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 \\ \vdots & & & & & \\ 1 & 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 \\ \vdots & & & & & \\ 0 & 0 & 0 & 1 & 1 & 0 \end{pmatrix}$$

Summary of Observations

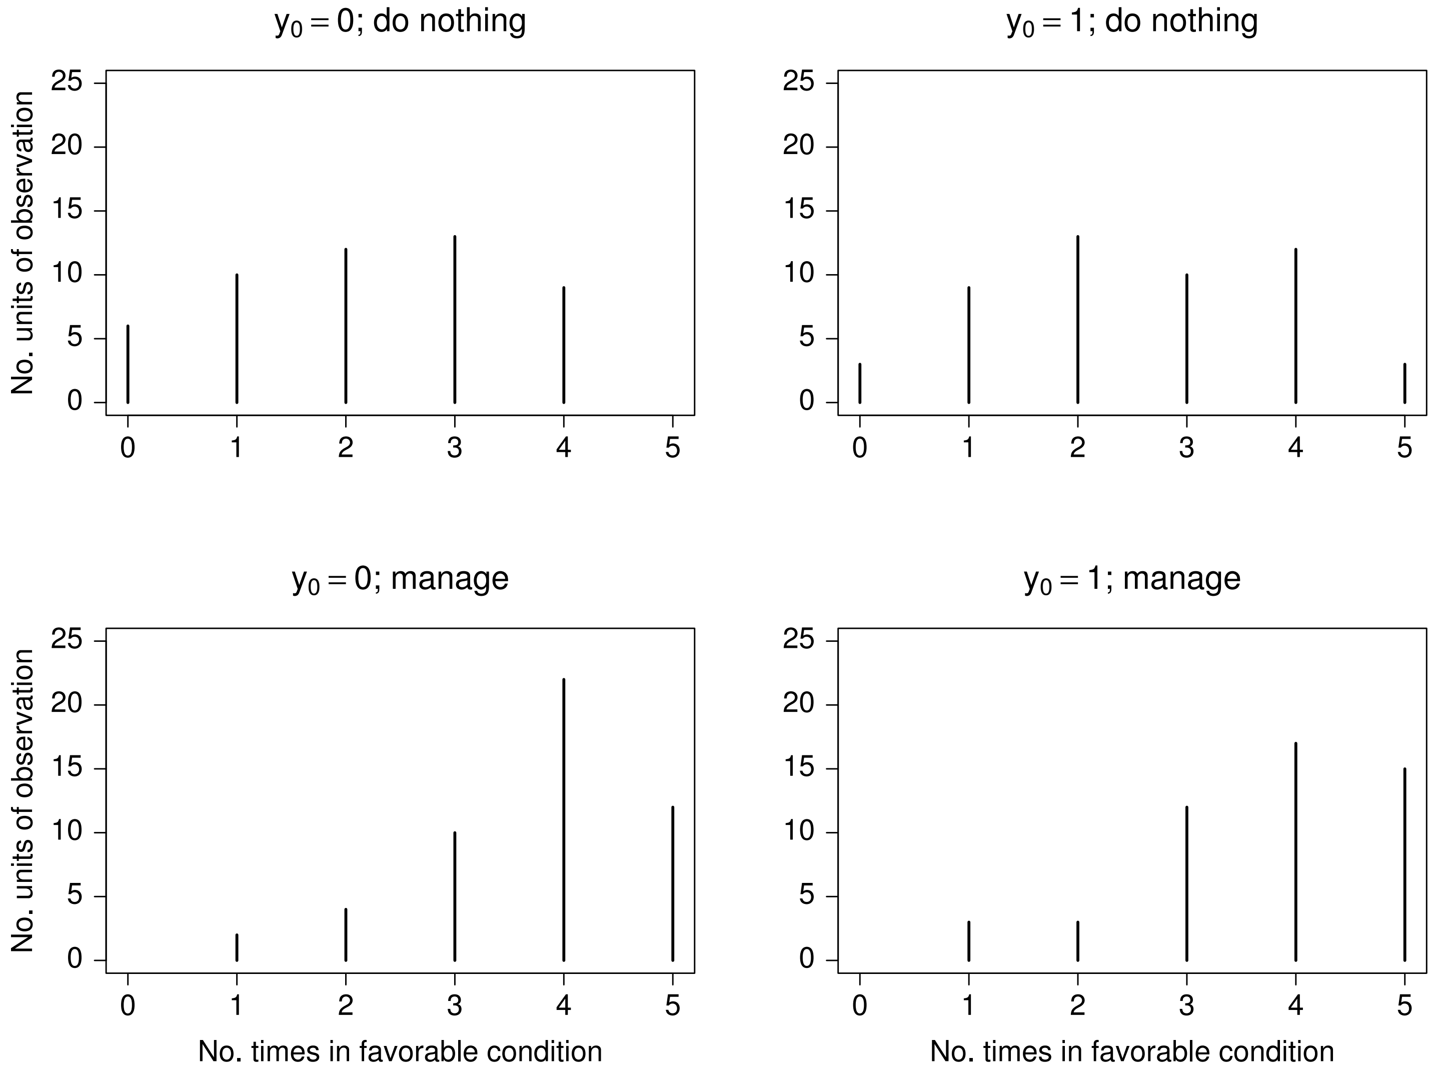

AR(1) Logistic Regression Model (Bonney 1987)

$$f(y_t \mid \boldsymbol{x}_{t-1}, y_{t-1}, \boldsymbol{\alpha}, \beta) = \frac{\exp(y_t \eta_{t-1})}{1 + \exp(\eta_{t-1})}$$

$$\eta_{t-1} = \boldsymbol{x}'_{t-1}\boldsymbol{\alpha} + (2y_{t-1} - 1)\beta$$

where $\boldsymbol{x}'_{t-1} \in \{(1,0),(0,1)\}$ and $\boldsymbol{x}'_{t-1}\boldsymbol{\alpha} \in \{\alpha_1, \alpha_2\}$

State- and Action-Dependent Transition Probabilities

Action		$y_t = 1$	$y_t = 0$
"do nothing"	$y_{t-1} = 1$	$\frac{\exp(\alpha_1+\beta)}{1+\exp(\alpha_1+\beta)}$	$\frac{1}{1+\exp(\alpha_1+\beta)}$
$\Updownarrow$			
$\boldsymbol{x}'_{t-1} = (1,0)$	$y_{t-1} = 0$	$\frac{\exp(\alpha_1-\beta)}{1+\exp(\alpha_1-\beta)}$	$\frac{1}{1+\exp(\alpha_1-\beta)}$
"manage"	$y_{t-1} = 1$	$\frac{\exp(\alpha_2+\beta)}{1+\exp(\alpha_2+\beta)}$	$\frac{1}{1+\exp(\alpha_2+\beta)}$
$\Updownarrow$			
$\boldsymbol{x}'_{t-1} = (0,1)$	$y_{t-1} = 0$	$\frac{\exp(\alpha_2-\beta)}{1+\exp(\alpha_2-\beta)}$	$\frac{1}{1+\exp(\alpha_2-\beta)}$

Bayesian Estimation and Inference

Likelihood

$$f(\boldsymbol{Y} \mid \boldsymbol{X}, \boldsymbol{\alpha}, \beta) = \prod_{i=1}^{n} \prod_{t=1}^{T} f(y_{i,t} \mid \boldsymbol{x}_{i,t-1}, y_{i,t-1}, \boldsymbol{\alpha}, \beta)$$

Prior, $\pi(\boldsymbol{\alpha}, \beta)$

$$\begin{aligned} \operatorname{logit}^{-1}(\alpha_j) &\sim \text{Uniform}(0,1) \\ \beta &\sim \text{Exponential}(10) \end{aligned}$$

Posterior

$$\pi(\boldsymbol{\alpha}, \beta \mid \boldsymbol{Y}, \boldsymbol{X}) \propto f(\boldsymbol{Y} \mid \boldsymbol{X}, \boldsymbol{\alpha}, \beta)\ \pi(\boldsymbol{\alpha}, \beta)$$

Bayesian Prediction and Expected Utility

Posterior predictive distribution

$$p(\tilde{\boldsymbol{y}}_{T+1}, \ldots, \tilde{\boldsymbol{y}}_{T+\tau} \mid \tilde{\boldsymbol{X}}_T, \ldots, \tilde{\boldsymbol{X}}_{T+\tau-1}, \boldsymbol{Y}, \boldsymbol{X}) = \int \left[\prod_{i=1}^{m} \prod_{t=T+1}^{T+\tau} f(\tilde{y}_{i,t} \mid \tilde{\boldsymbol{x}}_{i,t-1}, \tilde{y}_{i,t-1}, \boldsymbol{\alpha}, \beta)\right] \pi(\boldsymbol{\alpha}, \beta \mid \boldsymbol{Y}, \boldsymbol{X}) \, \partial\boldsymbol{\alpha} \, \partial\beta$$

Expected utility

$$\begin{aligned} \overline{U}(\tilde{\boldsymbol{X}}_T, \ldots, \tilde{\boldsymbol{X}}_{T+\tau-1} \mid \boldsymbol{Y}, \boldsymbol{X}) &\doteq \frac{1}{R} \sum_{r=1}^{R} \left[U(\tilde{\boldsymbol{X}}_T, \ldots, \tilde{\boldsymbol{X}}_{T+\tau-1}, \tilde{\boldsymbol{y}}^{(r)}_{T+1}, \ldots, \tilde{\boldsymbol{y}}^{(r)}_{T+\tau})\right] \\ &\doteq \frac{1}{R} \sum_{r=1}^{R} \left[\sum_{i=1}^{m} \sum_{t=T+1}^{T+\tau} U(\tilde{\boldsymbol{x}}_{i,t-1}, \tilde{y}^{(r)}_{i,t})\right] \end{aligned}$$

Simulated Annealing Algorithm

At kth iteration of algorithm, let

$$\tilde{\boldsymbol{X}}^{(k)} = (\tilde{\boldsymbol{X}}_T^{(k)}, \ldots, \tilde{\boldsymbol{X}}_{T+\tau-1}^{(k)})$$

$$\overline{U}(\tilde{\boldsymbol{X}}^{(k)}) = \overline{U}(\tilde{\boldsymbol{X}}_T^{(k)}, \ldots, \tilde{\boldsymbol{X}}_{T+\tau-1}^{(k)} \mid \boldsymbol{Y}, \boldsymbol{X})$$

Step 1 Randomly select an initial sequence $\tilde{\boldsymbol{X}}^{(0)}$ and evaluate $\overline{U}(\tilde{\boldsymbol{X}}^{(0)})$

Step 2 Assign an initial annealing "temperature" $T^{(0)} \gg \overline{U}(\tilde{\boldsymbol{X}}^{(0)})$

Step 3 Randomly select a new sequence $\tilde{\boldsymbol{X}}^{*}$ in the vicinity of $\tilde{\boldsymbol{X}}^{(k)}$ and evaluate $\overline{U}(\tilde{\boldsymbol{X}}^{*})$

- Randomly select one of the τ decision times for each of the m units
- At these decision times, randomly select one of the possible management actions

Step 4 If $\overline{U}(\tilde{\boldsymbol{X}}^{*}) \geq \overline{U}(\tilde{\boldsymbol{X}}^{(k)})$, accept $\tilde{\boldsymbol{X}}^{*}$ as new value of $\tilde{\boldsymbol{X}}^{(k)}$; otherwise accept $\tilde{\boldsymbol{X}}^{*}$ with probability $\exp\left((\overline{U}(\tilde{\boldsymbol{X}}^{*}) - \overline{U}(\tilde{\boldsymbol{X}}^{(k)}))/T^{(k)}\right)$

Simulated Annealing Algorithm (continued)

Step 5 Repeat steps 3–4, $N = 500$ times and keep track of the number of s successful moves to a new point in the decision space.

Step 6 If $s > 0$, slowly decrease annealing "temperature" $T^{(k+1)} = \rho T^{(k)}$ (say, $\rho = 0.9$) and go to step 3;
otherwise stop (i.e., maximum has been found at current temperature).

Posterior Distribution of Model Parameters

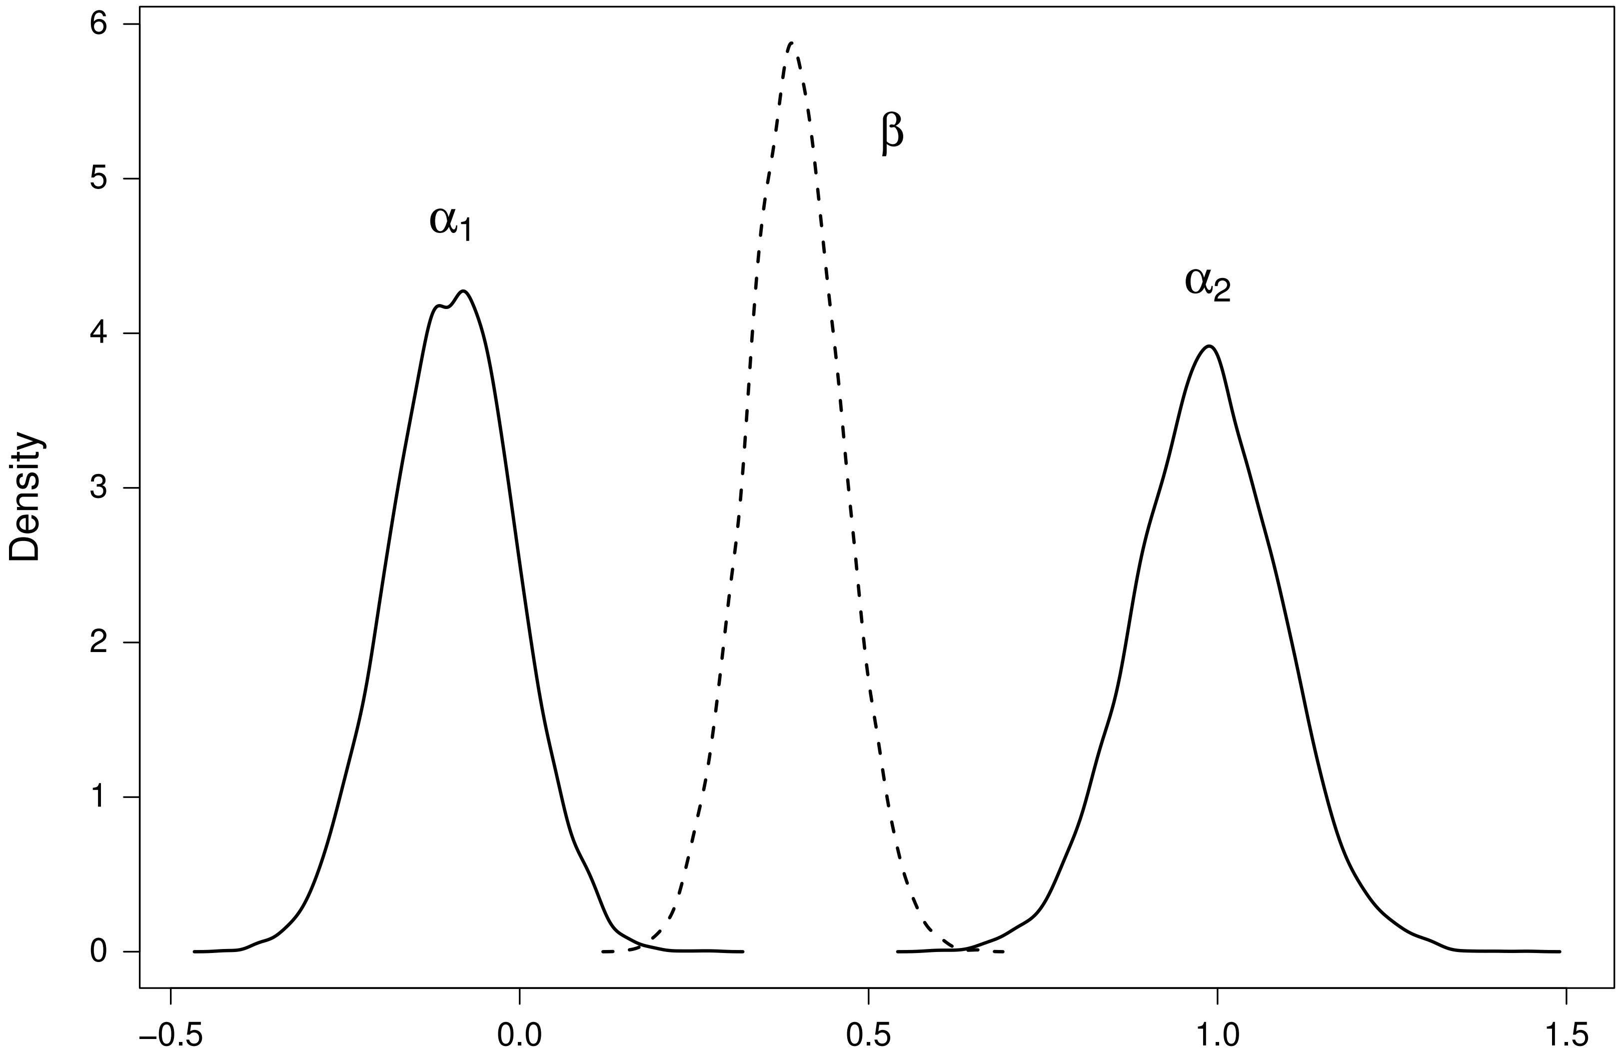

Expected Utilities for a Single Unit of Observation

Given $y_5 = 0$, optimal sequence of management actions is $\tilde{\boldsymbol{X}} = (1, 1, 1, 1, 1)$

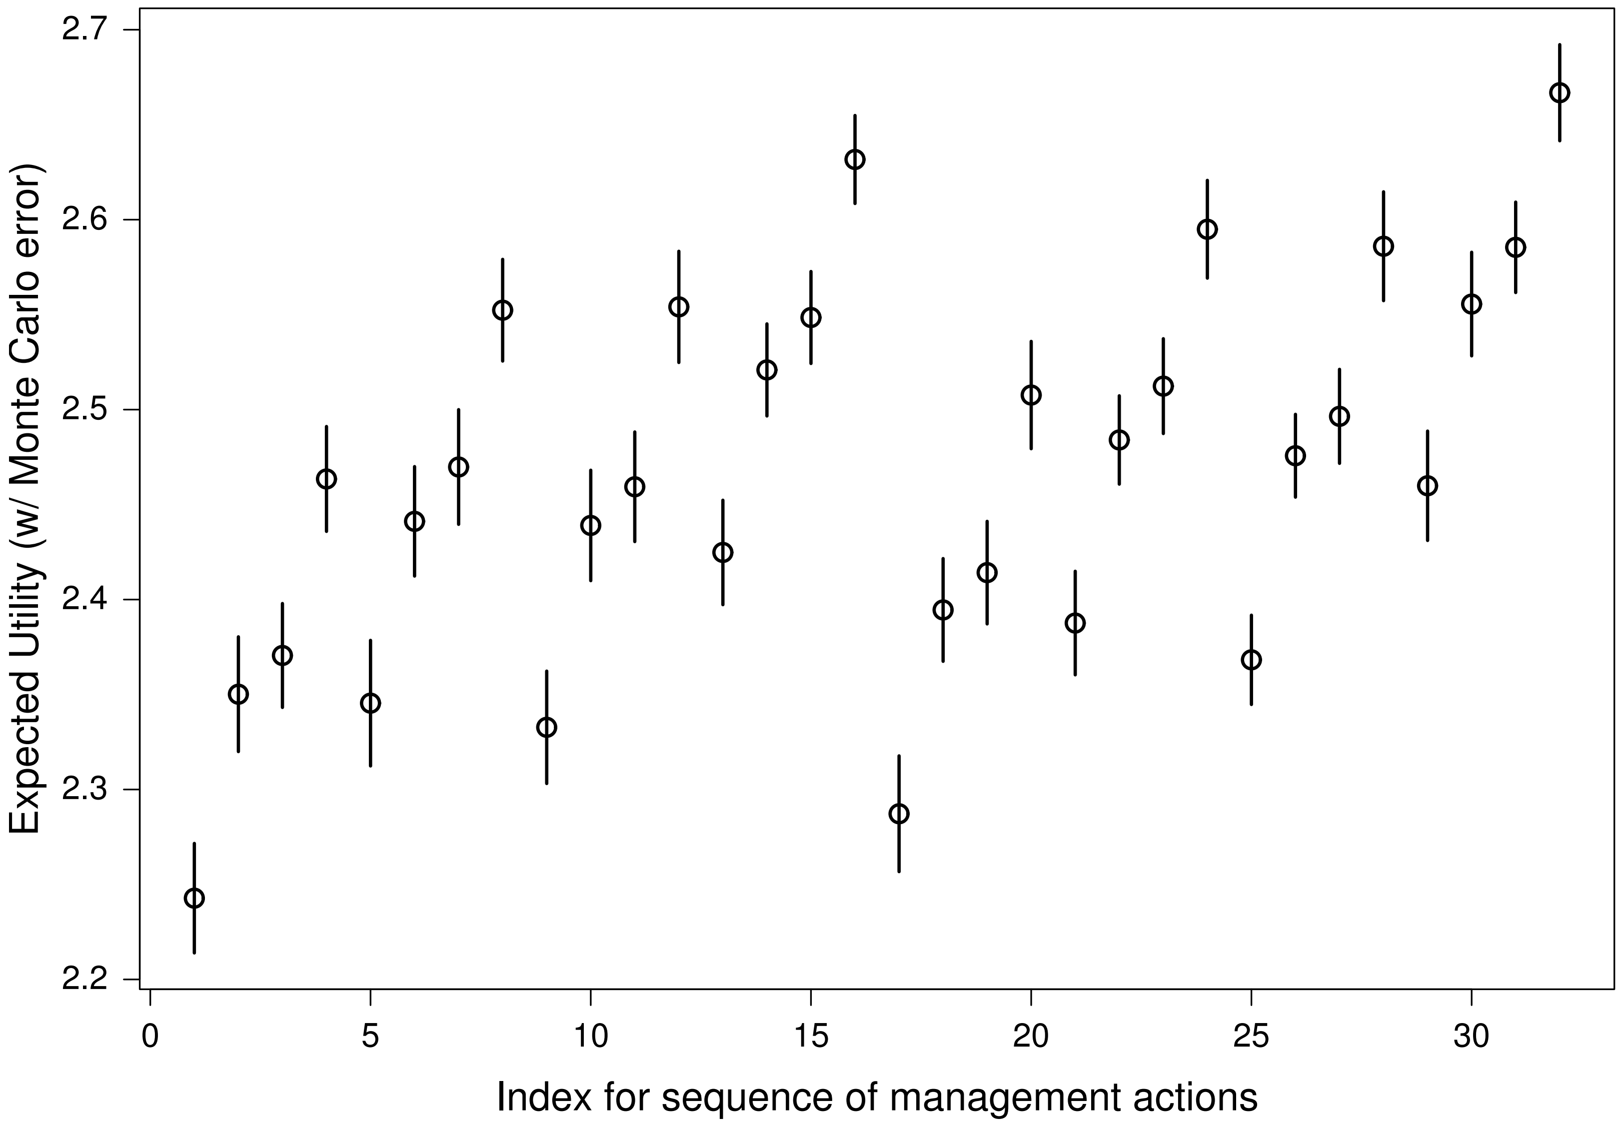

Expected Utilities for 2 Units of Observation

Given $\boldsymbol{y}_5 = (0, 1)'$, optimal sequences of management actions are

$$\tilde{\boldsymbol{X}} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix}$$

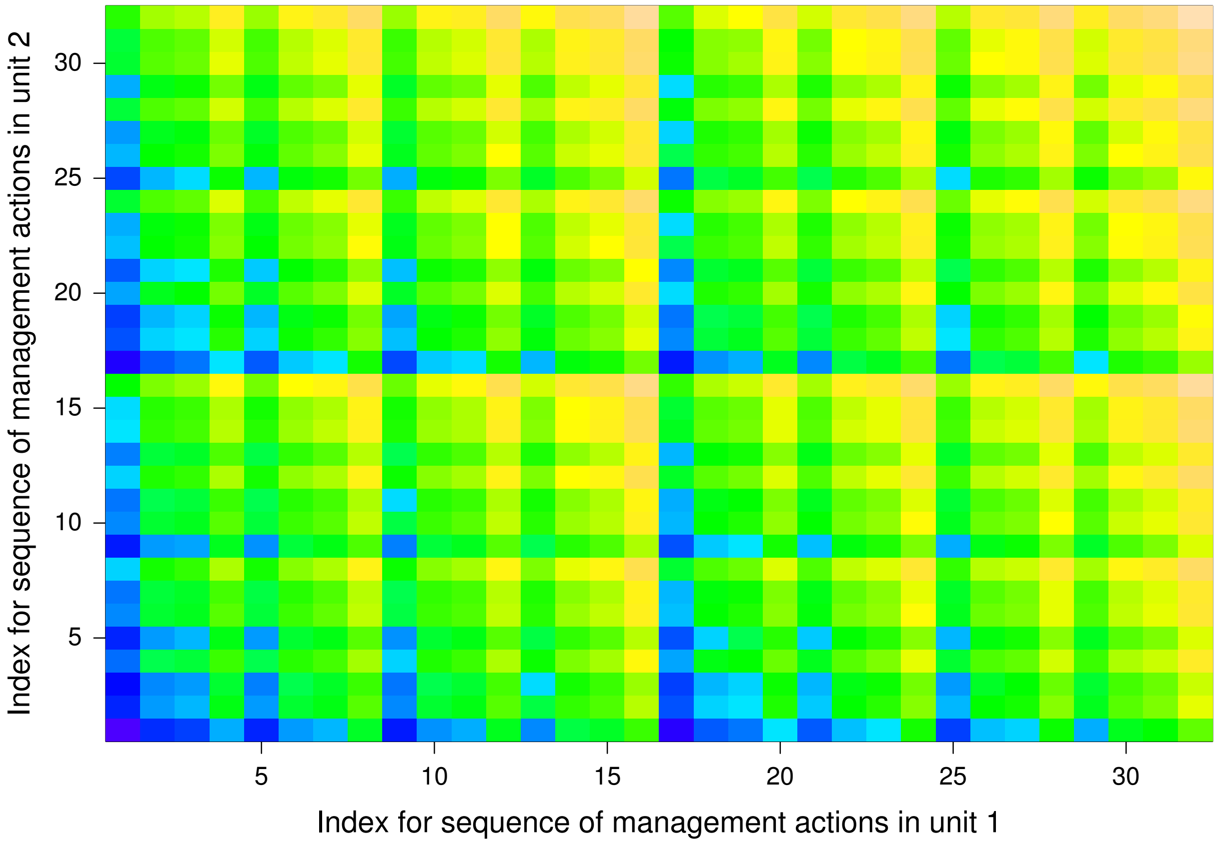

Summary of a Second Set of Observations

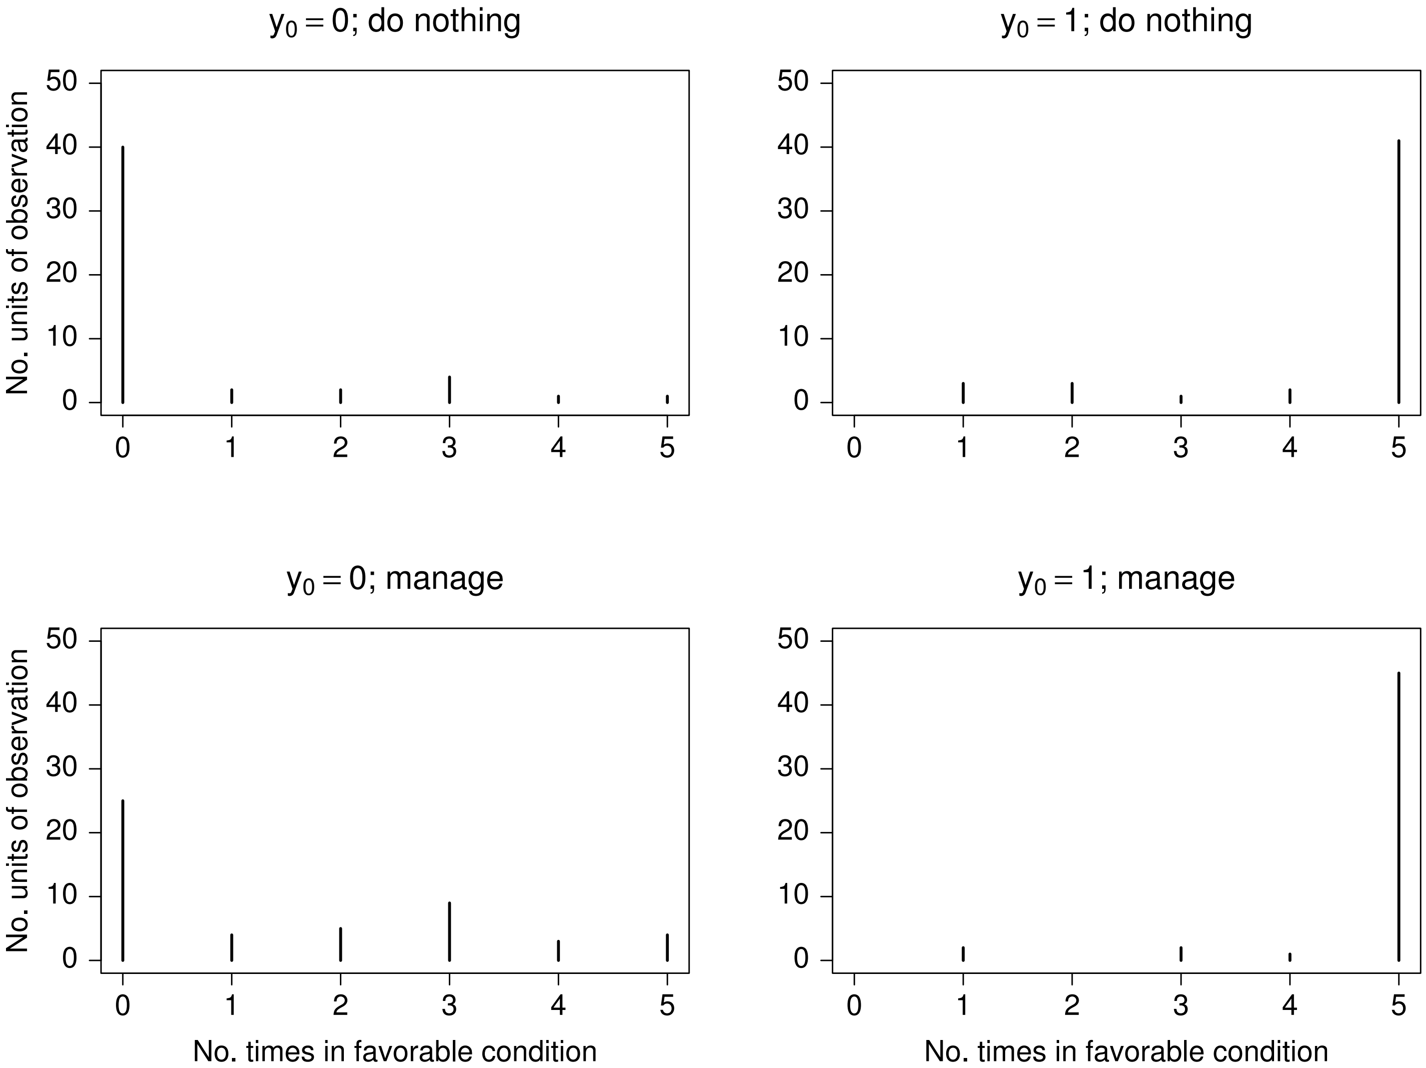

Posterior Distribution of Model Parameters

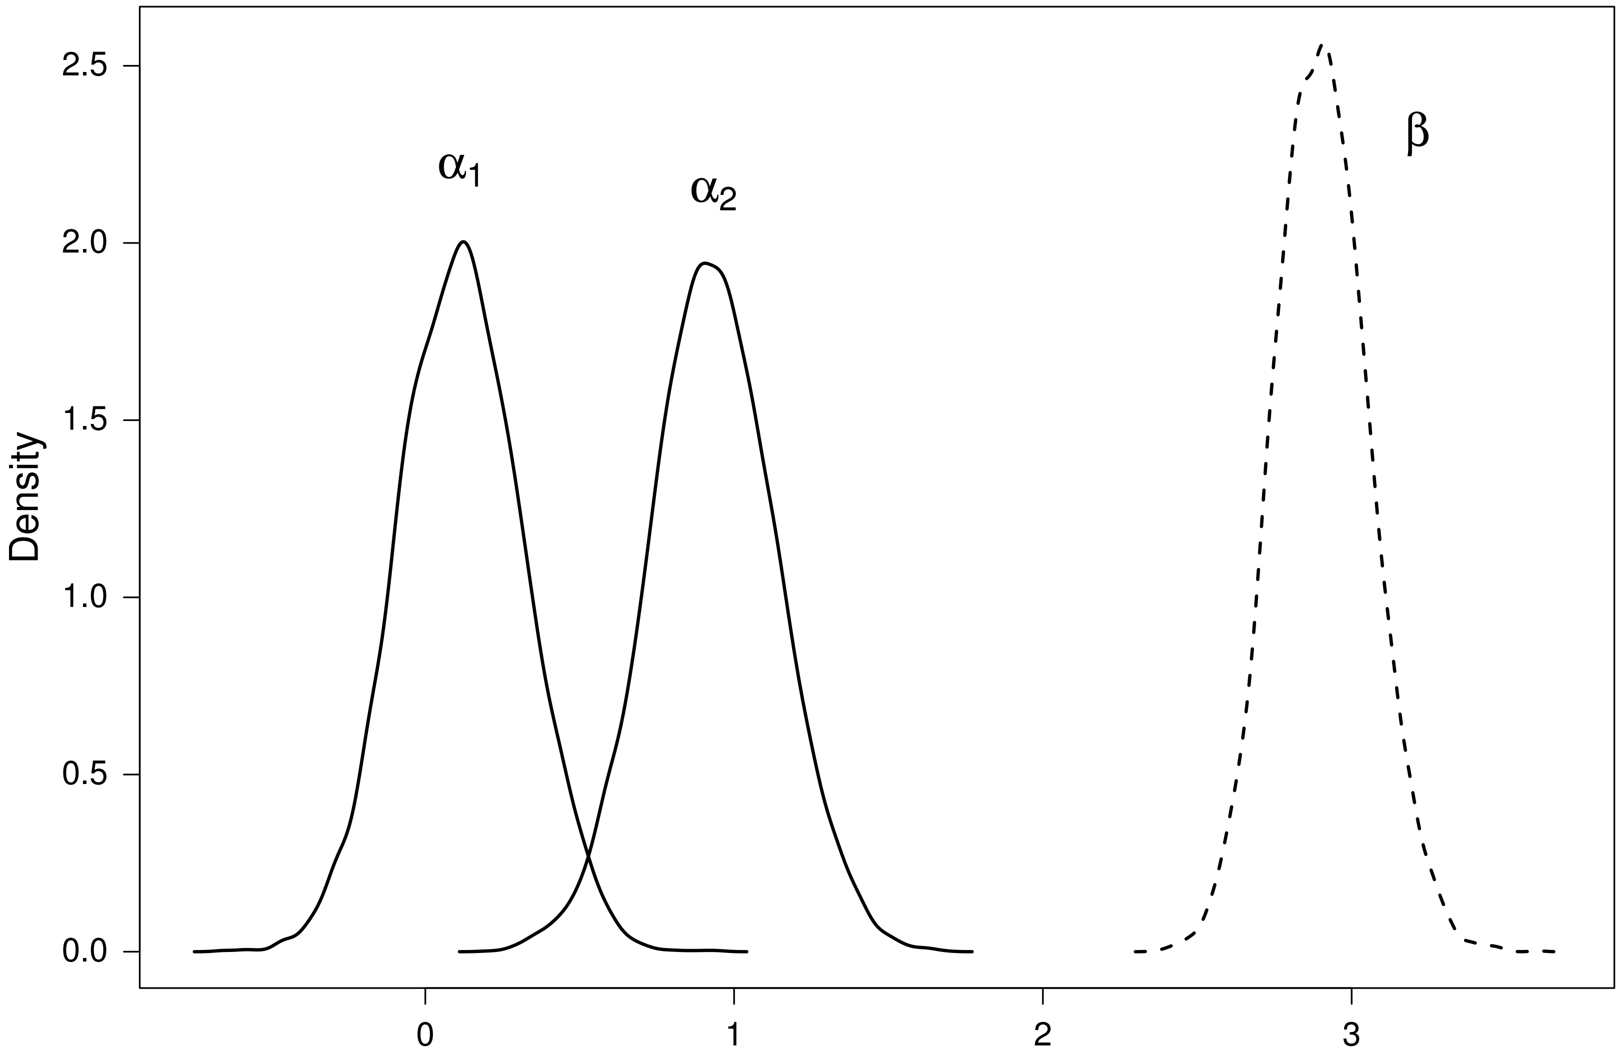

Expected Utilities for a Single Unit of Observation

Given $y_5 = 0$, optimal sequence of management actions is $\tilde{\boldsymbol{X}} = (1, 1, 0, 0, 0)$

Given $y_5 = 1$, optimal sequence of management actions is $\tilde{\boldsymbol{X}} = (0, 0, 0, 0, 0)$

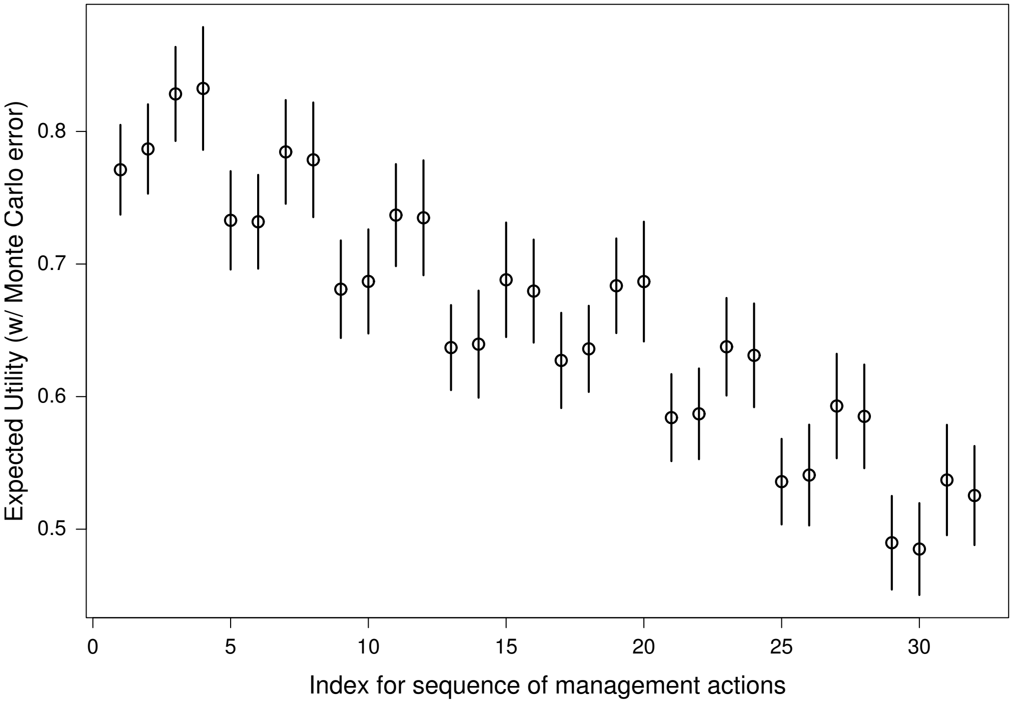

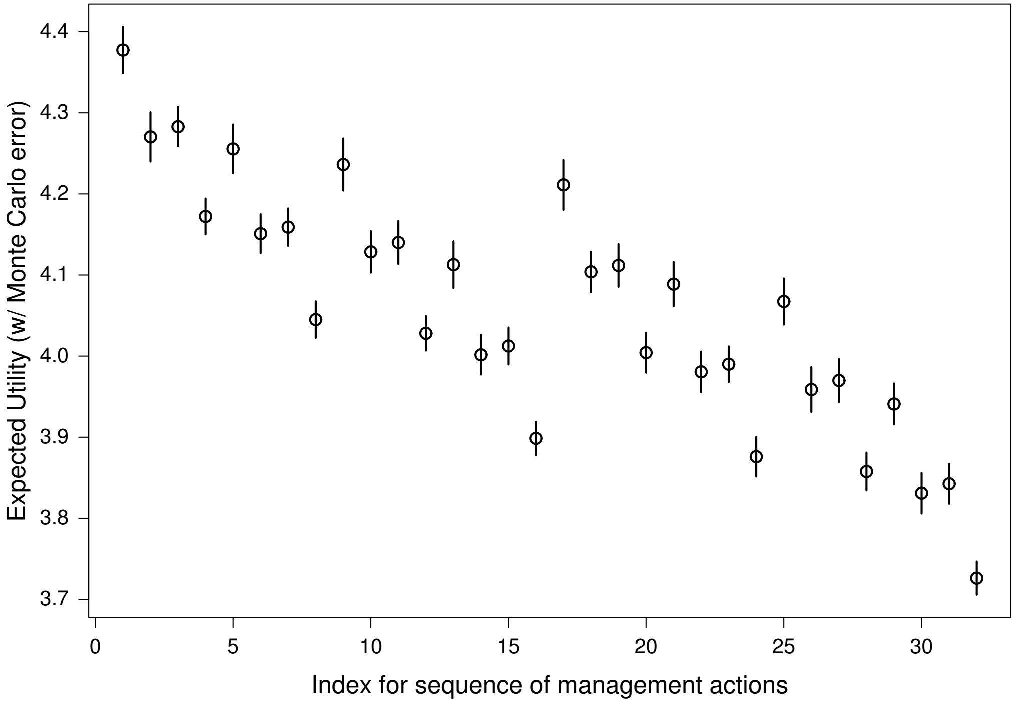

Effect of Sample Size on Convergence/Annealing Time

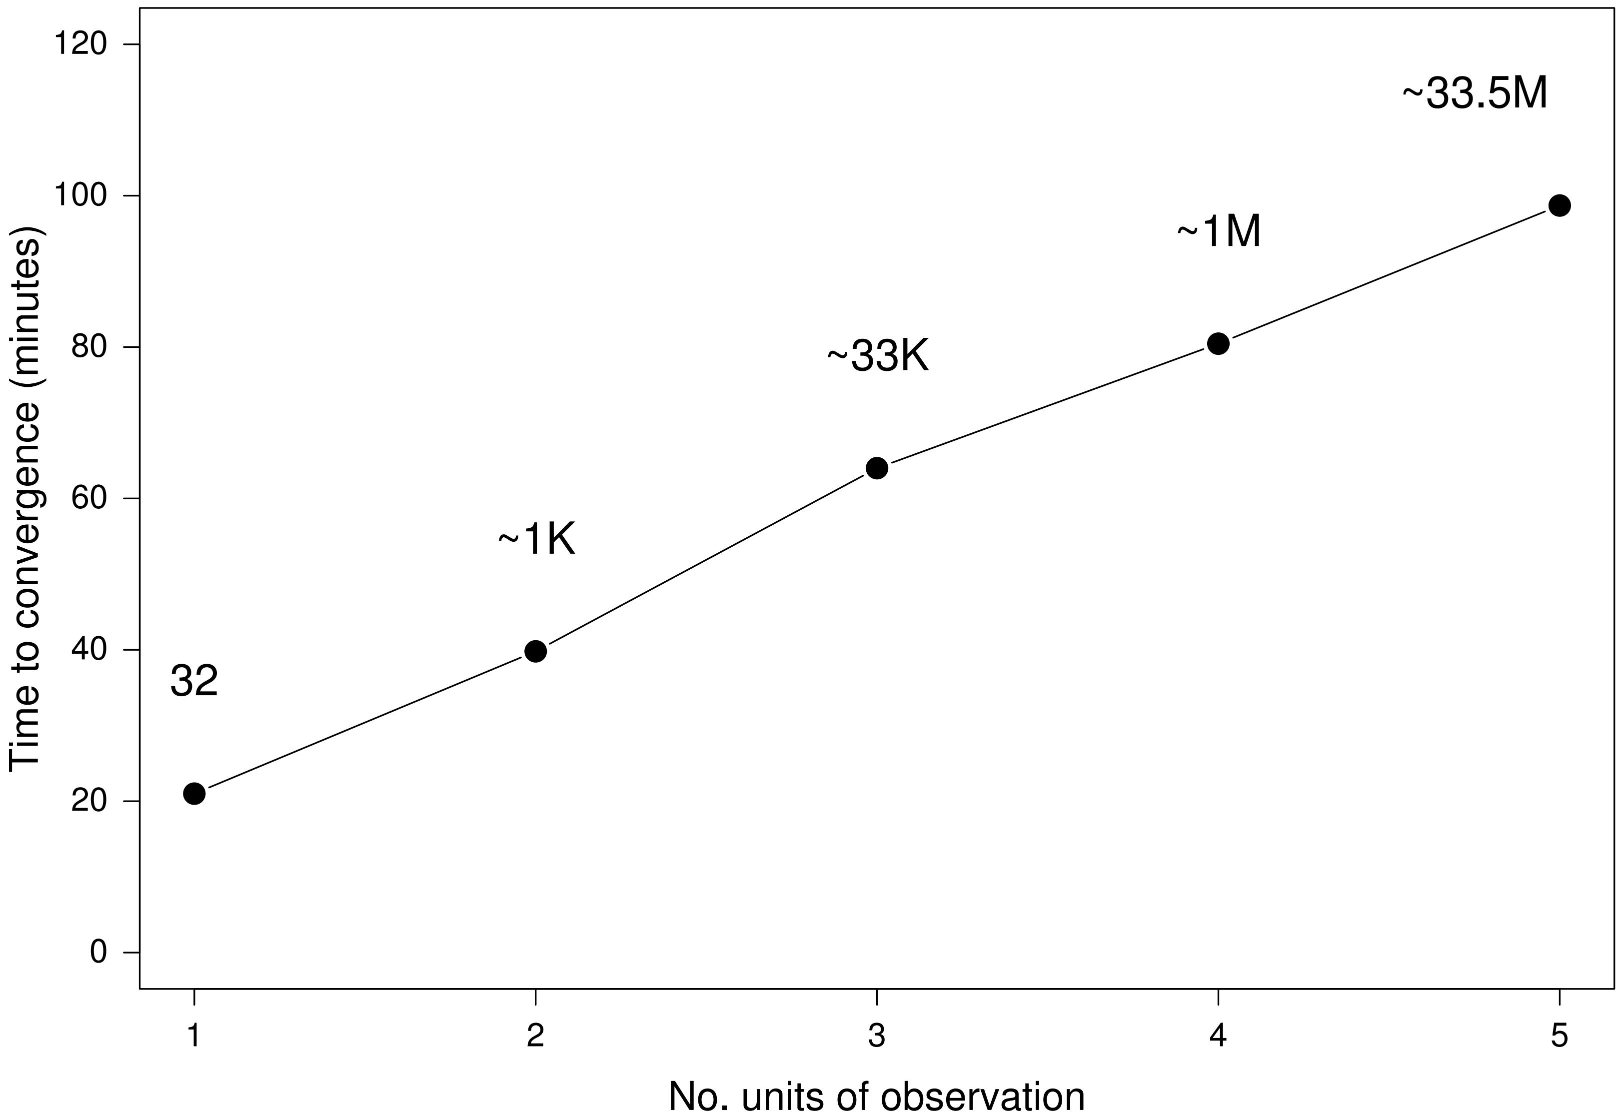

Tentative Conclusions and Suggestions for Additional Research

- Still more work to be done, but simulated annealing appears to be computationally feasible for optimization problems of "moderate" size
- How large is "moderate"?
- Performance comparisons are needed: DP vs SA vs GA vs RL
- Can additional gains in computational efficiency be achieved by combining stochastic search algorithms and backward induction (in problems/models where this is possible)?
- We have some powerful tools. Need to demonstrate their utility in *real problems*:
  - dependence *among* units of observation (e.g., spatially correlated observations)
  - alternative forms of temporal dependence *within* units of observation (e.g., long "memory " in communities of plant species)
  - errors in sampling, measurement, or in application of management actions
  - complex utility functions that include predictions of observable system state and model parameters (e.g., active adaptive management problems)