Learning

ALC3176: Adaptive Management:
Structured Decision Making for Recurrent Decisions
Chapter 7

Presented by: Scott Boomer, USFWS DMBM

Module Developed by:
Bill Kendall and Jim Nichols
USGS Patuxent Wildlife Research Center

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?

- 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.

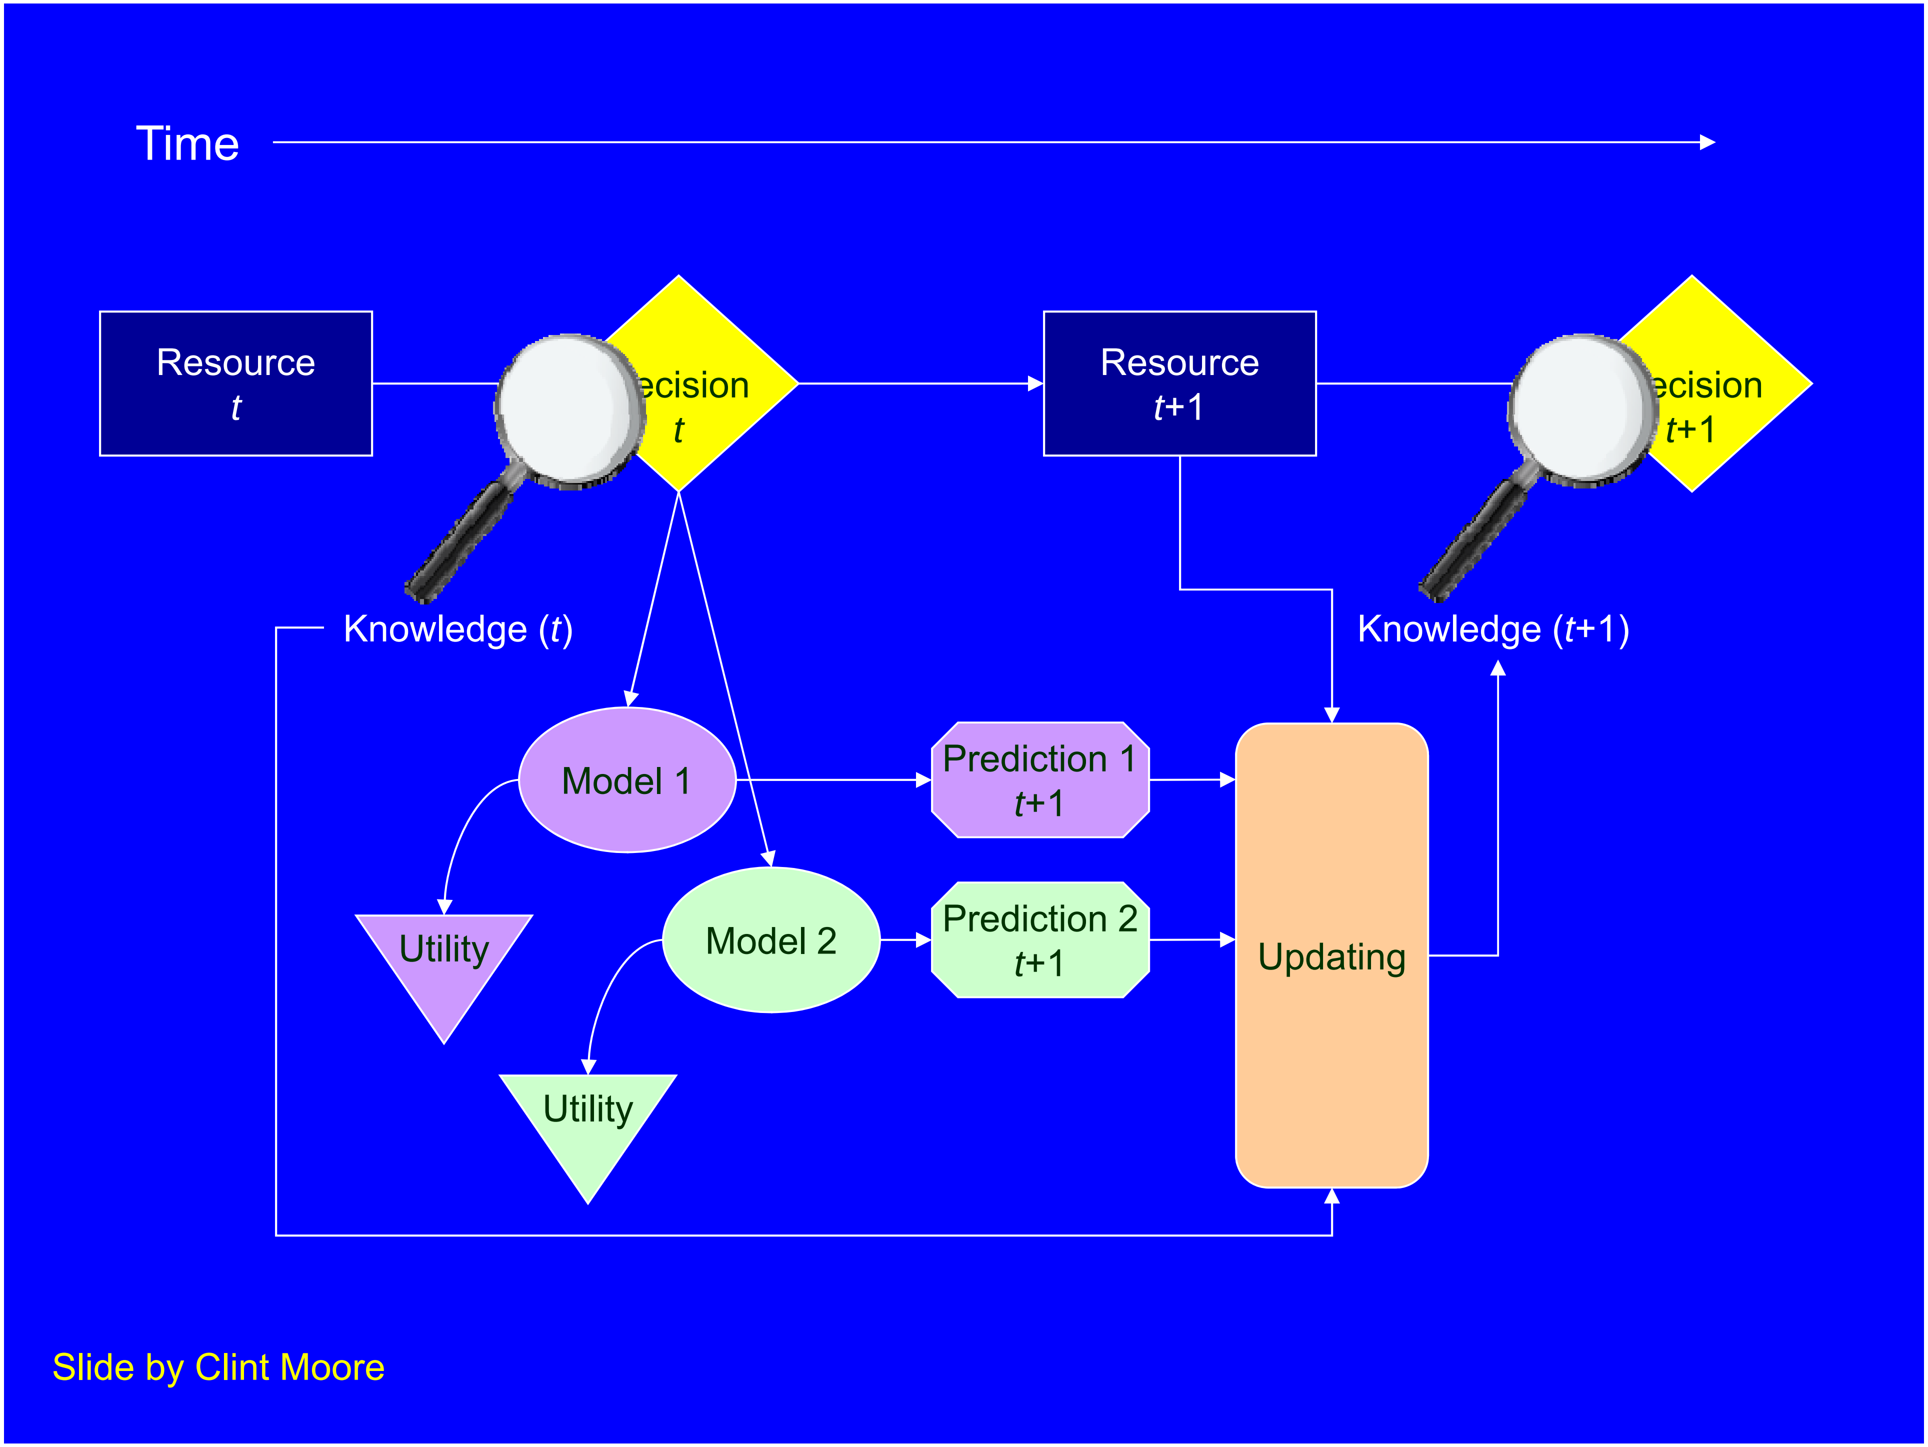

Time
Resource t
ecision t
Resource t+1
ecision t+1
Knowledge (t)
Knowledge (t+1)
Model 1
Model 2
Prediction 1 t+1
Prediction 2 t+1
Utility
Utility
Updating
Slide by Clint Moore

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

New weight of *model i* $\propto$

*(Old weight of model i) **

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

Bayes' 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 furthers 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

- 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 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* state (e.g., population size) with resulting *observed* state.

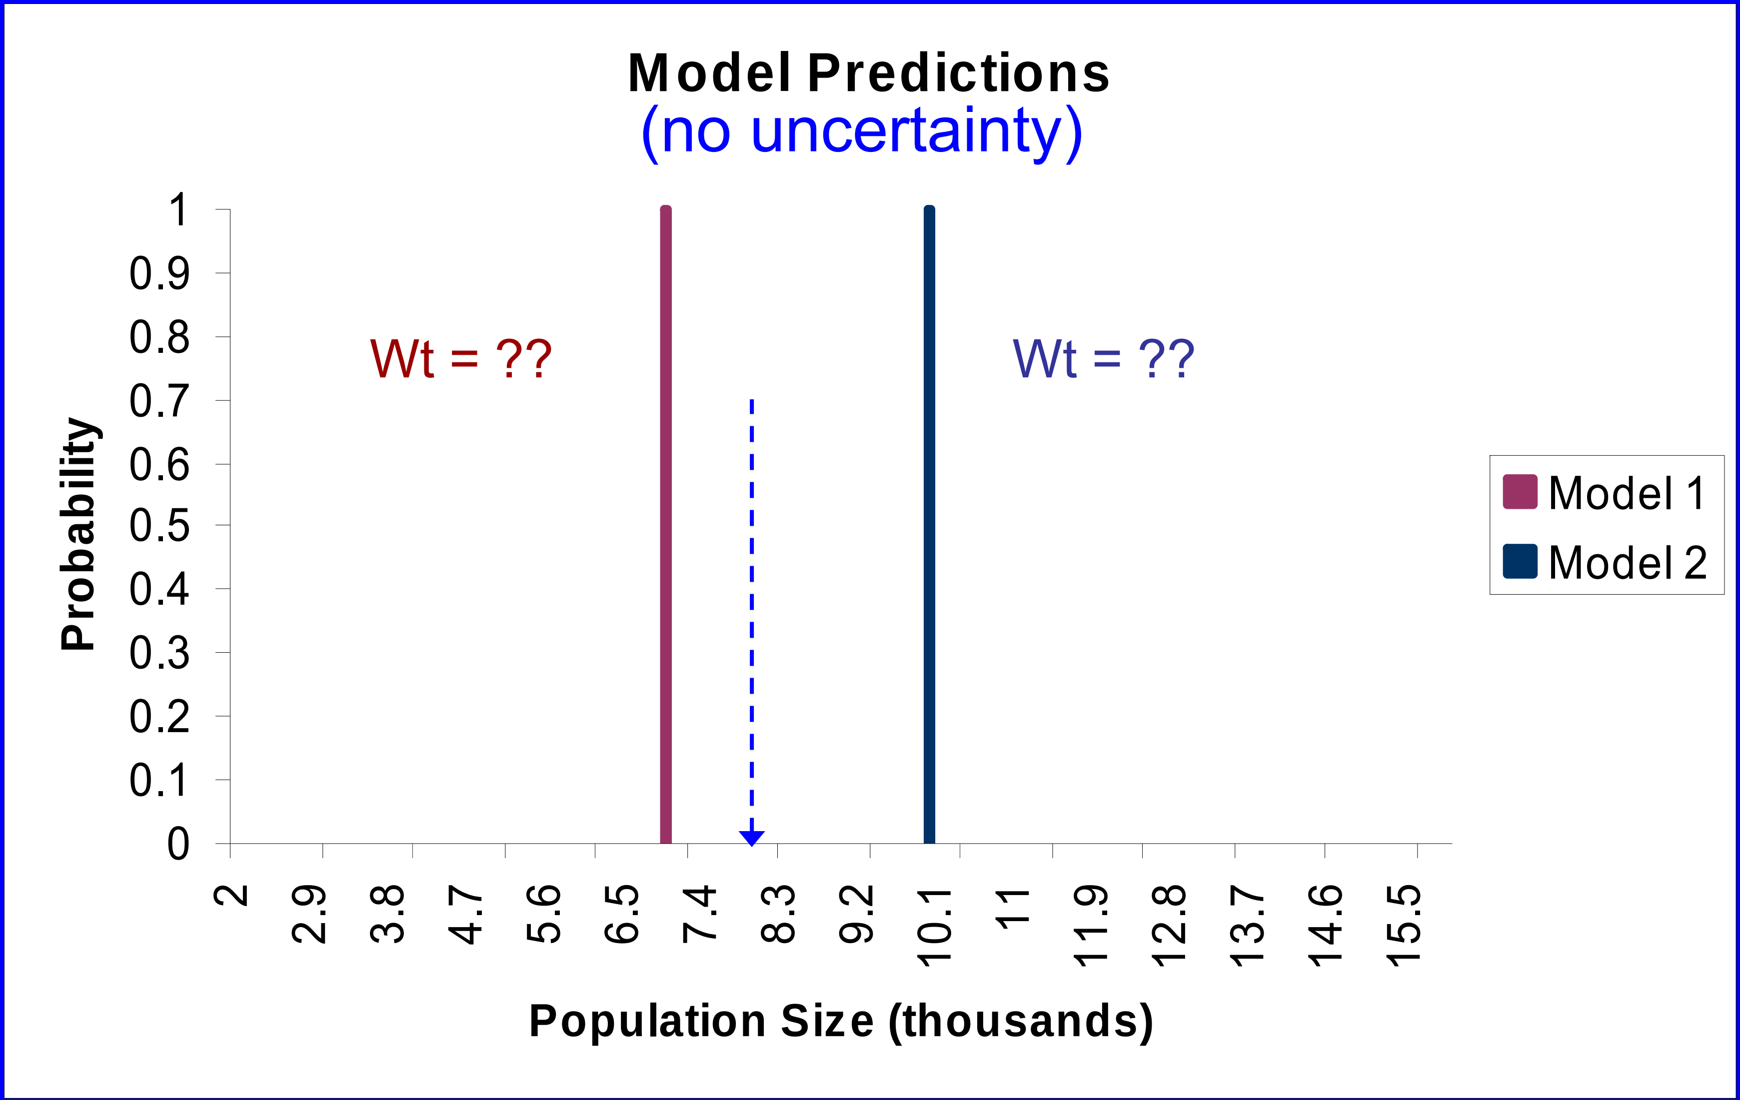
Model Predictions
(no uncertainty)
Probability
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
Wt = ??
Wt = ??
Model 1
Model 2
2
2.9
3.8
4.7
5.6
6.5
7.4
8.3
9.2
10.1
11
11.9
12.8
13.7
14.6
15.5
Population Size (thousands)

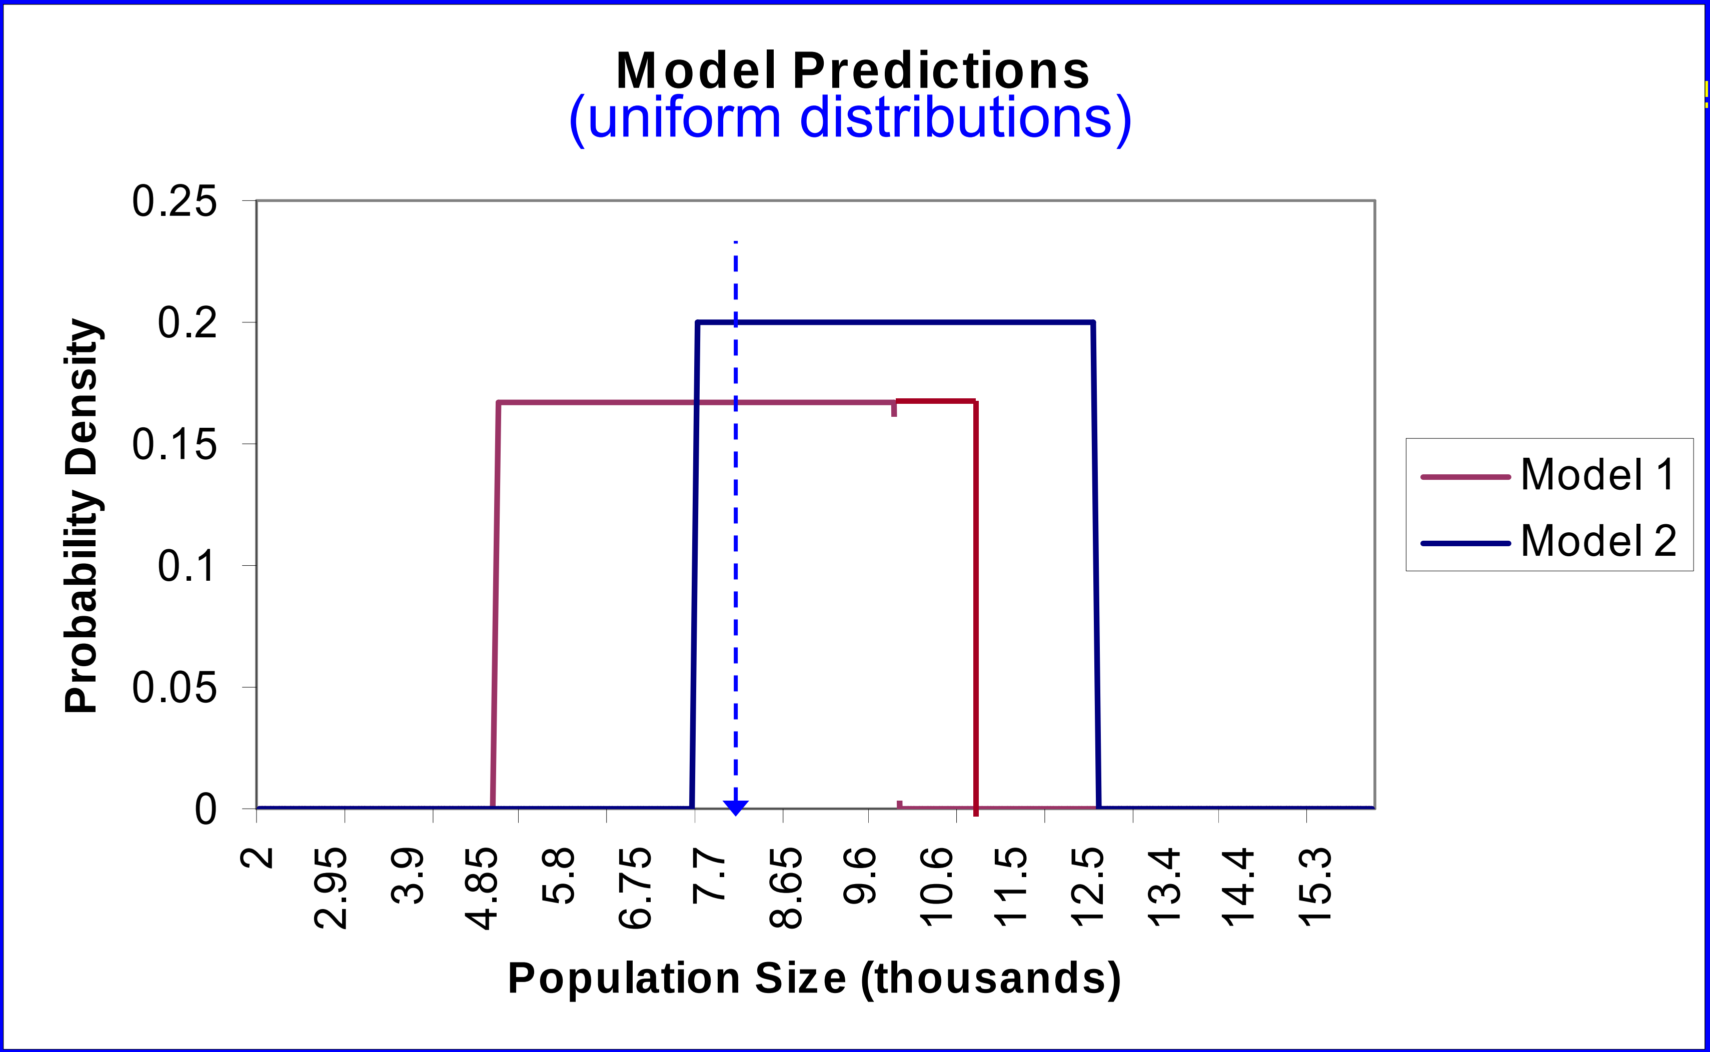

Model Predictions
(uniform distributions)
0.25
0.2
0.15
0.1
0.05
0
Probability Density
2
2.95
3.9
4.85
5.8
6.75
7.7
8.65
9.6
10.6
11.5
12.5
13.4
14.4
15.3
Population Size (thousands)
Model 1
Model 2

Bayes' Formula

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

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

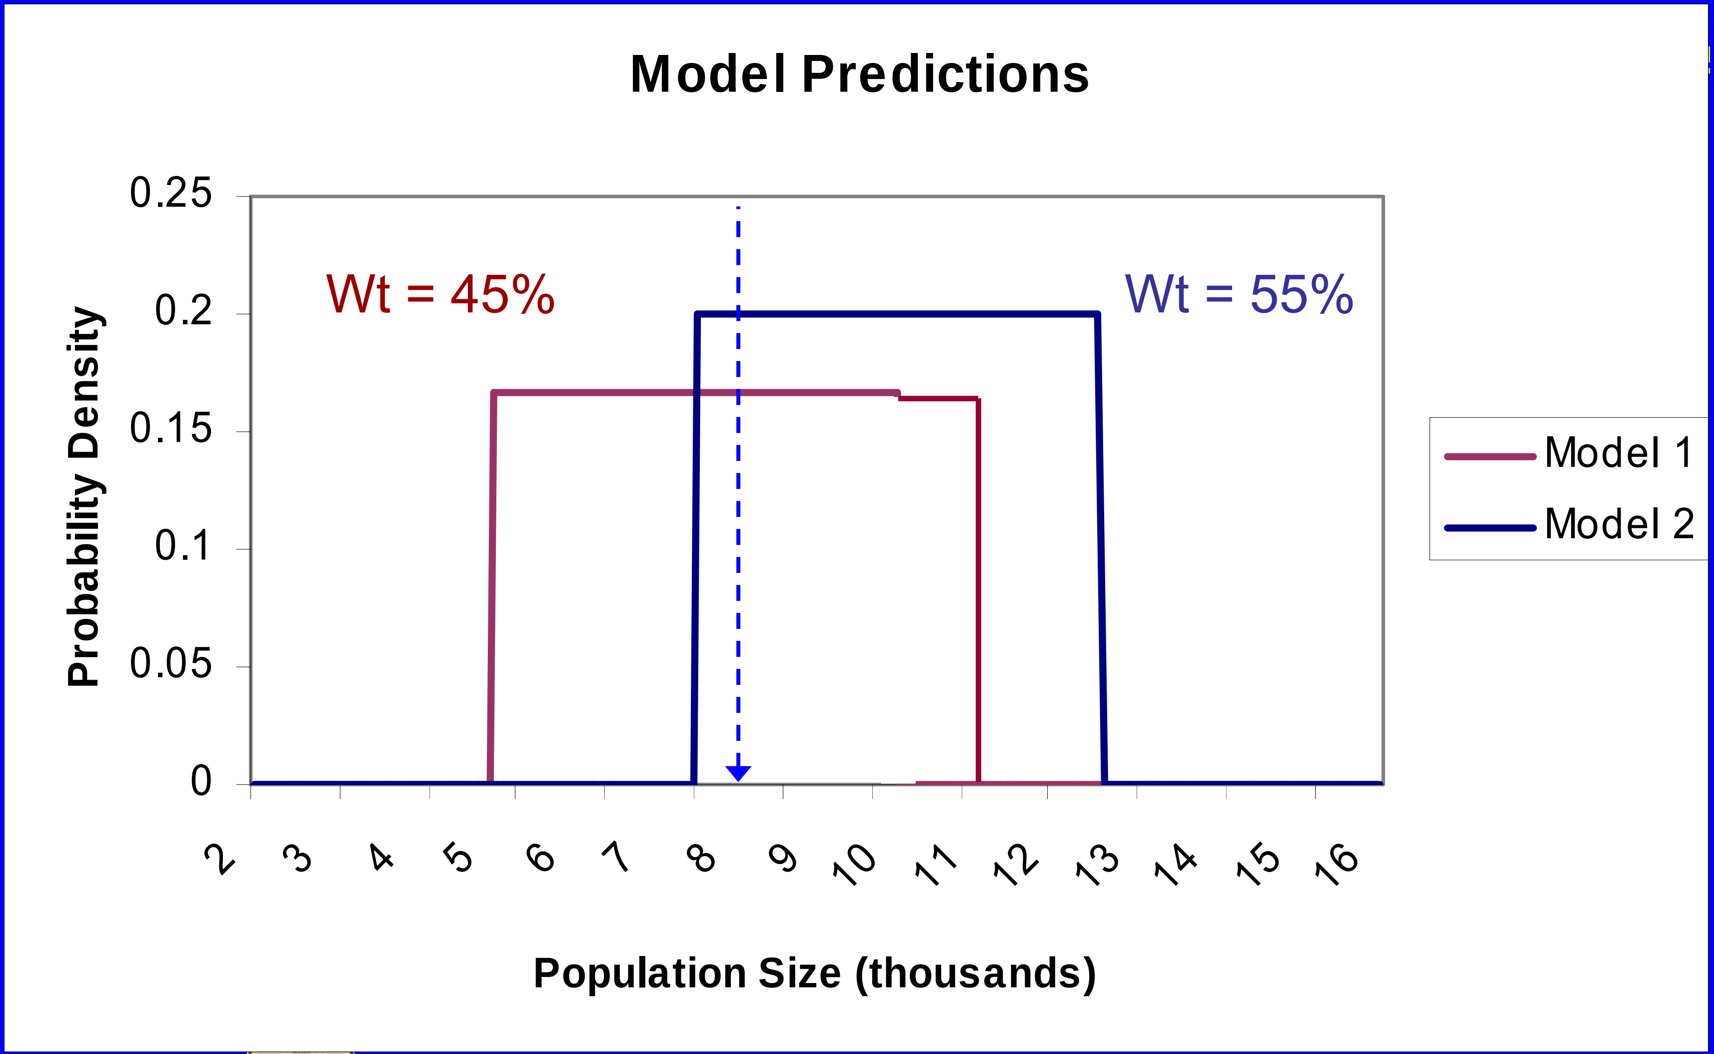

Model Predictions
Wt = 45%
Wt = 55%
Probability Density
0.25
0.2
0.15
0.1
0.05
0
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
Population Size (thousands)
Model 1
Model 2

Model Predictions

(normal distributions - some uncertainty)

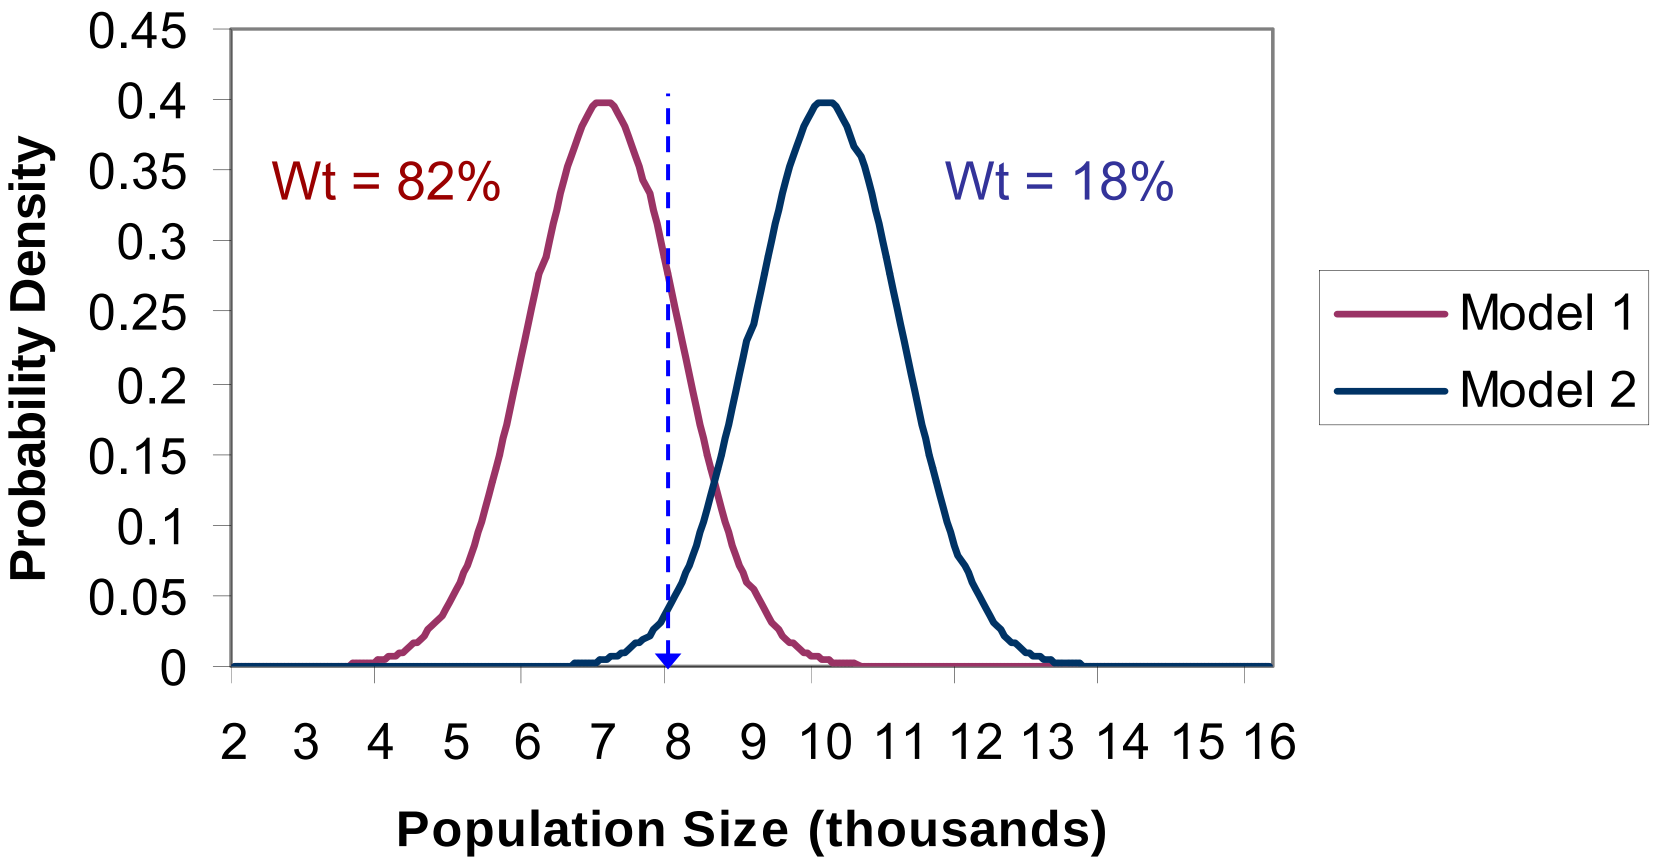

Bayesian Updating- model "set" defined by key parameter of a single model

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

Practical aspects of Bayesian updating

- Conjugacy: “The property that the posterior distribution follows the same parametric form as the prior distribution” (Gelman et al. 2000)
  - E.g., a Normal prior and likelihood yields a Normal posterior

Practical aspects of Bayesian updating

- 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

Beta Densitys

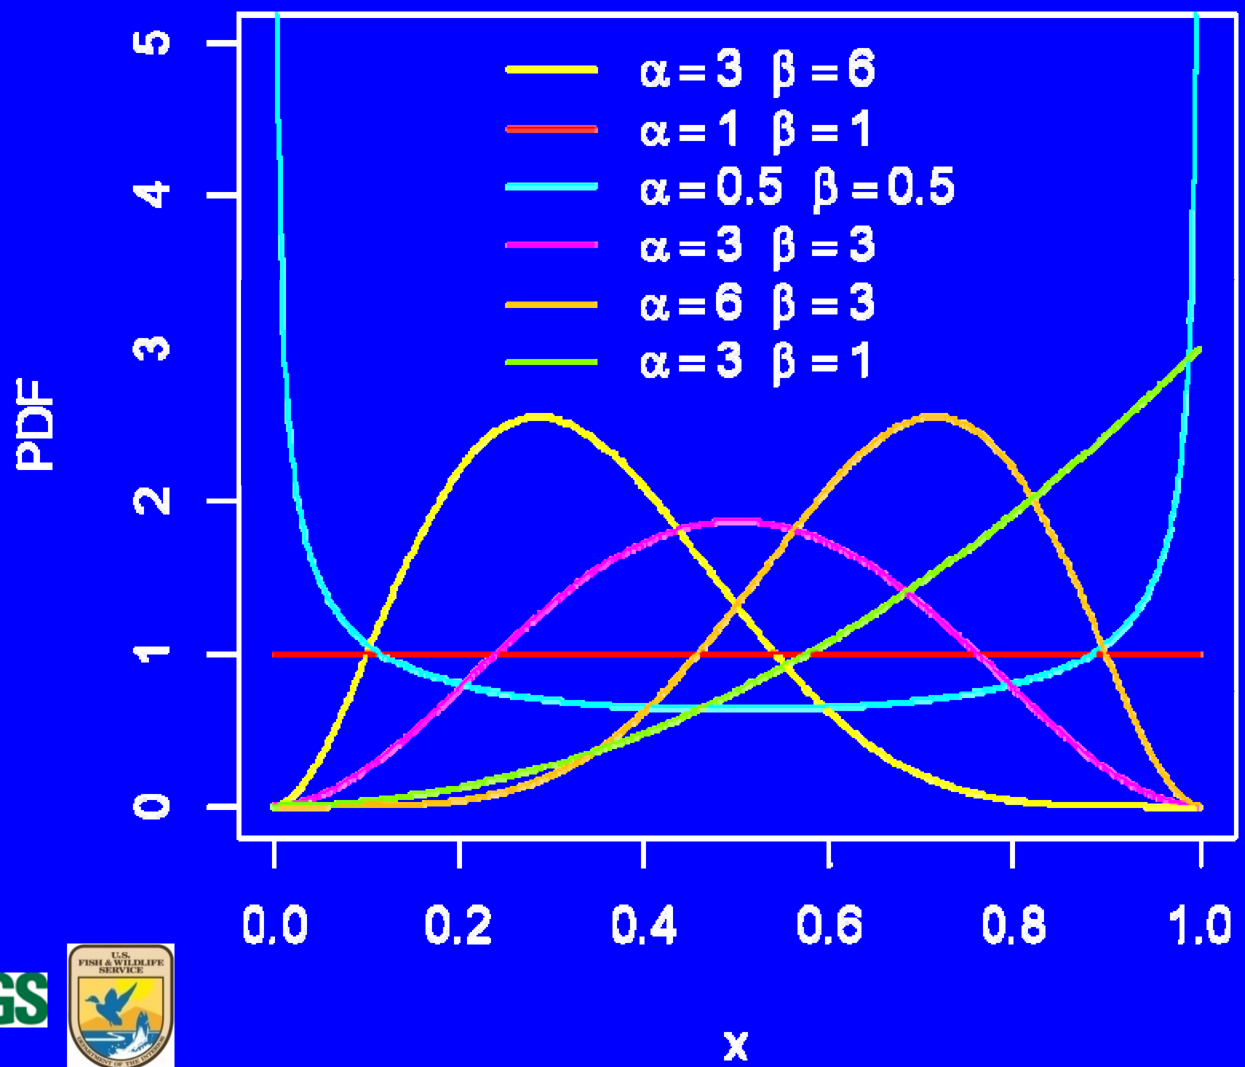

Beta-Binomial example: Estimating survival rates with prior information where S = 0.7, S(SE) = 0. 1

- Use method of moments to specify prior

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

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

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

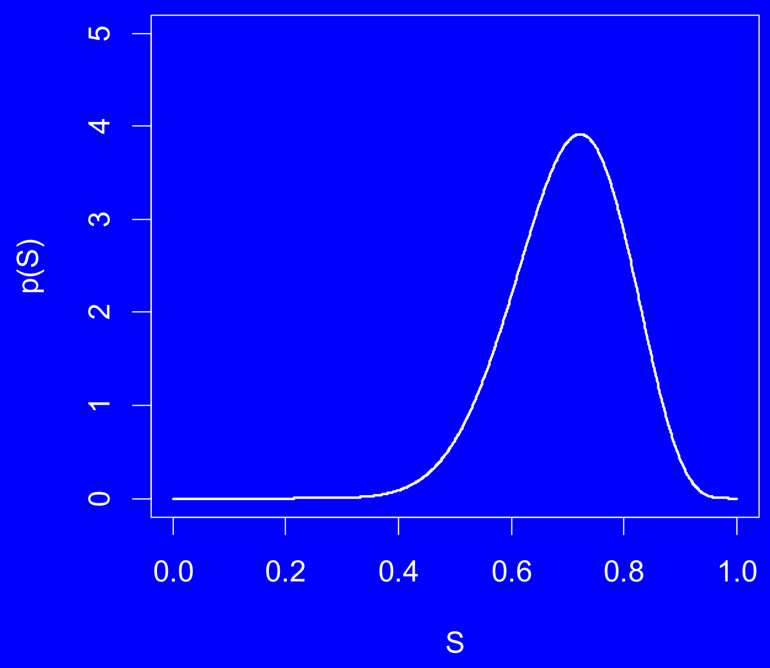

Apply n = 20 transmitters in year 1, y = 18 survive to year 2

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

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

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

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

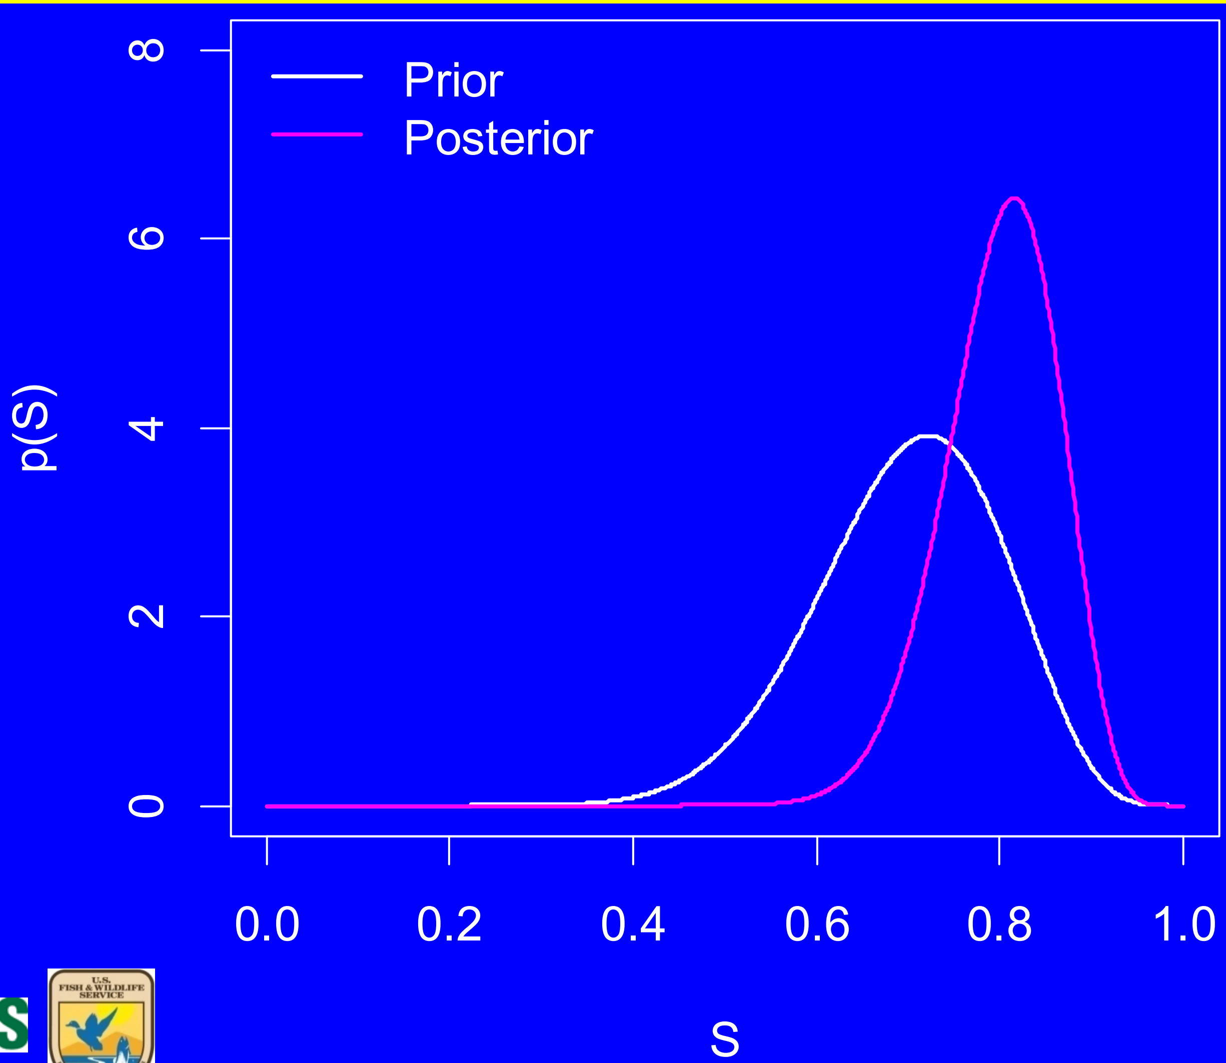

Actual Example

Adaptive Harvest Management for Mid-continent Mallard Ducks

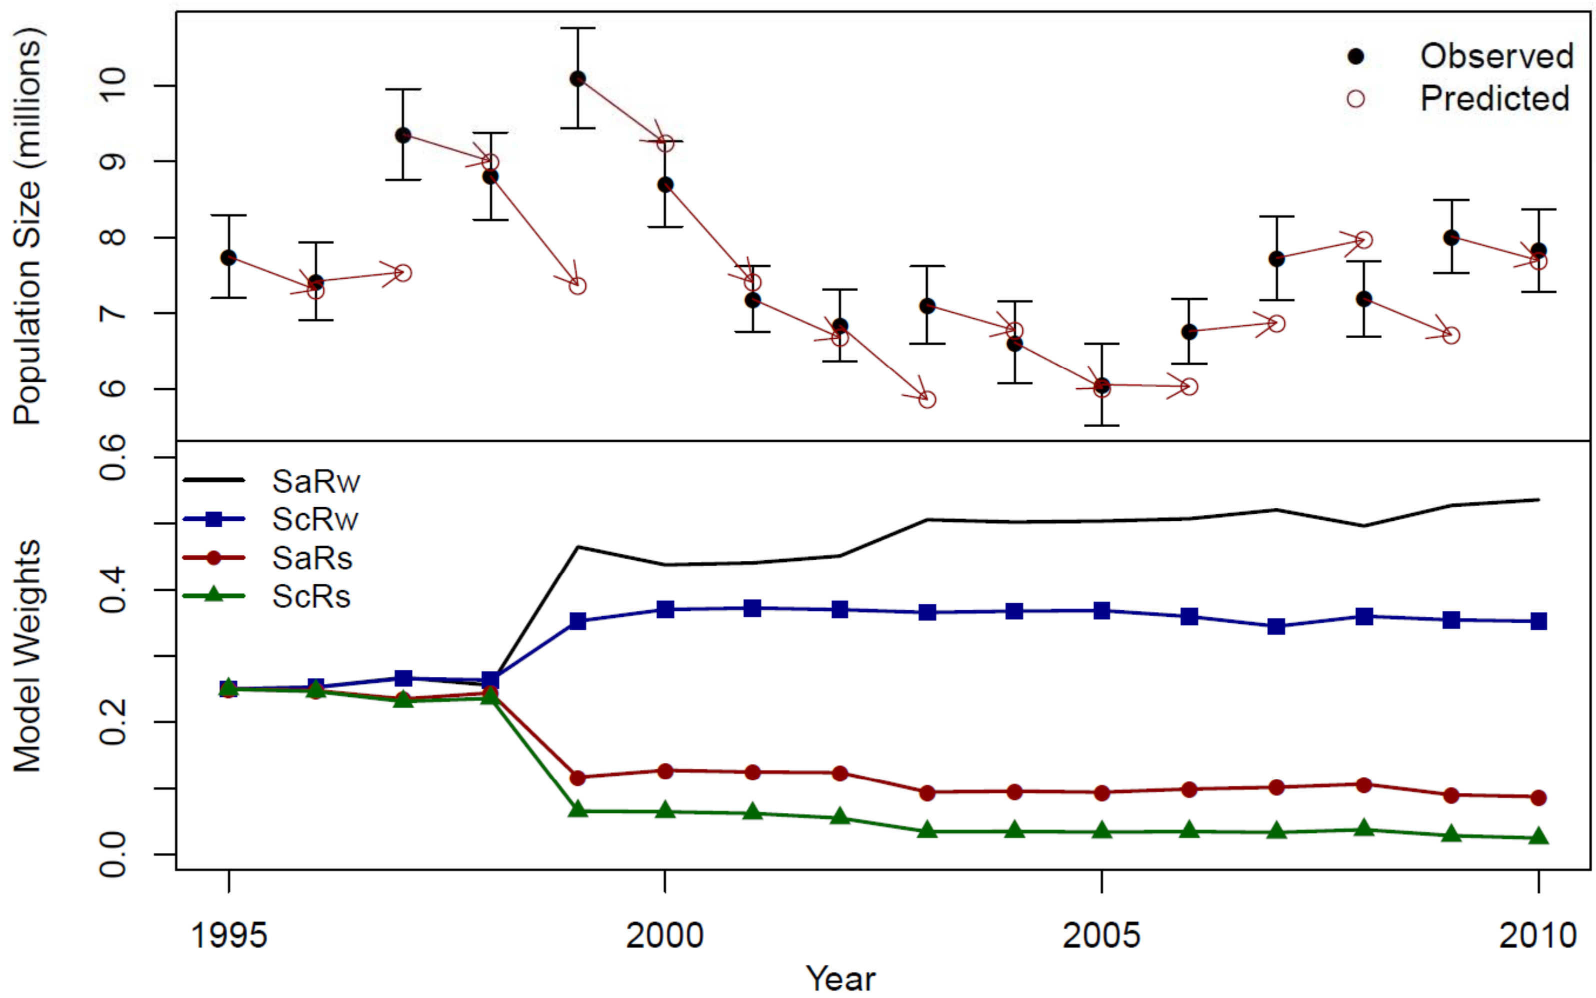

Observed
Predicted
Population Size (millions)
10
9
8
7
6
SaRw
ScRw
SaRs
ScRs
Model Weights
0.6
0.4
0.2
0.0
1995
2000
2005
2010
Year

How fast does learning occur?

What affects the speed of learning?

- Model structure, parameter values
  - Does the set include a good approximating model?
  - 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

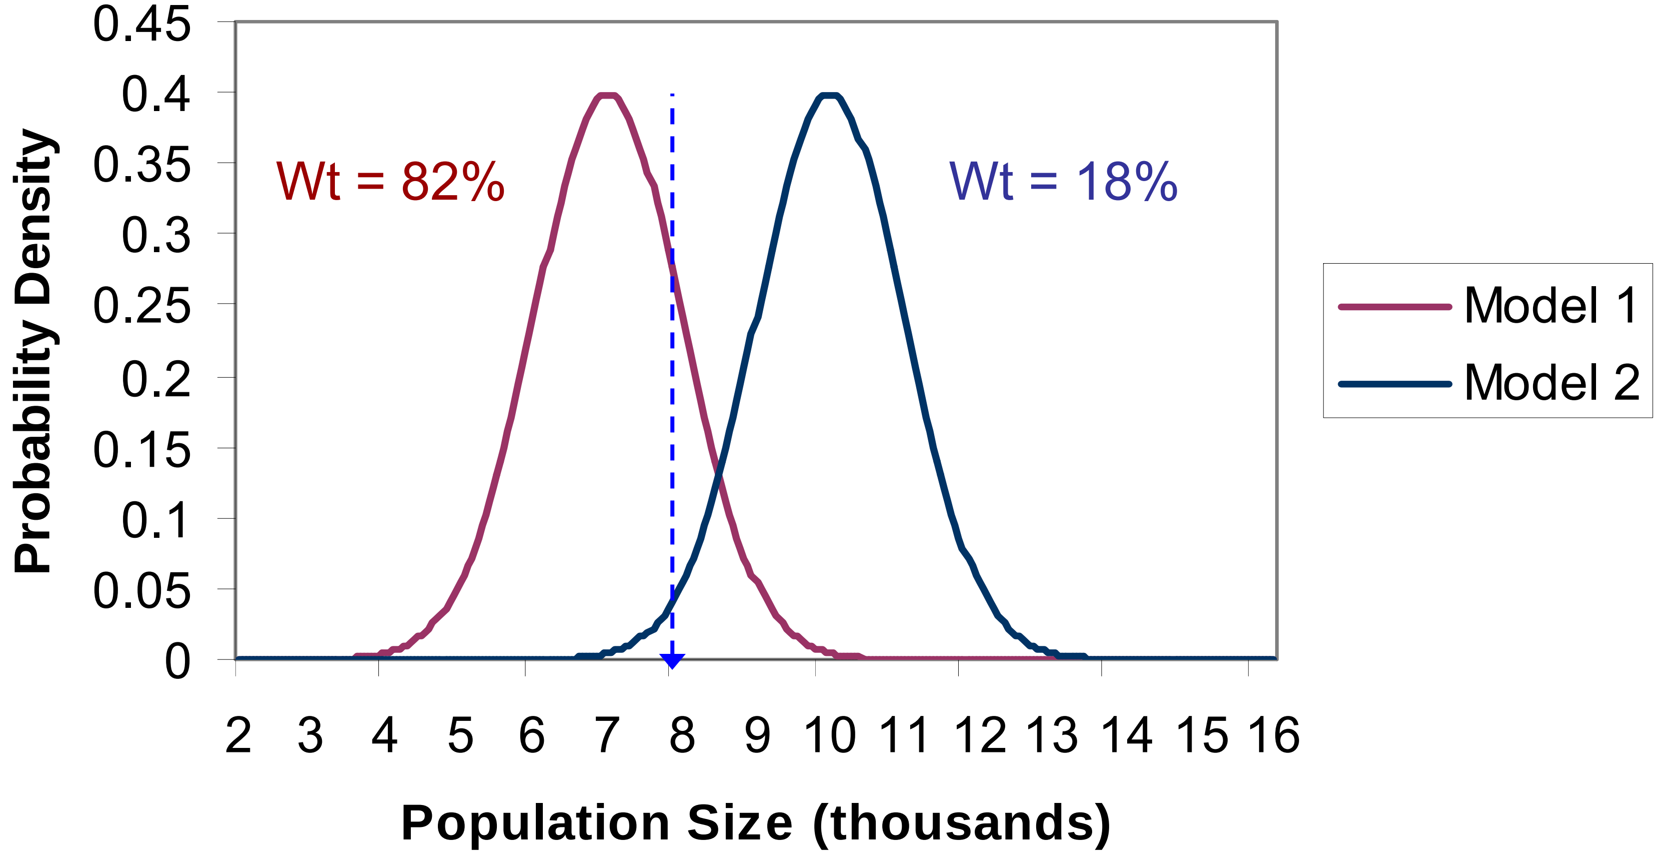

Model Predictions

(Due to greater imprecision in model or monitoring data, or stochasticity)

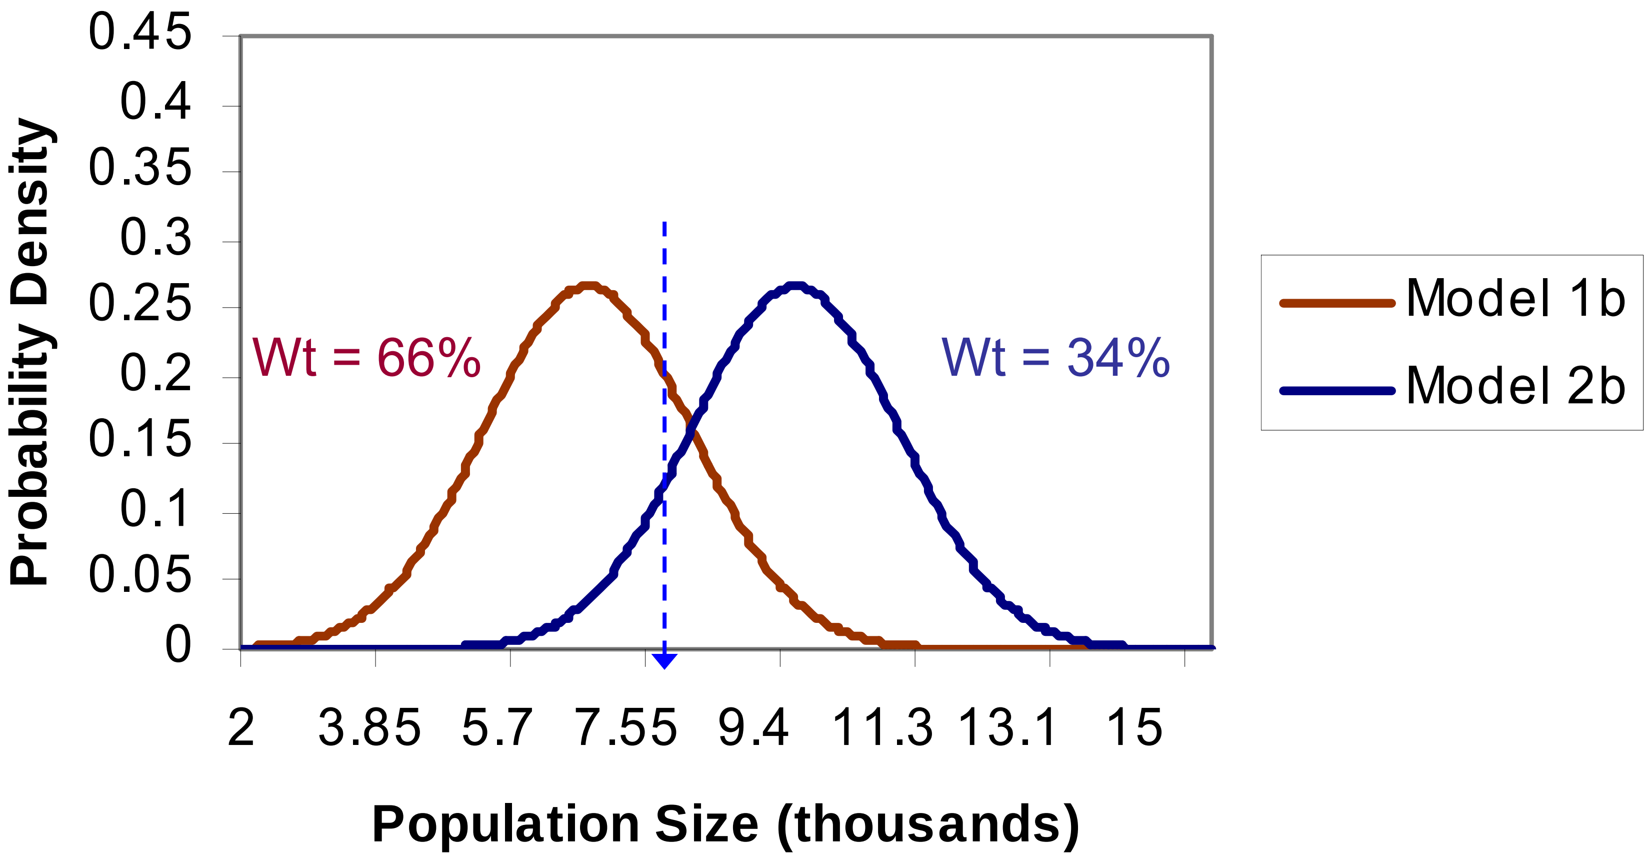

Model Weights

(Predictions for Model 1 negatively biased)

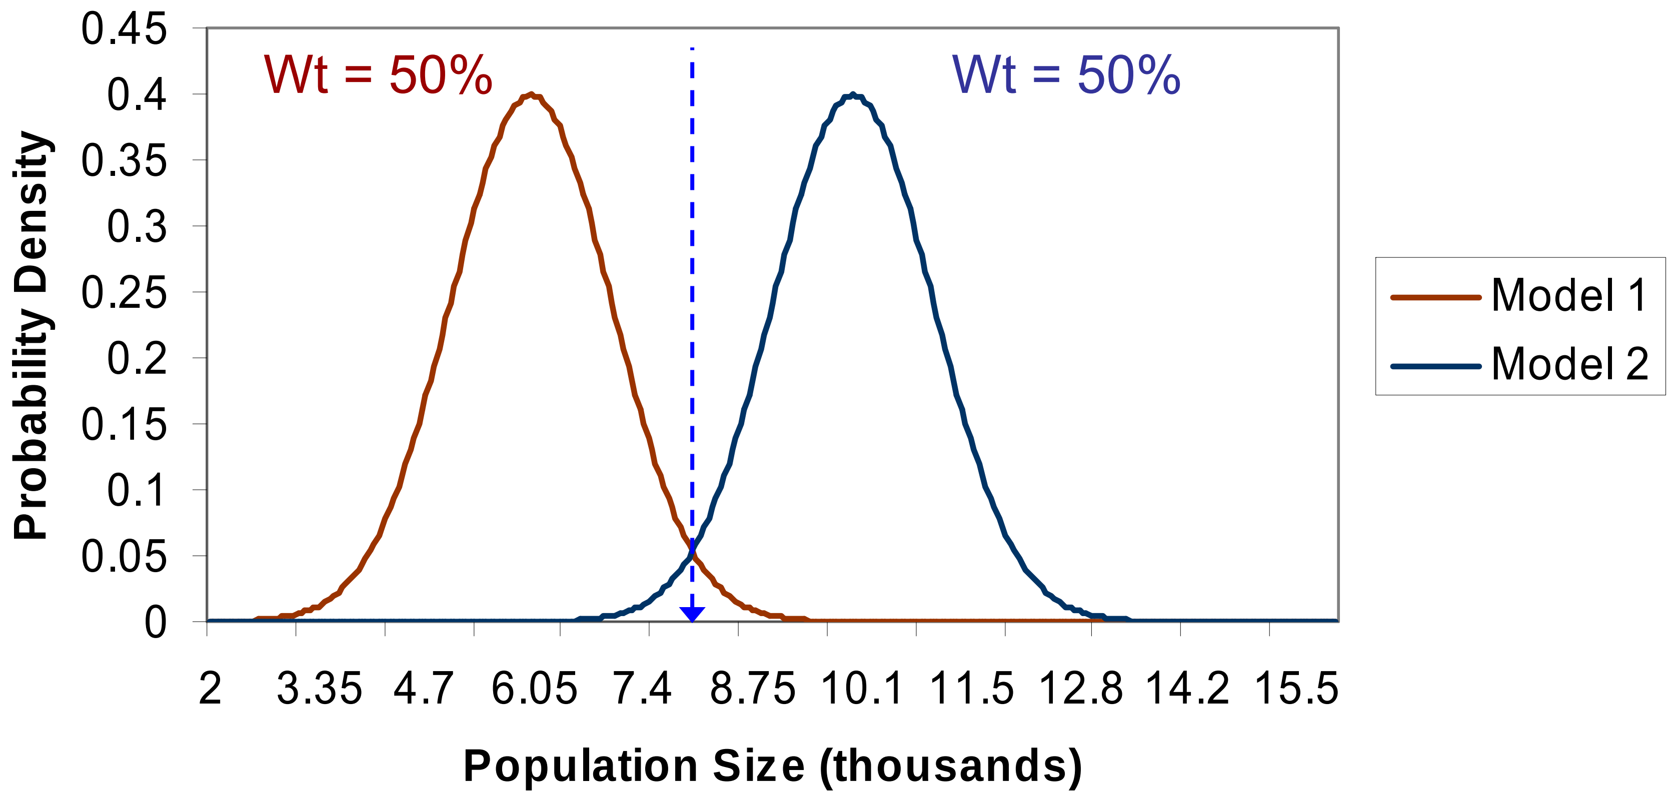

Model Weights

(Predictions for both models negatively biased)

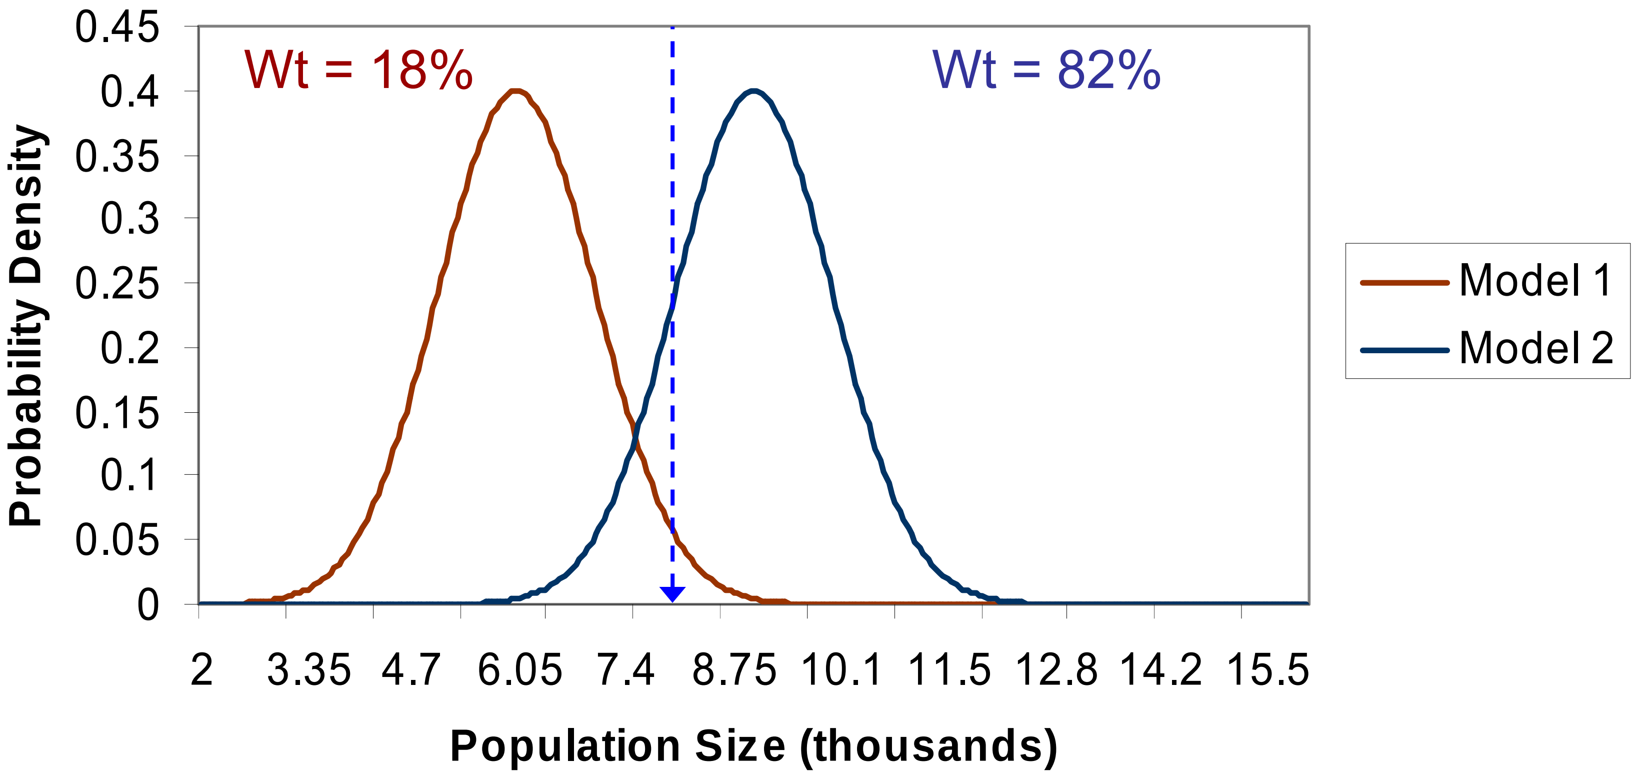

Can we measure the cost of poor monitoring?

Driving in Fog

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

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

USGS
science for a changing world

State-specific decision making

System states
- Current physical and biological conditions of a managed system

Decisions
- Candidate management actions

Rewards
- Expected management gain for given decision and system state

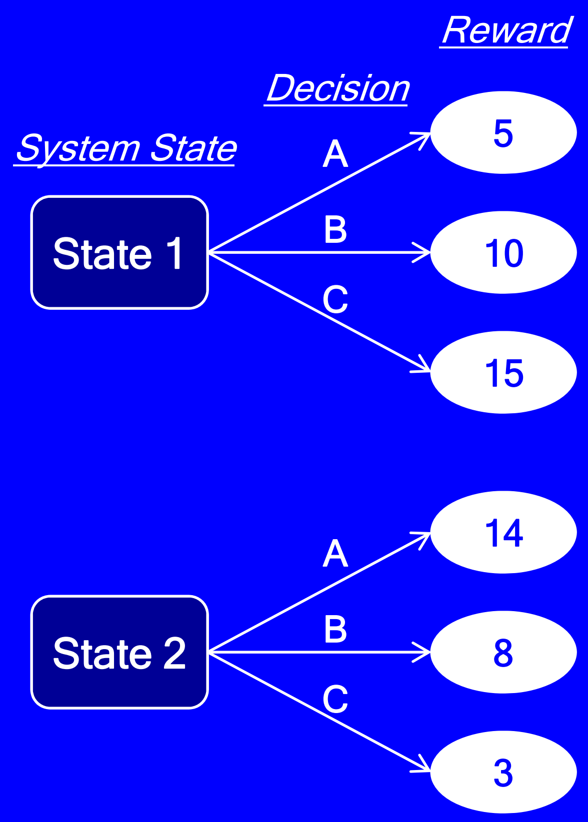

Moore and Kendall, TWS 2006

State-specific decision making

The decision you make depends on how you see the system

<u>Case 1: True state is observable</u>

Decision “A” is best

Decision gain of 14 units

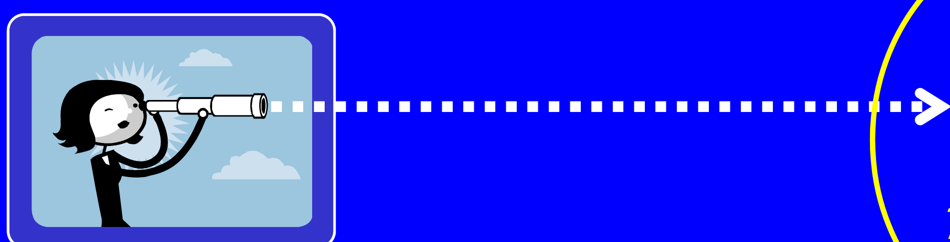

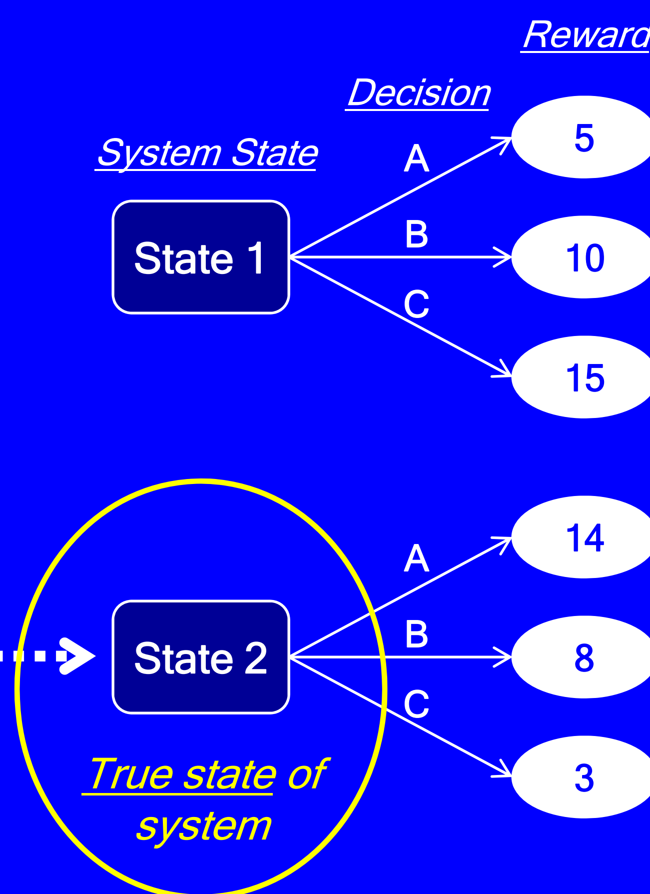

Moore and Kendall, TWS 2006

State-specific decision making

The decision you make depends on how you see the system

<u>Case 2: True state *not* observable</u>

Decision “C” is apparently best

Decision *cost* of 14 - 3 = 11 units

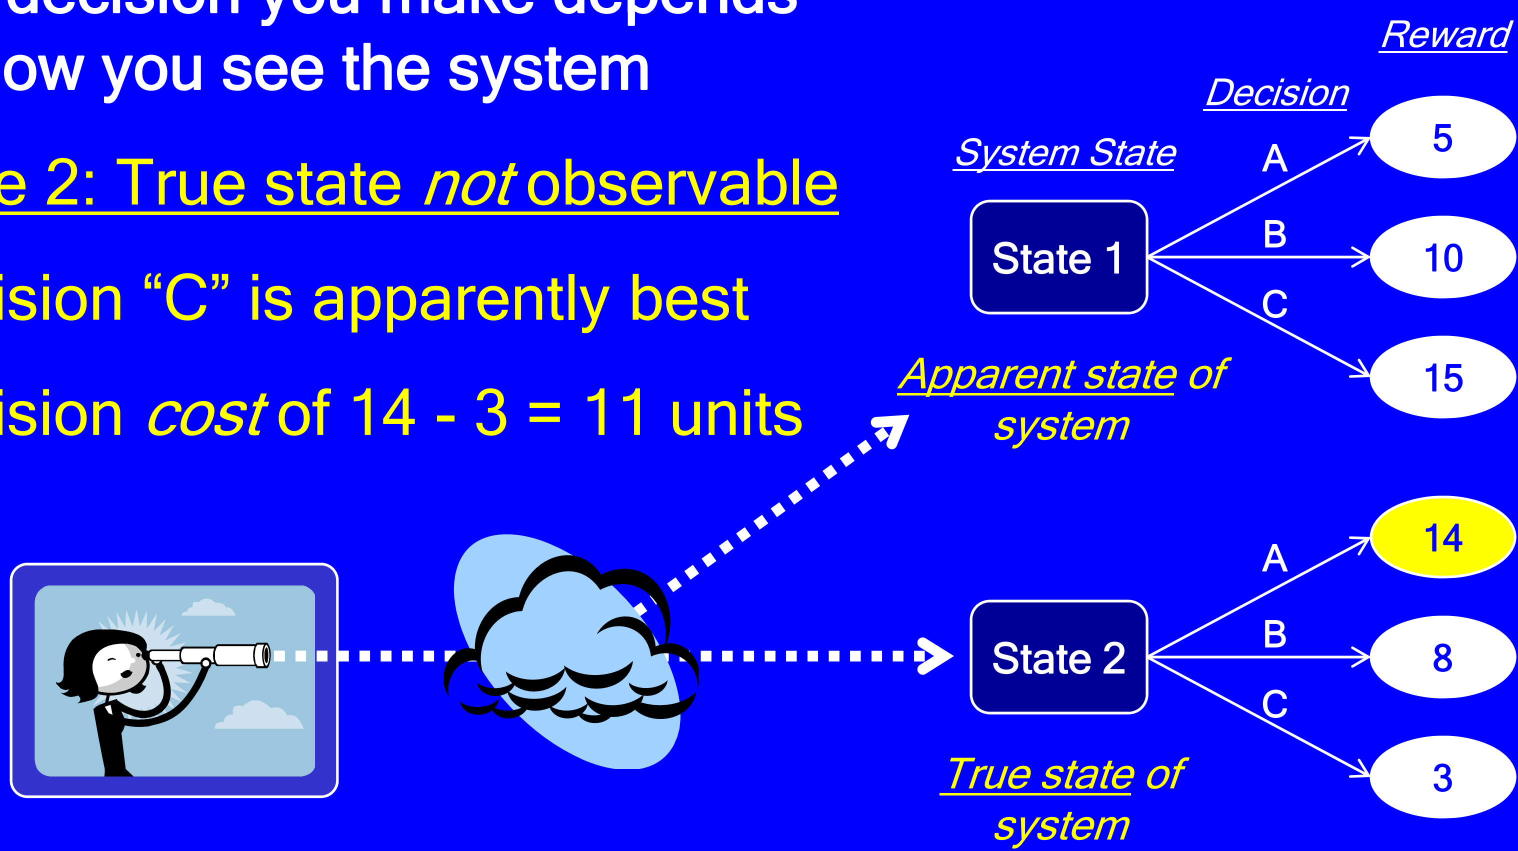

Moore and Kendall, TWS 2006

Partial observability I

- 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

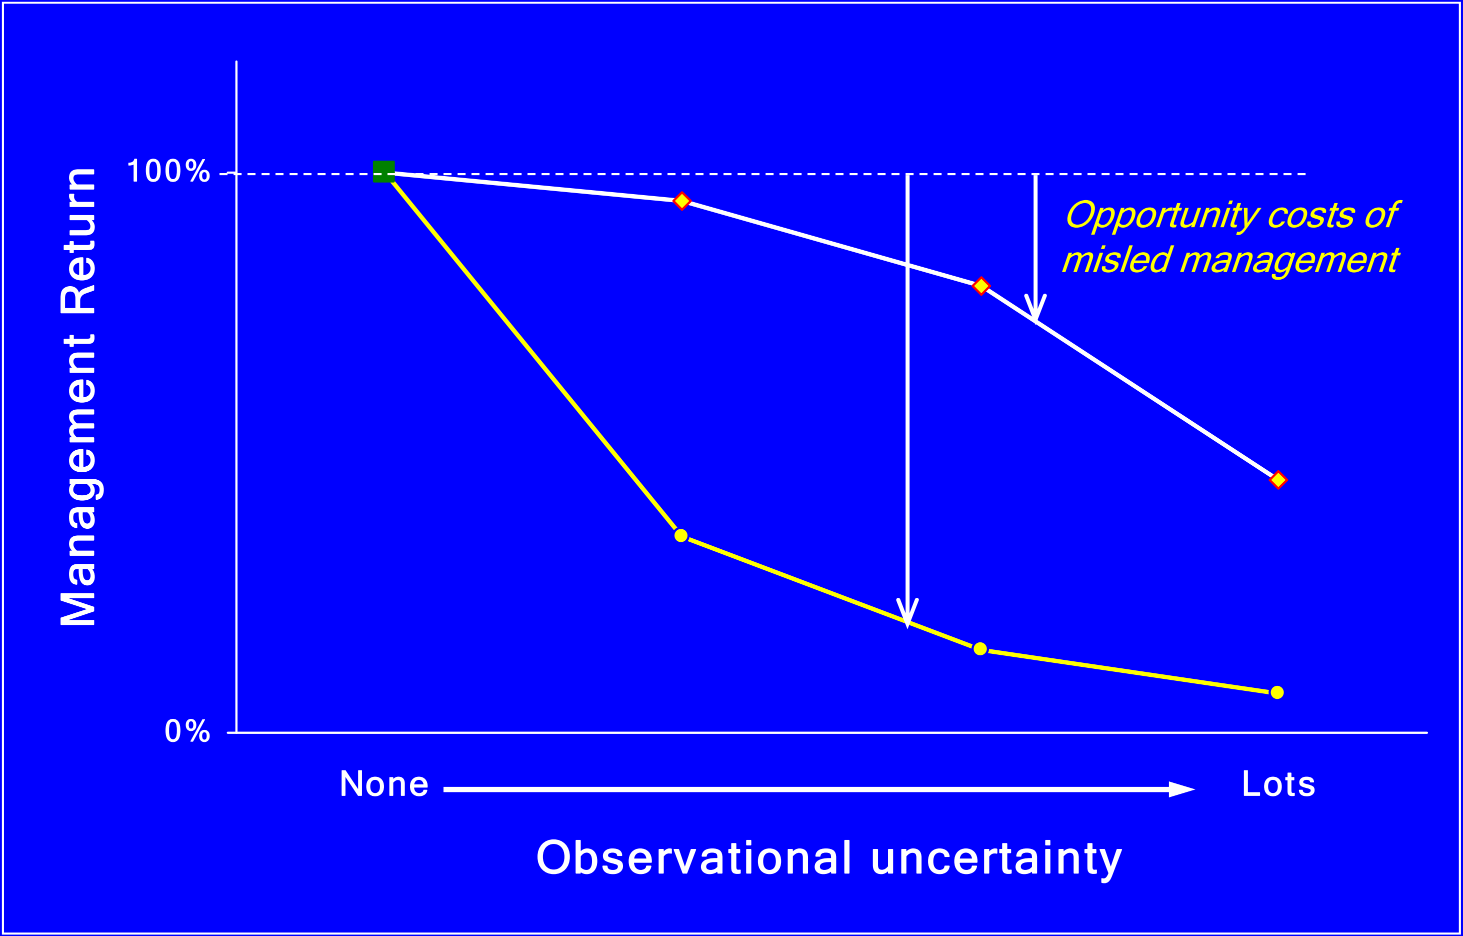

Moore and Kendall, TWS 2006

Partial observability II

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

Moore and Kendall, TWS 2006

Partial observability and model identification

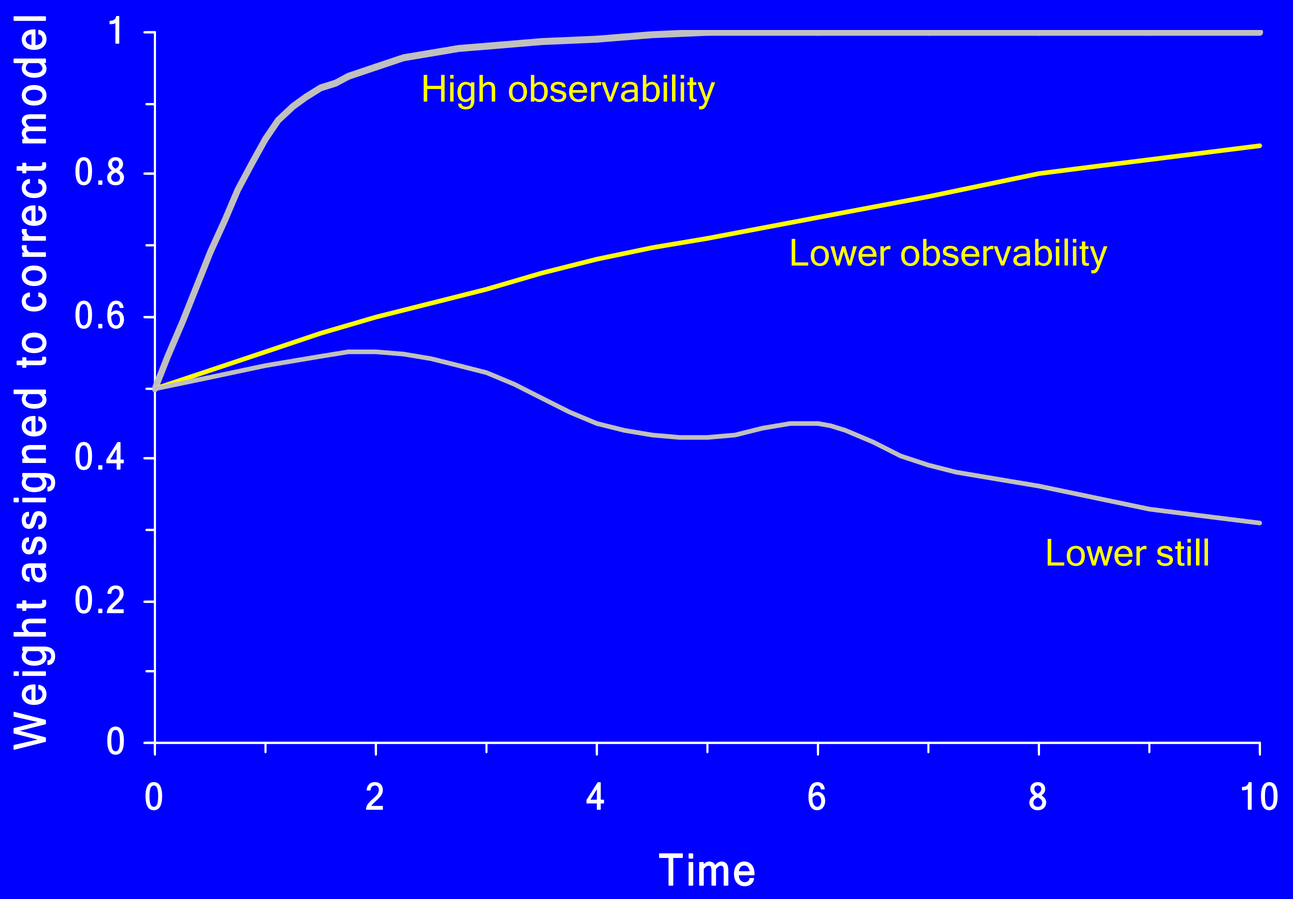

Moore and Kendall, TWS 2006

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 (Oz)

- 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
- Could lead to adaptive monitoring design:
  - Value of reducing uncertainty is high → monitoring intensity increases
  - Value of reducing uncertainty is low → monitoring intensity decreases

Moore and Kendall, TWS 2006

Approach to Optimization

Approaches to optimization

- 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
    - (2) Passive ARM under "best" model(s) based on (1)

USGS

U.S. FISH & WILDLIFE SERVICE

Speed of learning also function of objectives, approach to optimization

low

focus on learning

high

Passively Adaptive

Actively Adaptive

(anticipates benefit of learning to mgmt. objectives)

Experimentation

high

focus on mgmt. objectives

low

Can learn faster with spatial replication

Mgmt Area 1	Mgmt Area 2
Action A	Action B

Mgmt Area 3	Mgmt Area 4
Action B	Action A

Robustness: model vs. 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 About Learning: I

- 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

Conclusions About Learning: II

- 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)

Why bother to learn in ARM?

- Expected value of perfect information (EVPI) compares:
  - weighted average of model-specific maximum values, across models (omniscience)
  - maximum of an average of values (based on average model performance; value under best nonadaptive decision)

Why bother to learn in ARM?

$$EVPI = \sum_i p_i(t) \left\{ \max_{A_t} V_i(A_t|x_t) \right\}$$

$$- \max_{A_t} \sum_i p_i(t) V_i(A_t|x_t)$$