Learning in Adaptive Management

Chapter 7

Developed by: William L. Kendall, James D. Nichols, Clinton T. Moore, and G. Scott Boomer

<u>Session Objectives:</u> **By the end of this session, participants will be able to:**

- Discuss what learning means in Adaptive Resource Management (ARM)
- Describe how learning is accomplished

Outline

- Learning
  - What does it mean in ARM?
  - How is it done?

- What affects rate of learning?
  - Models, monitoring, approach to optimization

What is learning in an uncertain world?

- Dictionary definitions usually include “acquiring knowledge”
- Scientific definition might include “accumulation of faith (or lack of faith) associated with the predictions of competing hypotheses and their corresponding models”
- Science is a progressive endeavor that depends on learning (e.g., Descartes 1637)

Models and Learning

- Basic criterion by which a management model is judged is its ability to predict system response to management actions
- In case of multiple discrete models: develop model "weights" reflecting relative degrees of faith in the models of the model set

- For a given model set
  - Weights assigned to each model add to 1.0 (thus relative credibility)
  - Models with higher weight have greater credibility and will have more influence over future management decisions
  - If a robust predictive model is in the set its weight should go to 1.0 over time.

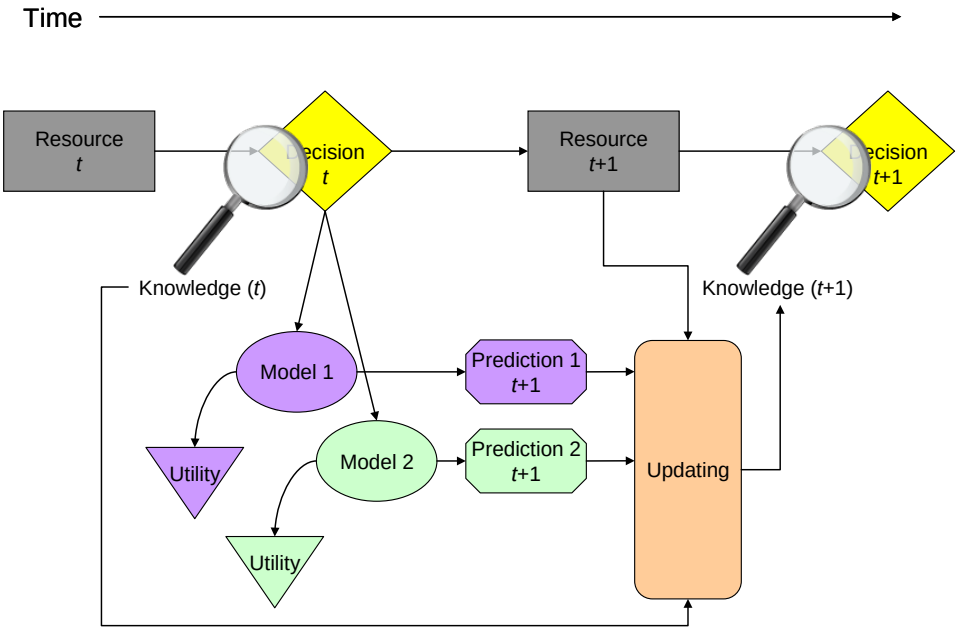

Slide by Clint Moore

What is learning in ARM?

- Reduction of structural uncertainty; i.e., discriminating among competing models of system response to management actions

- Accomplished by comparing model-based predictions against estimates of state variables and rate parameters (from monitoring program)

Why bother to learn in ARM?

- Structural uncertainty frequently reduces returns that are possible for a managed system

- For any system, this reduction can be assessed (EVPI)

- If EVPI is large, then learning is important, as greater returns can be realized if uncertainty is reduced

How does learning occur within ARM?

- For a given action, predictions made under each model

- The system response of the implemented decision is monitored

- Model weights are updated via Bayes' Formula

Initial weight values

- Subjectively
  - Politically
  - Based on expert opinion

- Based on historical data, e.g.,
  - AIC weights (Burnham and Anderson 2002)
  - Pick previous date, start with equal weights, and update to present time

Weights updated as function of

- The current weight (*prior probability*)

- New information (i.e., the difference between model predictions and what actually occurs, based on monitoring results).

- The new weight is called a *posterior probability*

Bayes' Formula

In words:

New weight of *model i* $\propto$

(Old weight of model i) * (*likelihood of new data according to model i*)

Formula:

$$p_{t+1}(\text{model } i \mid \text{response}_{t+1}) =$$

$$\frac{p_t(\text{model } i)\, P(\text{response}_{t+1} \mid \text{model } i)}{\sum_j p_t(\text{model } j)\, P(\text{response}_{t+1} \mid \text{model } j)}$$

Process advances learning when

- A good approximating model is in the model set (i.e., a model that predicts well across the state space)

- Predictions from each model fairly represent the idea that generated them

- An adequate monitoring program is in place for model comparison/discrimination

Model predictions should:

- Be unbiased under the ecological hypothesis they represent
  - Bias could change direction of weight changes and lead to erroneous conclusion of poor predictive ability

- Include all pertinent uncertainties
  - Model-based stochastic variation,
  - Parametric uncertainty – sampling variation due to estimation.
  - Partial observability of resulting state (monitoring bias/imprecision)

Real World Examples – two models

- Compensatory versus additive mortality in hunted species
- Density independence/dependence in recruitment, survival, abundance
- Wood thrush abundance as a linear vs. logistic function of habitat quantity
- Shorebird use of impoundment dependent on percent that is mudflat?
- Beaver trapping effort as a function of gas prices?

Generic Example – two models

- Assume equal priors
  - each model gets weight of 0.50
- Compute posteriors (i.e., update weights) based on comparison of *predicted* observed state (e.g., population size) with resulting *observed* state.

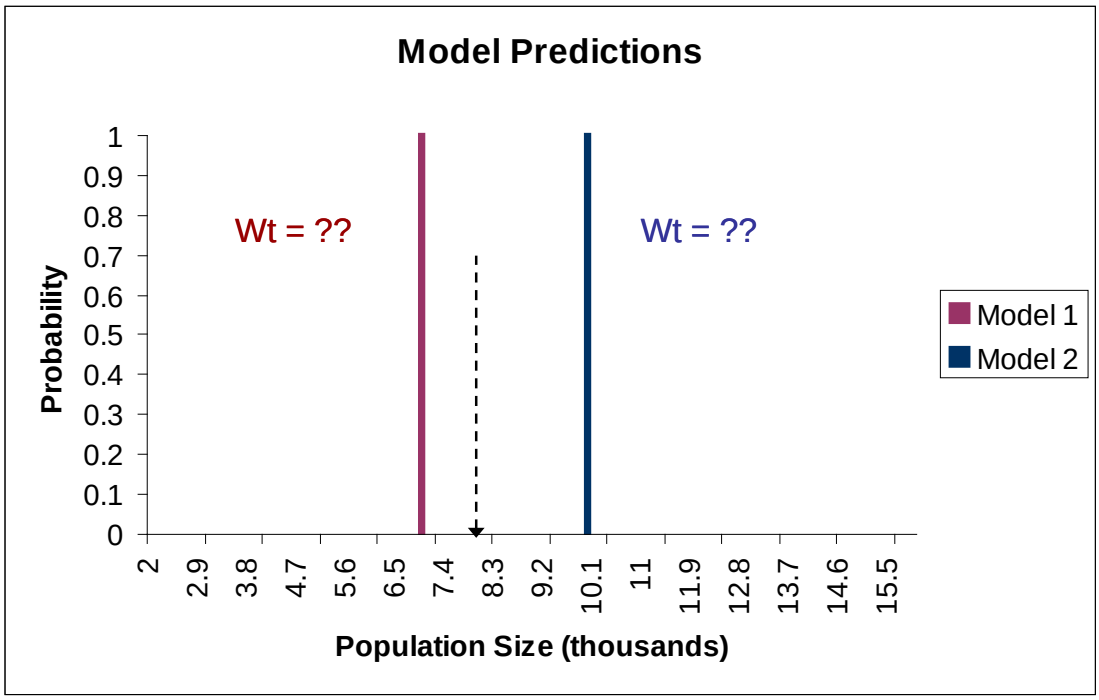
Model Predictions
Probability
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Wt = ??
Wt = ??
2
2.9
3.8
4.7
5.6
6.5
7.4
8.3
9.2
10.1
11
11.9
12.8
13.7
14.6
15.5
Population Size (thousands)
Model 1
Model 2

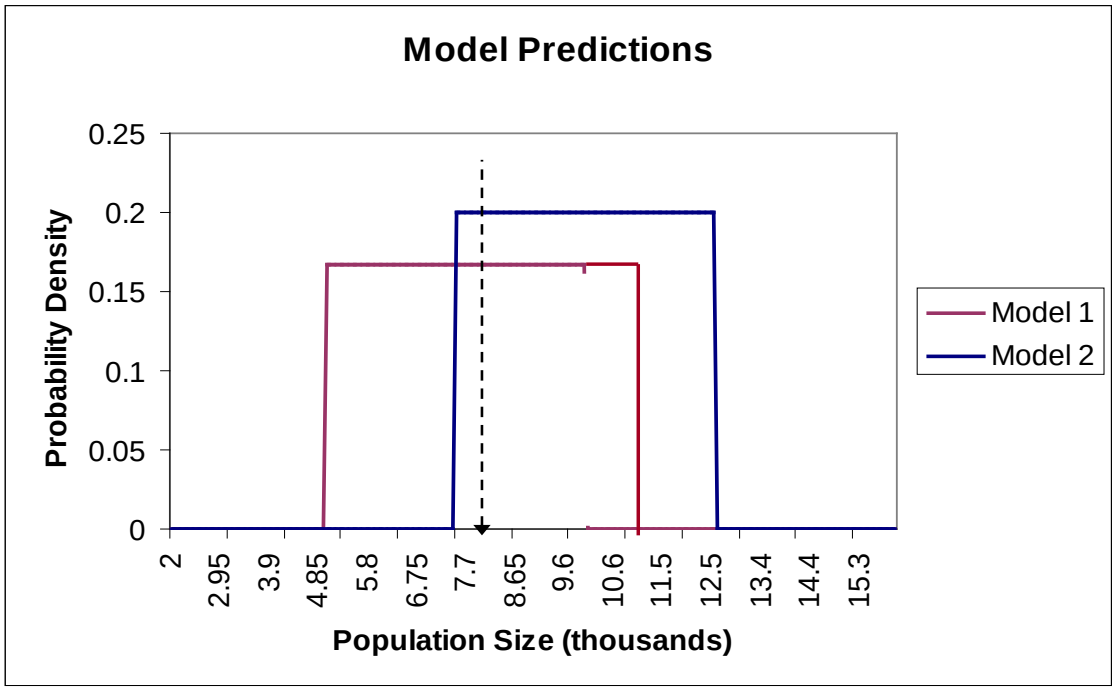
Model Predictions
Probability Density
0.25
0.2
0.15
0.1
0.05
0
2
2.95
3.9
4.85
5.8
6.75
7.7
8.65
9.6
10.6
11.5
12.5
13.4
14.4
15.3
Population Size (thousands)
Model 1
Model 2

Bayes' Formula

p_{t+1}(model i | response$_{t+1}$) =

$$\frac{0.5 * 1/6}{0.5 * 1/6 + 0.5 * 1/5} = 0.45$$

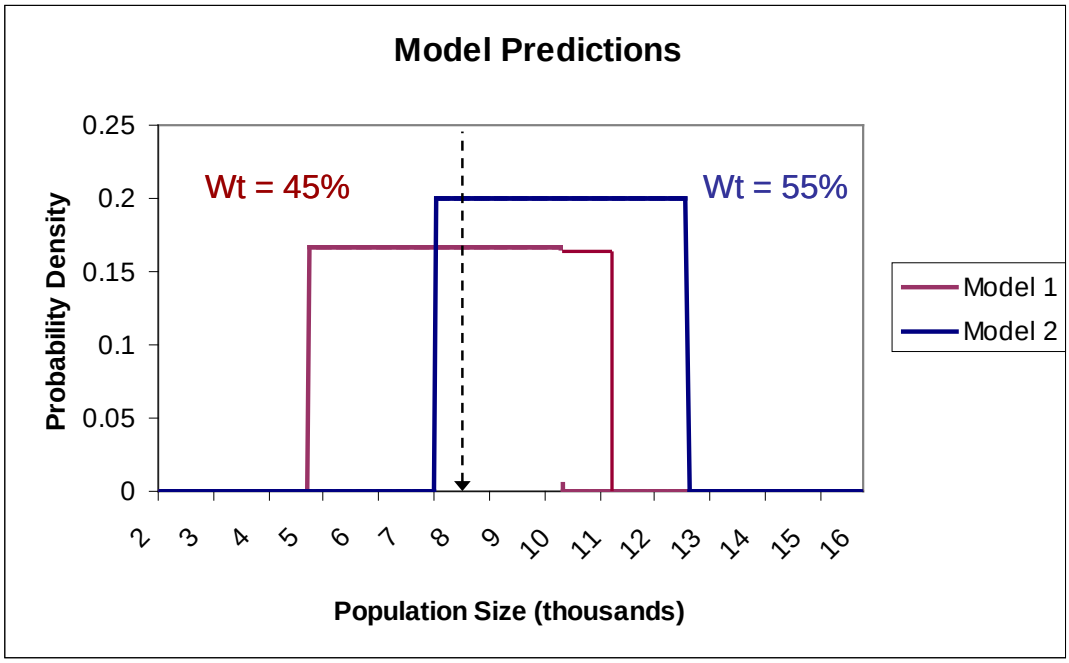

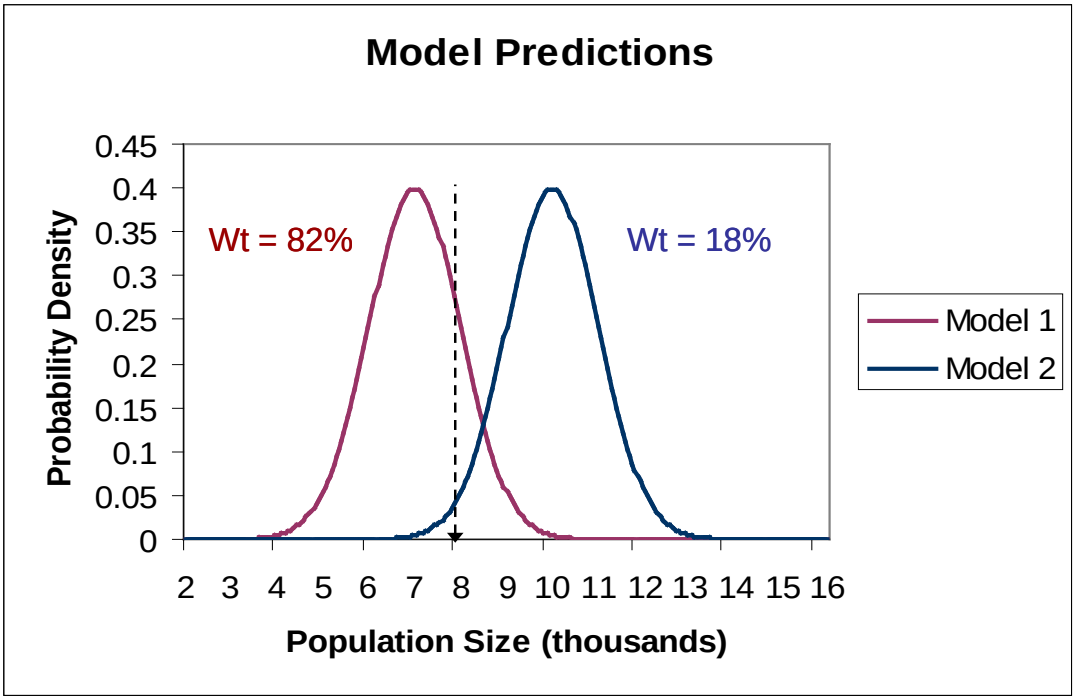

Bayesian Updating – continuous case

- Model "set" defined by key parameter of a single model

$$p_{t+1}(\theta \mid \text{data}_{t+1}) = \frac{p_t(\theta)\, P(\text{data}_{t+1} \mid \theta)}{\int_\theta p_t(\theta)\, P(\text{data}_{t+1} \mid \theta)\, d\theta}$$

- Practical aspects of Bayesian updating
- Conjugacy: "The property that the posterior distribution follows the same parametric form as the prior distribution" (Gelman 2000)
  - E.g., a Normal prior and likelihood yields a Normal posterior
  - A beta prior with binomial likelihood yields a beta posterior
  - A gamma prior with Poisson likelihood yields a gamma posterior
- Often the form of the prior in combination with the likelihood results in a posterior that cannot be solved analytically and other methods are required for evaluation; e.g., MCMC (Markov chain Monte Carlo) methods

Beta distribution

$$f(x|\alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}$$

- For $\alpha>0$, $\beta>0$, f(x) restricted to 0 - 1 interval
- Useful for modeling proportions (e.g., survival or harvest rates)
- Conjugate prior for the Binomial distribution where x is the probability of success

$$E(x) = \frac{\alpha}{\alpha+\beta}$$

$$var(x) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$$

Beta distribution

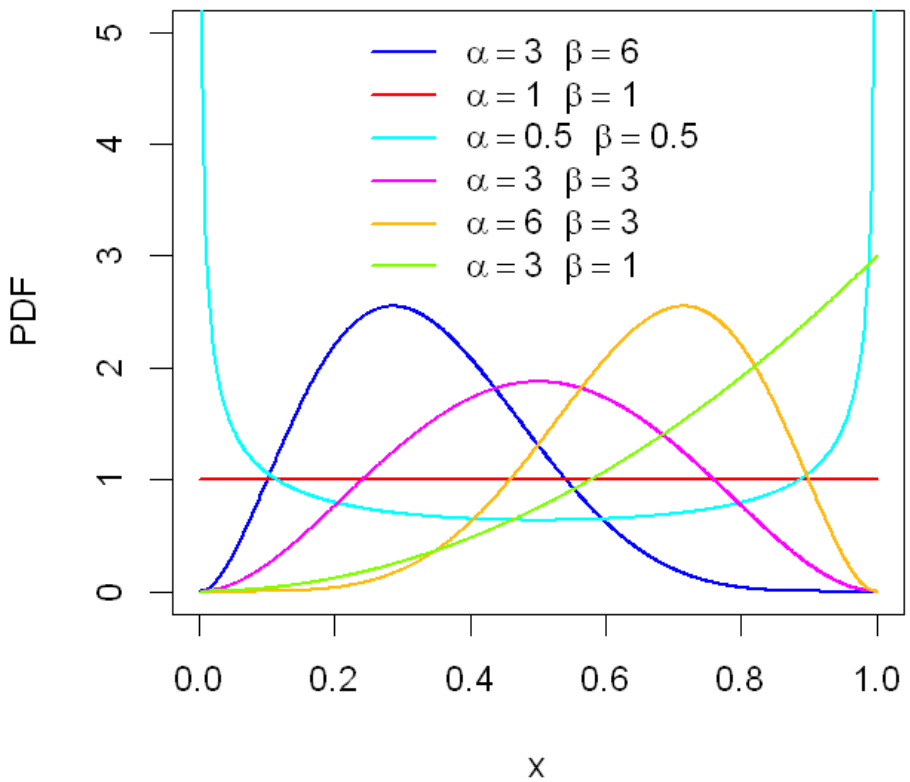

Beta-Binomial example:

Estimating survival rates with prior information where S = 0.7, SE_S = 0.1

- Use method of moments to specify prior distribution

$$\alpha = \mu_S\left(\frac{\mu_S(1-\mu_S)}{\sigma_S^2} - 1 \right)$$

$$\beta = (1-\mu_S)\left(\frac{\mu_S(1-\mu_S)}{\sigma_S^2} - 1 \right)$$

$$p(S) \sim Beta(\alpha = 14, \beta = 6)$$

Prior Distribution

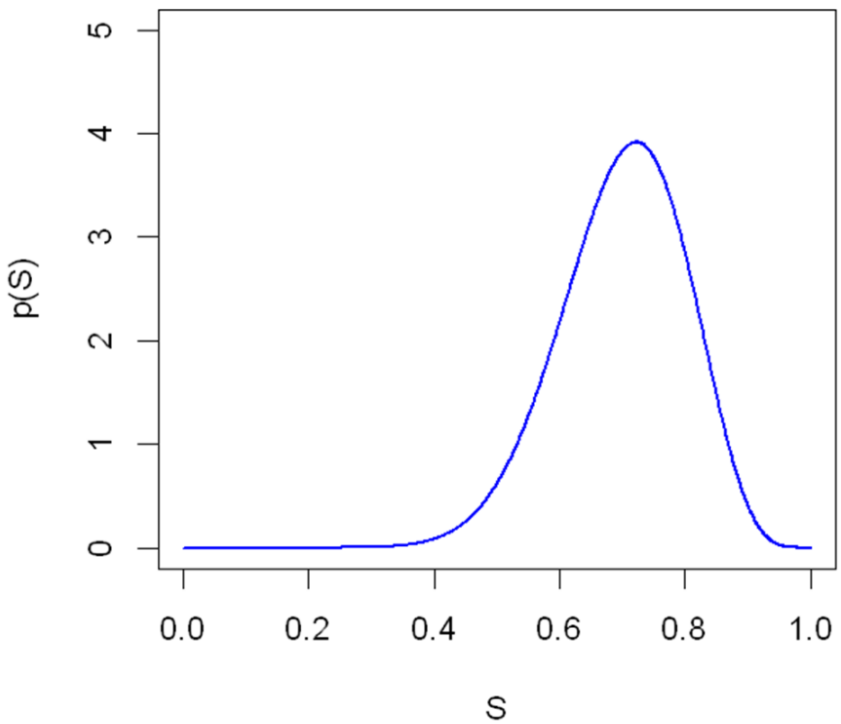

Do Experiment

- Apply n = 20 transmitters in year 1,
- Observe y = 18 survive to year 2,
- Assume y ~ Binomial (n=20, S)
- Specify Posterior

$$p(S|y,n) \propto \underbrace{\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} S^{\alpha-1}(1-S)^{\beta-1}}_{\text{Beta prior}} \underbrace{S^{y}(1-S)^{(n-y)}}_{\text{Binomial likelihood}}$$

Combine terms

$$p(S|y,n) \propto \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} S^{y+\alpha-1}(1-S)^{n-y+\beta-1}$$

Let $\alpha_{new} = y + \alpha$ and $\beta_{new} = n - y + \beta$

$$p(S|y,n) \propto Beta(\alpha_{new} = 18 + 14, \beta_{new} = 20 - 18 + 6)$$

Evaluate Posterior

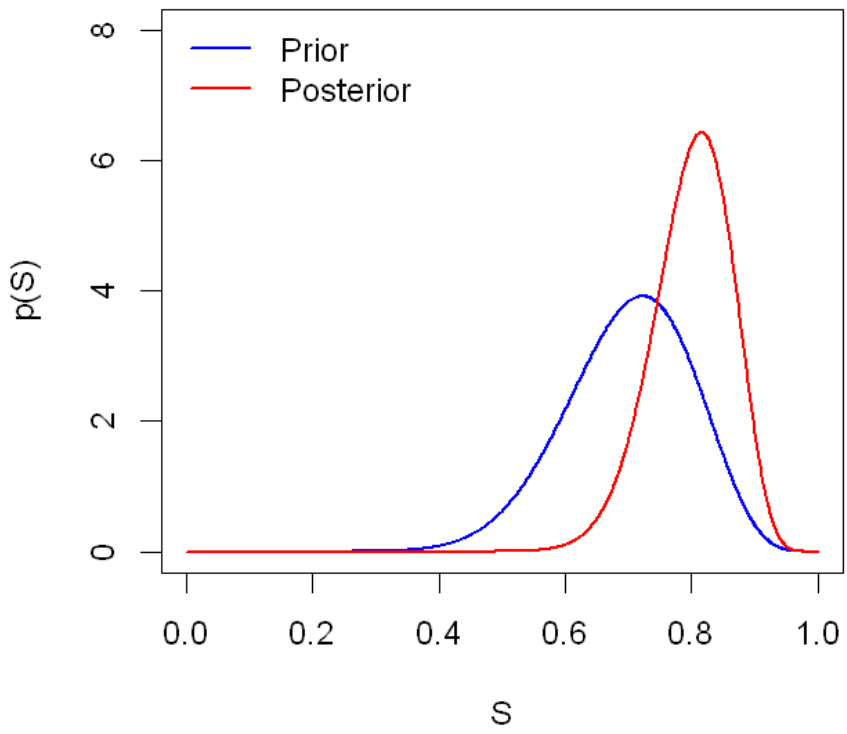

$$E(S|y,n) = 0.8$$

$$var(S|y,n) = .004$$

Actual Example:
Adaptive Harvest Management for Mid-continent Mallards

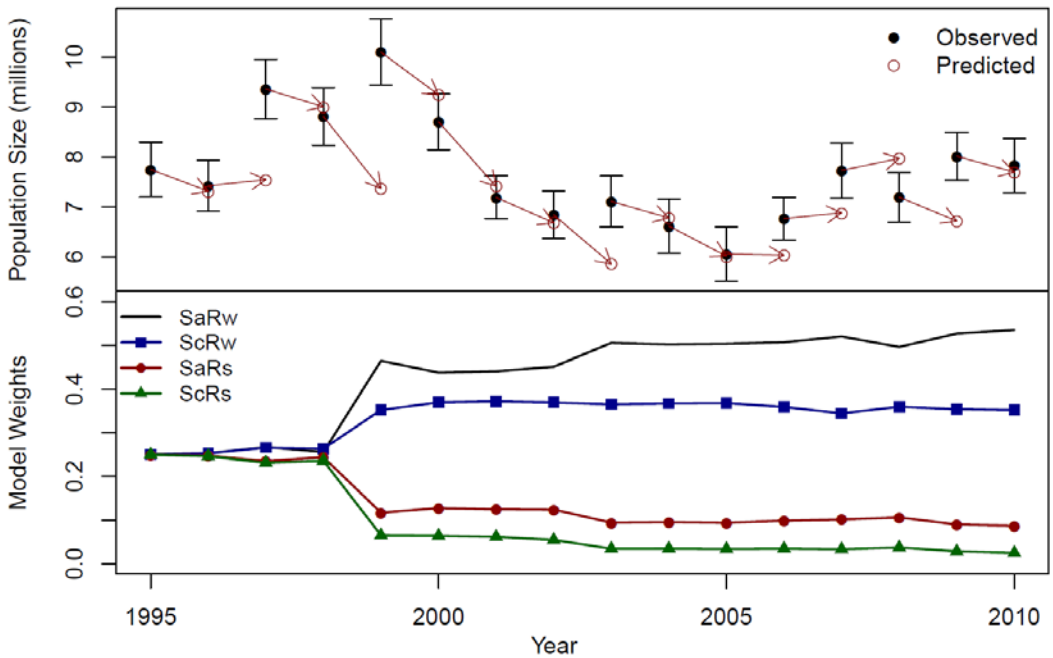

How fast does learning occur?

What affects the speed of learning?

- Model structure, parameter values
  - Does the set include a good approximating model?
    - i.e., a robust predictor over the state space
  - Are parameter estimates
    - Precise?
    - Unbiased?
- Amount of noise (stochasticity) in the system
- Partial observability
  - Bias and precision in monitoring
- Approach to optimization
- Spatial replication

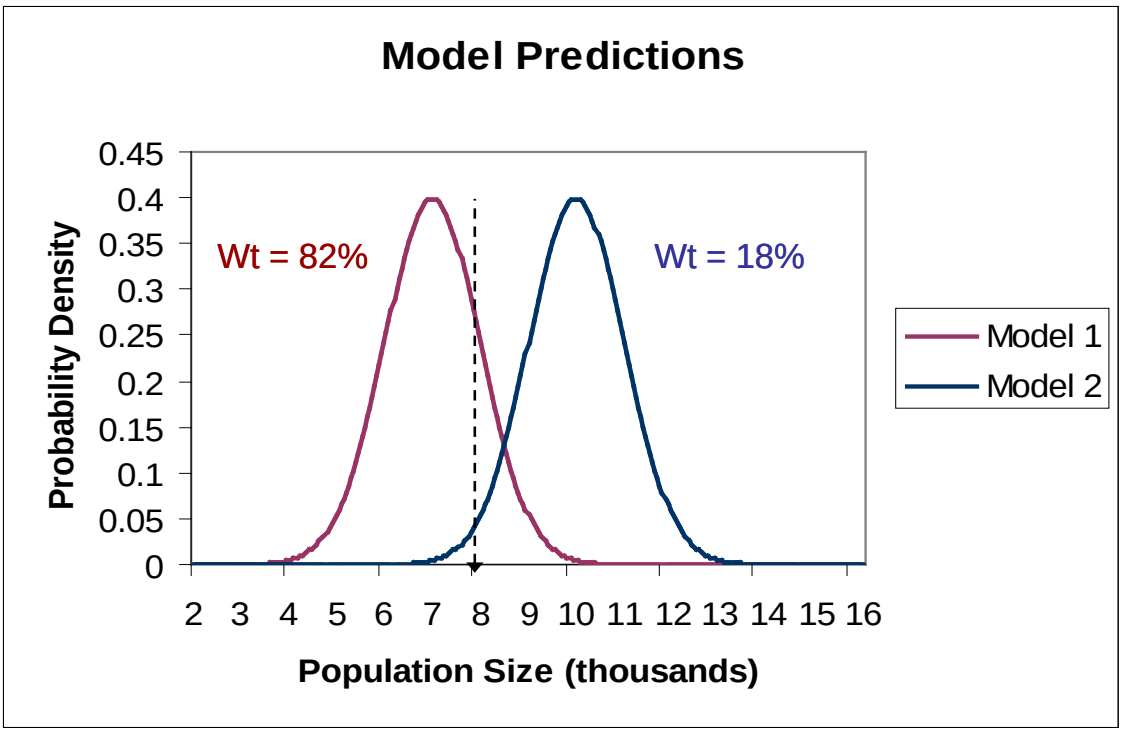
Model Predictions
Probability Density
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
Wt = 82%
Wt = 18%
Model 1
Model 2
2 3 4 5 6 7 8 9 10 11 12 13 14 15 16
Population Size (thousands)

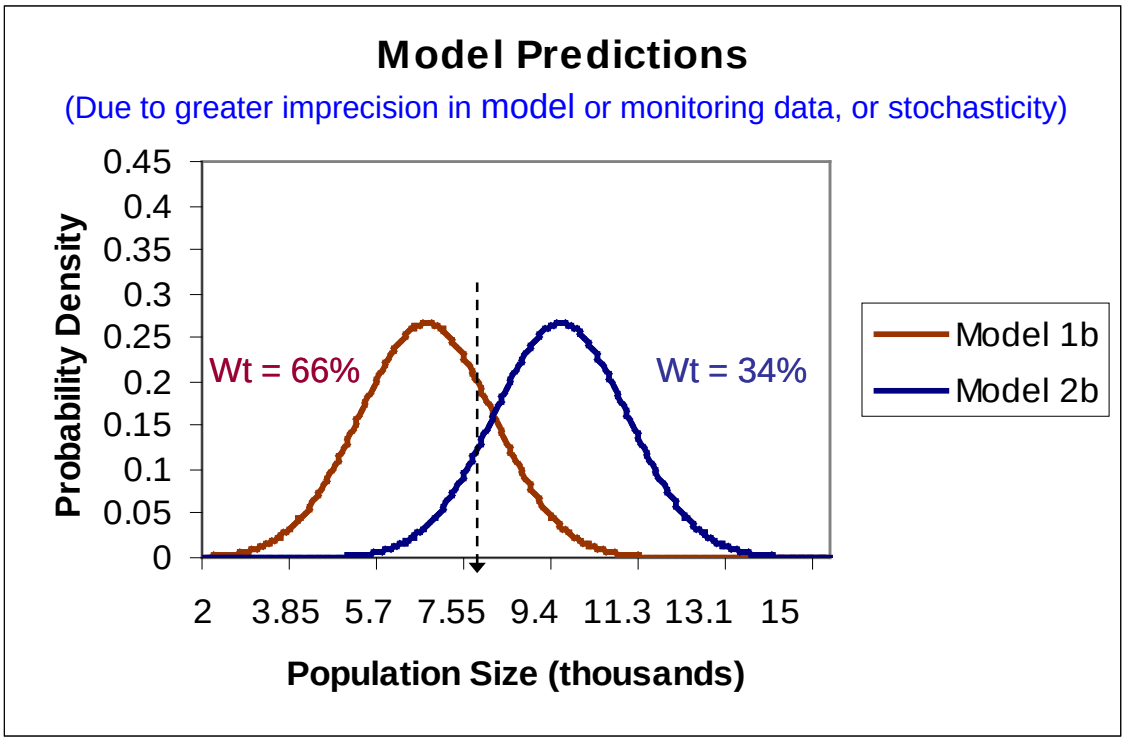
Model Predictions
(Due to greater imprecision in model or monitoring data, or stochasticity)
Probability Density
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
Wt = 66%
Wt = 34%
Model 1b
Model 2b
2 3.85 5.7 7.55 9.4 11.3 13.1 15
Population Size (thousands)

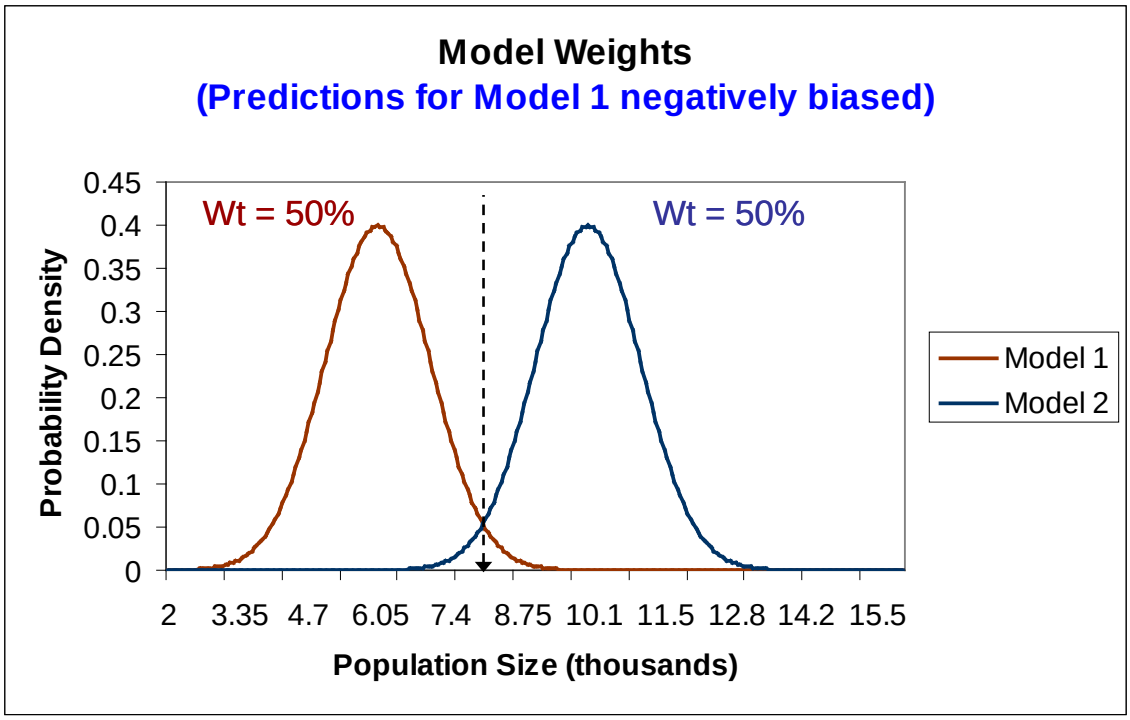
Model Weights
(Predictions for Model 1 negatively biased)
Wt = 50%
Wt = 50%
Probability Density
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
2
3.35
4.7
6.05
7.4
8.75
10.1
11.5
12.8
14.2
15.5
Population Size (thousands)
Model 1
Model 2

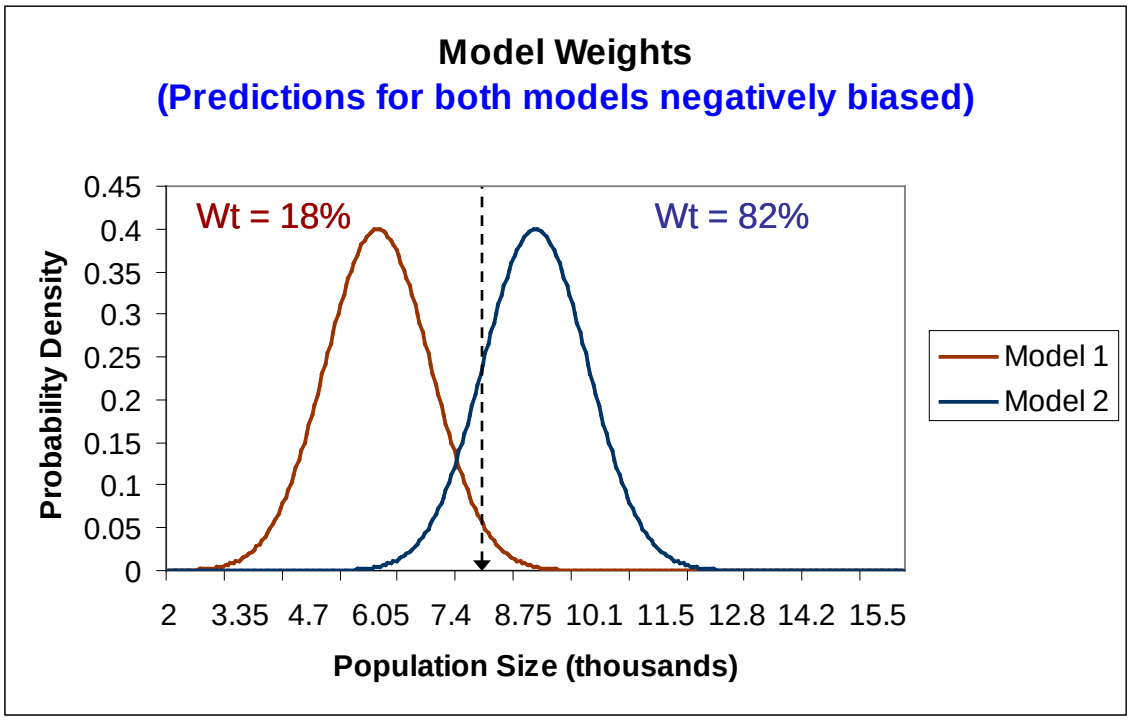
Model Weights
(Predictions for both models negatively biased)
Wt = 18%
Wt = 82%
Probability Density
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
2
3.35
4.7
6.05
7.4
8.75
10.1
11.5
12.8
14.2
15.5
Population Size (thousands)
Model 1
Model 2

Can we measure the cost of poor monitoring?

Driving in Fog: Accounting for the hidden costs of measurement uncertainty in wildlife decision-making through adaptive monitoring design

Clinton T. Moore and William L. Kendall
USGS Patuxent Wildlife Research Center

State-specific decision making

System states

- Current physical and biological conditions of a managed system
  - e.g., number of bird pairs and number of nest cavities

Decisions

- Candidate management actions
  - e.g., treatments to increase nest cavity abundance

Rewards

- Expected management gain for given decision and system state
  - e.g., brood production

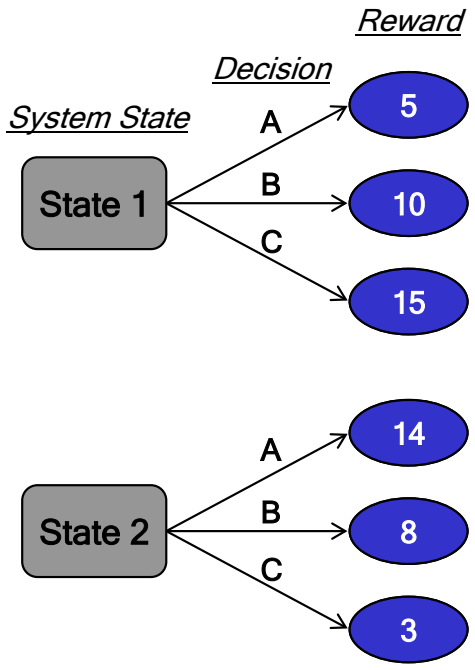

The decision you make depends on how you see the system

Case 1: True state is observable

Decision "A" is best

Decision gain of 14 units

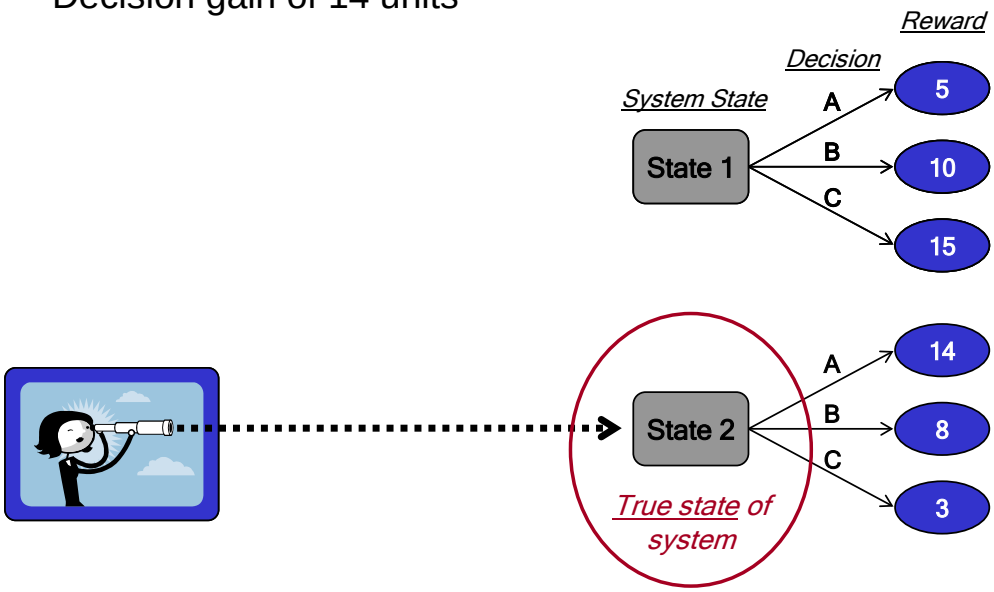

The decision you make depends on how you see the system

Case 2: True state *not* observable

Decision "C" is apparently best

Decision *cost* of 14 - 3 = 11 units

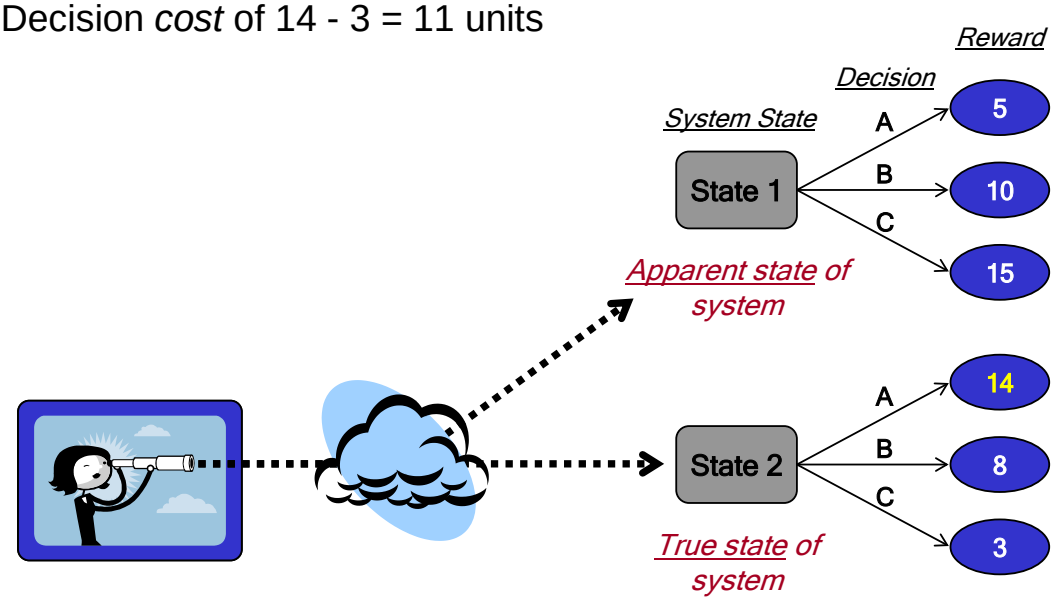

Partial observability

- Leads to reduction in management returns
  - Best decision for apparent state differs from that for true, unknown state
  - Management *opportunity cost* of partial observability
    - Measurable in units of the resource

Partial observability and management return

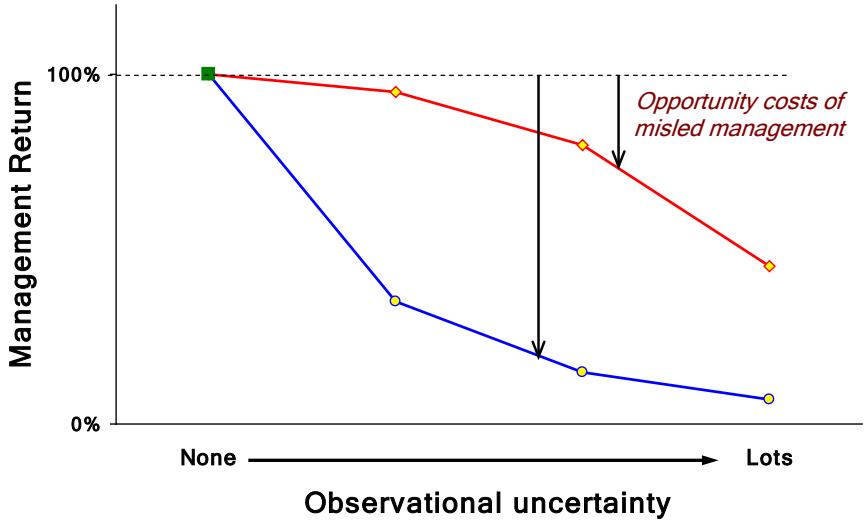

Partial observability

- Leads to reduction in management returns
  - Under structural (model) uncertainty, partial observability can interfere with ability to resolve model uncertainty and improve management

Partial observability and model identification

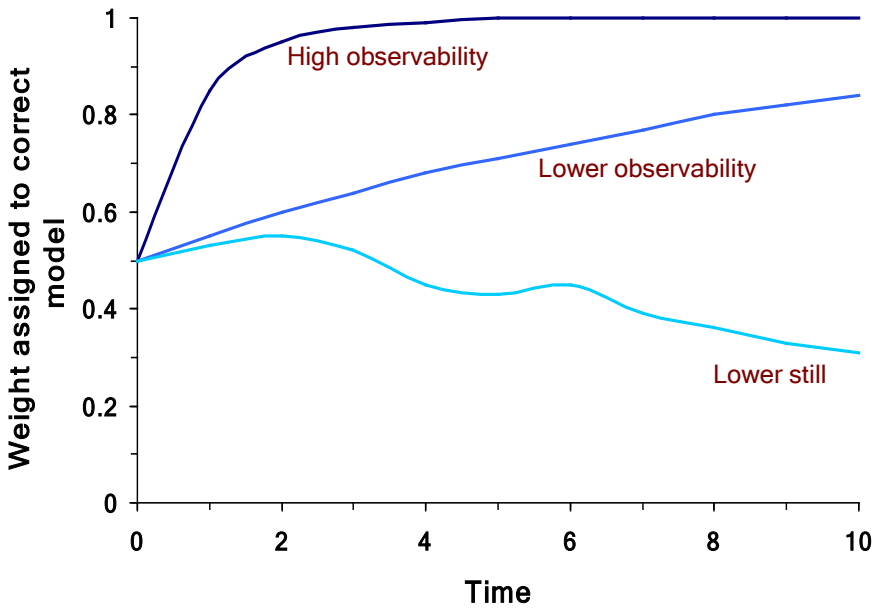

Monitoring Program Costs: Considerations

- Cost of monitoring

- Costs of not monitoring or monitoring poorly:
  - Poor estimates of state, for decisions
  - Slow/improper resolution of structural uncertainty
  - Poor estimation of model parameters

Monitoring effort should be formally cast as a management decision variable?

- See Hauser et al. (2006, 2009), Moore & McCarthy (2010)
- Recurring decisions about:
  1. Management action
  2. Monitoring intensity

- Objective:
  - Include both resource conservation returns and survey costs via use of common currency, utility thresholds, whatever

- Adaptive monitoring design
  - Value of reducing uncertainty is high
    → monitoring intensity increases
  - Value of reducing uncertainty is low
    → monitoring intensity decreases

Approach to Optimization

<u>Approaches to optimization</u>

- Passive ARM
  - Decision made based on management objectives and current information state (i.e., model weights)

- Active ARM
  - Simultaneous/concurrent Active ARM
    - Decision made based on management objectives, current information state and anticipated benefit of learning (Dual Control).
  - Sequential Active ARM
    - (1) Experimentation (learn quickly for a set of steps), with little consideration for resource returns, followed by
    - (2) Passive ARM under "best" models based on (1)
    - e.g., McCarthy & Parris (2008)

<u>Speed of learning also artifact of objectives, approach to optimization</u>

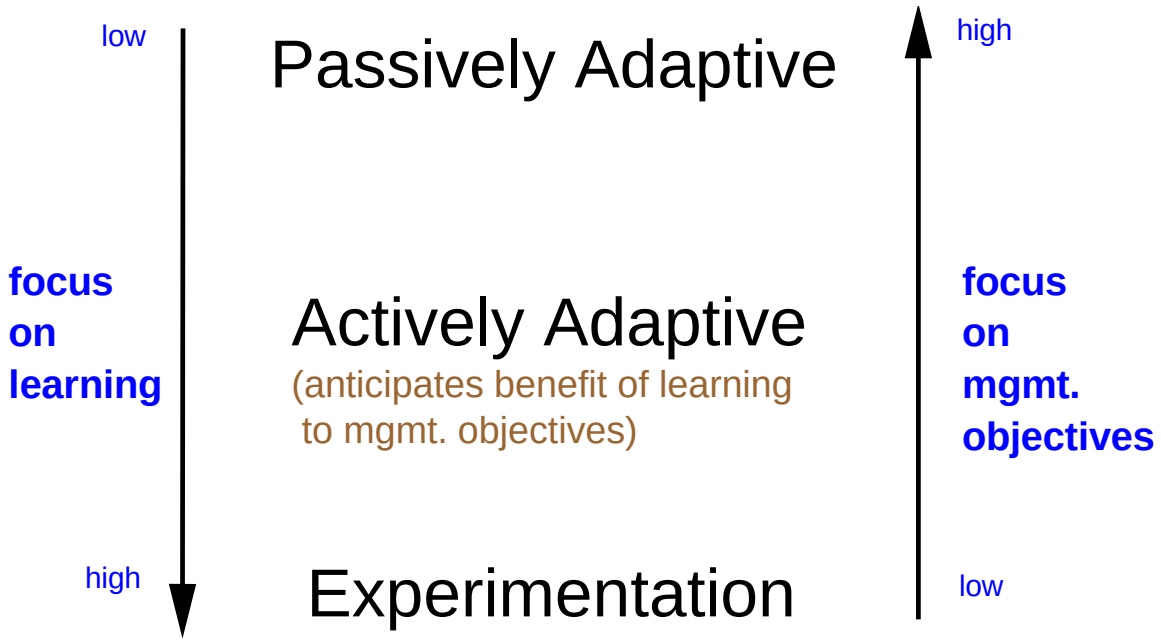

Can learn faster with spatial replication

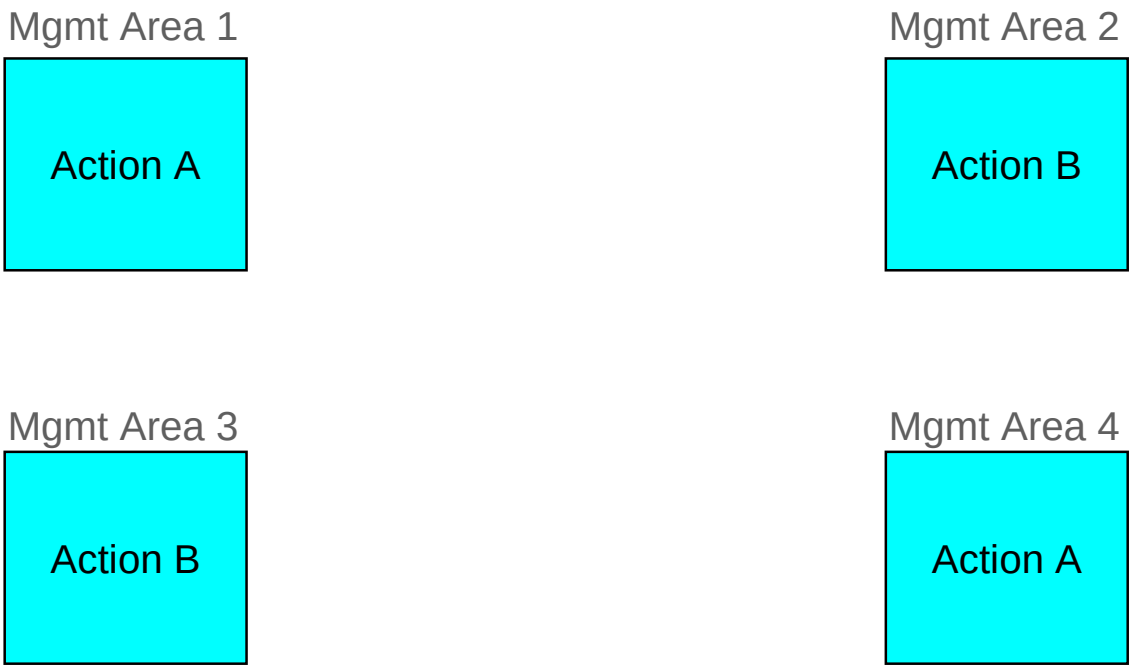

Robustness of a model vs. robustness of the model set

- Suggestion: don't discard a hypothesis too quickly based on poor model predictive performance (model may not properly capture hypothesis, may be constructed with poor parameter estimates, etc.)
- If weights are ambiguous (e.g., bouncing around over time) but model set predicts well, then no need to panic
- If model set predicts poorly, then really need to revise or add models (double-loop learning)

Conclusions

- Learning is hallmark of ARM

- It is not appropriate to label a management program as "adaptive" without a clear mechanism for incorporation of learning to improve subsequent management

- The purpose of learning in ARM is to provide increased returns by improving predictions across entire state space

- Bayes' formula is natural vehicle for "learning" in ARM (and in science)

- Rate of "learning" depends on many factors, e.g.,
    - Stochastic variation of model predictions
    - Variation among model-based predictions for members of model set
    - Partial observability
    - Approach to optimization

  - True learning depends on how well at least one member of the model set captures underlying mechanisms (so we still need to think)