Session 4 - how to learn in the context of AM

So far we have assumed that we were able to characterise the dynamics of the system ..

In many domains, we do not have access to the dynamics but still need to make decisions

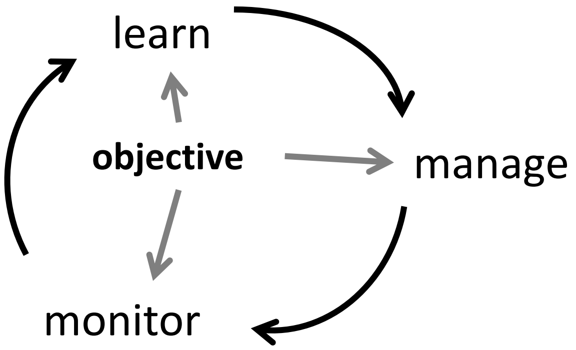

Adaptive management provides a solution. Adaptive management is “learning by doing”. Decisions are selected to achieve a **management objective** while simultaneously gaining information to **improve future management outcomes** (Walters and Hilborn 1976).

Main questions to address before undertaking an AM approach

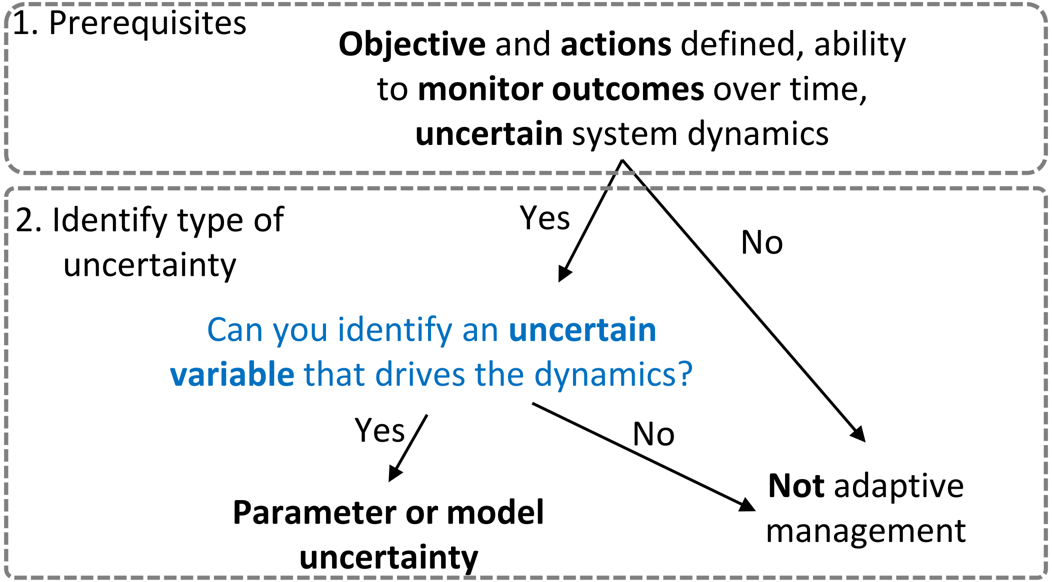

Adaptive management deals with two types of "structural" uncertainty

1) **Parameter uncertainty**:
   e.g. survival, growth, probability of success

2) **Model uncertainty:**
   e.g. competing scenarios, Sea Level Rise, expert opinions, population dynamics.

Can you provide an example of uncertain information for your system? Unknow parameter? Unknow model?

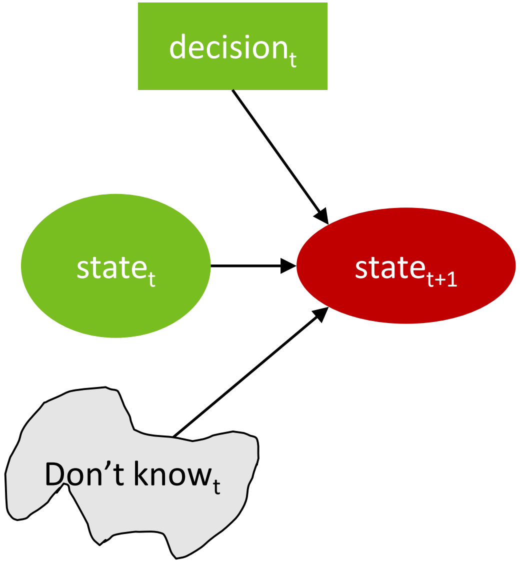

Main questions to address before undertaking an AM approach

Evaluate the need for AM -

is the value of information >0 ?

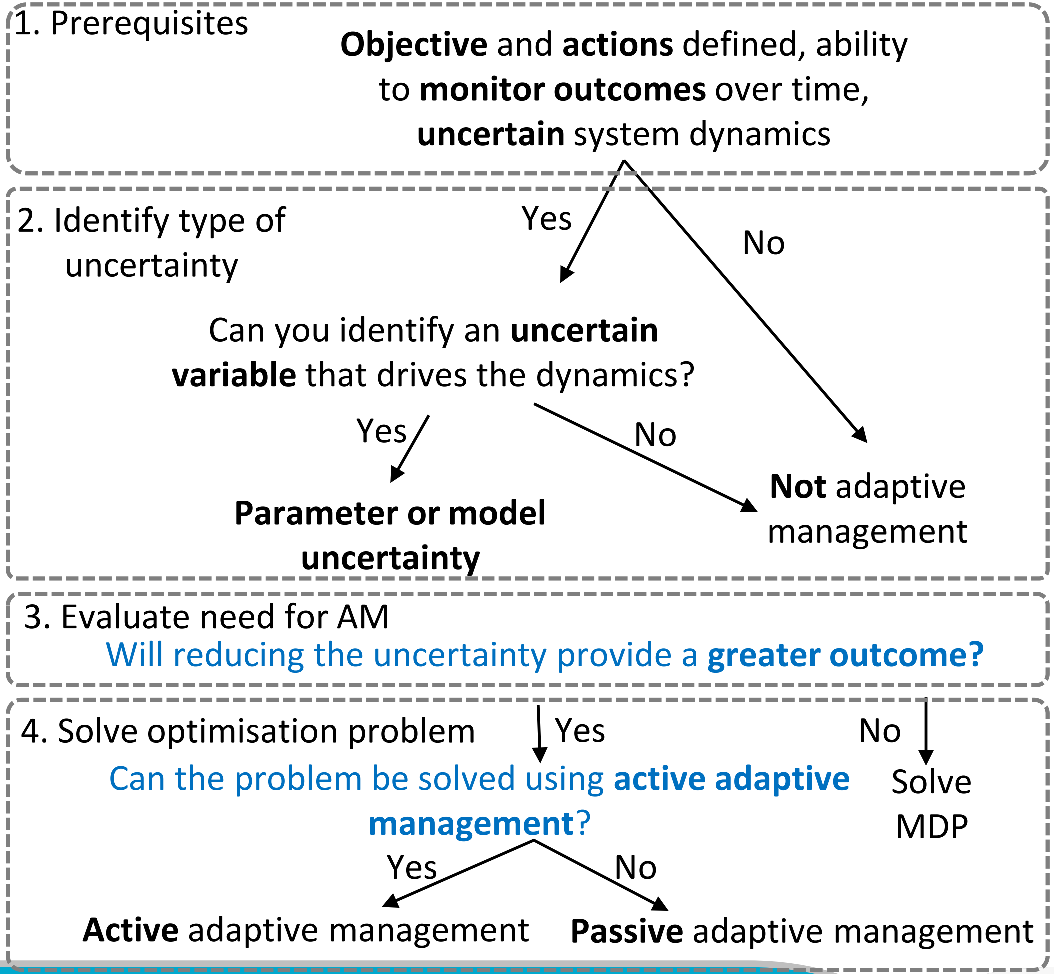

Passive adaptive management provides the best actions given our current knowledge … Learning occurs independently.

Bayes theorem

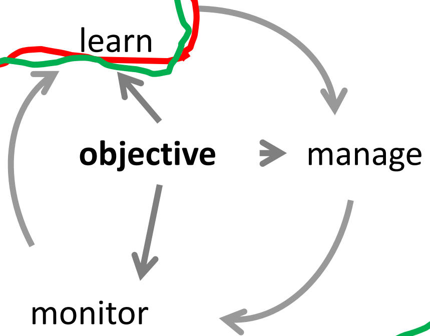

R. Bellman

Stochastic dynamic programming

heuristics
easier to solve
(certainty equivalence principles)

Active adaptive management provides the best actions given our current knowledge … AND what we will learn in the future

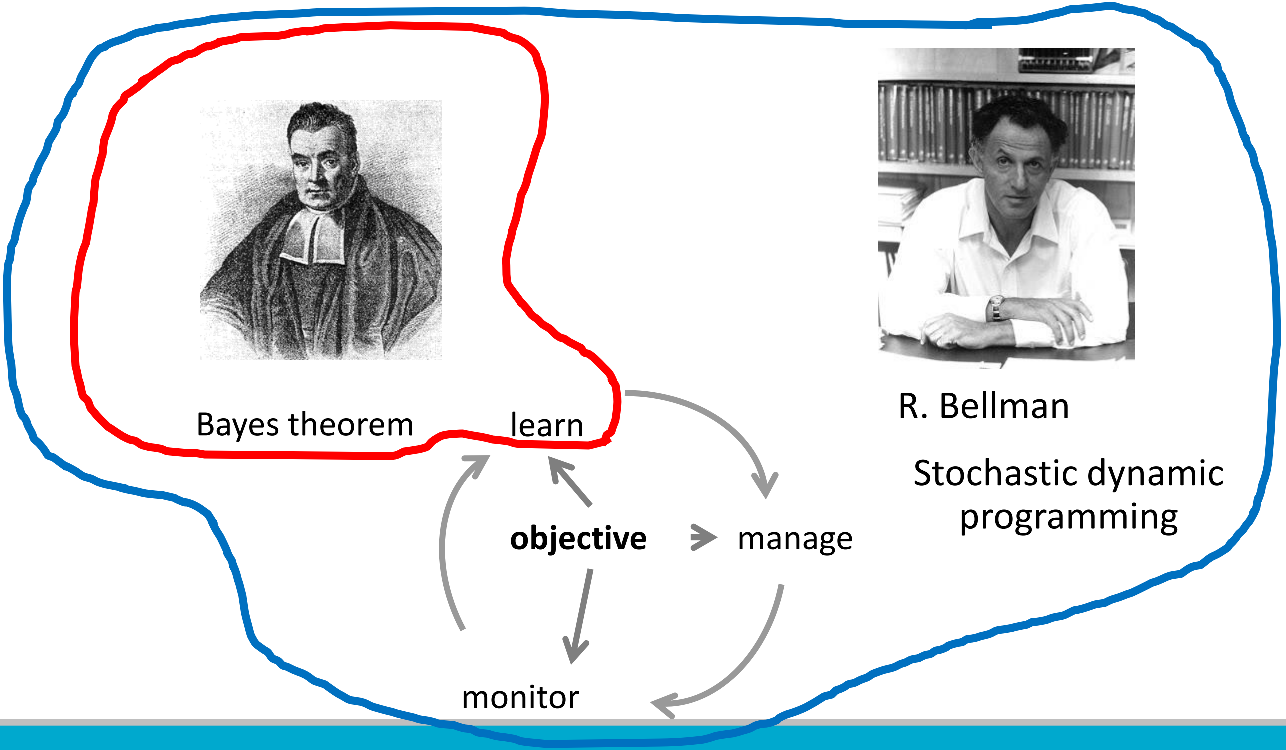

Optimal but difficult to solve

We can optimise the way we perform adaptive management or 'learning by doing'- but how?

How to represent our current knowledge and how we can 'learn':

1. Choose a sufficient statistic
2. Update sufficient statistic

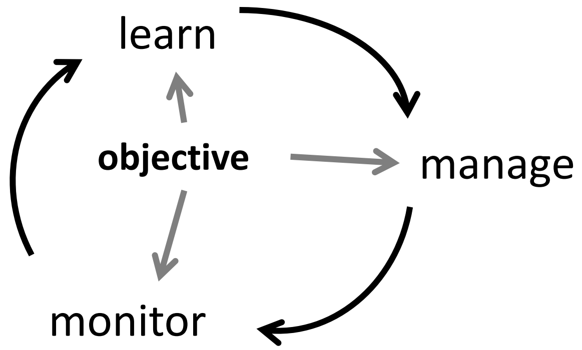

Adaptive management problems differ from classical MDPs because:

- The value of a parameter or the true model is **hidden from the decision maker (Y_t)**
- The value of a parameter or the true model **influences the dynamics of the system (S_t) and best action (a_t)**.

The optimal policy π* (strategy A*(S)) depends on both the observable state variable and the value of the hidden variable.

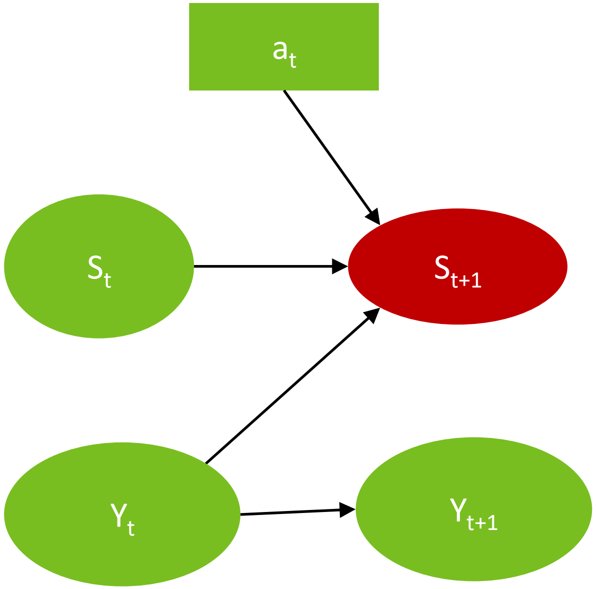

Sufficient statistics are key to Adaptive management

The value of the hidden variable must be estimated using the history of observations and actions:

$s_{t-n}, a_{t-n}, s_{t+1-n}, a_{t+1-n} \ldots s_t \rightarrow a_t$

Because it is not feasible to remember the complete past history of observations and actions, **sufficient statistics** are used (Bertsekas 1995, p 251; Fisher 1922).

Sufficient statistics allow us to retain data without losing any important information.

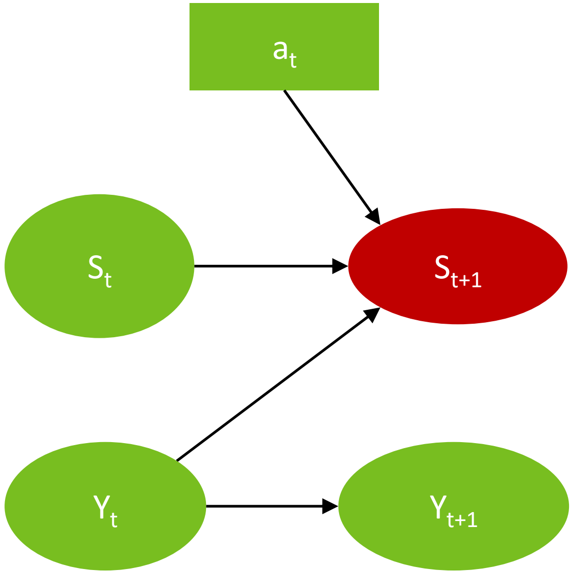

Desirable properties of sufficient statistics

To be useful in adaptive management problems, sufficient statistics must:

- obey the Markov property;
- easy to represent;
- easy to update.

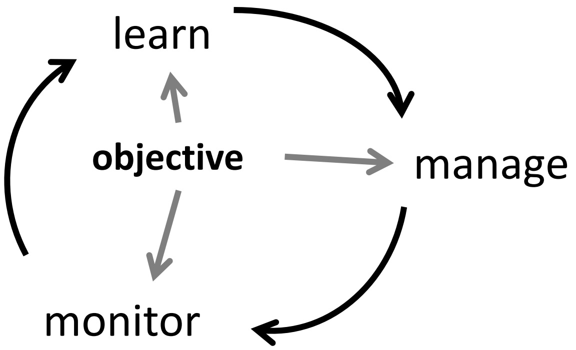

Finding sufficient statistics that best represent uncertain variables is a central and long standing challenge of adaptive management (Walters 1986).

Hidden variables can take finite or infinite number of values

1) **Parameter uncertainty**:
   e.g. survival rate, growth rate, probability of success

   ~takes infinite number of values

2) **Model uncertainty:**
   e.g. competing scenarios, SLR, expert opinions

   ~takes finite number of values

	Finite # values	Infinite # values
Survival rate		✓
Growth rate		✓
Prob. of success		✓
Competing scenarios	✓	
Expert opinions	✓	

When hidden variables take finite values: belief states.

For problems with hidden variables that can take **finite** values:

- **belief states** are widely used sufficient statistics.

Belief states are probability distributions over finite quantities

and can be updated using Bayes' theorem.

Domain	Objective	Belief states over
Sustainable harvest (Williams et al. 1996)	Maximize long-term cumulative harvest of waterfowl, above a certain density threshold	Two alternative models of population response to harvest and survival
Conservation (Moore et al. 2011)	Maximize time-discounted plant population size across years without burning	Two models describing the juvenile plant stage response to burning
Climate change, conservation (Nicol et al. 2015)	Maximize migratory shorebirds populations across space and subject to seal level rise	Three models representing alternative responses to management under sea level rise

CSIRO

An application

Climate change, conservation (Nicol et al. 2015)	Maximize migratory shorebirds populations across space and subject to seal level rise	Three models representing alternative responses to management under sea level rise

S= *breeding population*

Y={0m, 1m,2m}; //uncertain

Belief_t=[b(0m), b(1m), b(2m)]

b_t=[1/3, 1/3, 1/3];

b_{t+1}=[0.4, 0.2, 0.2];

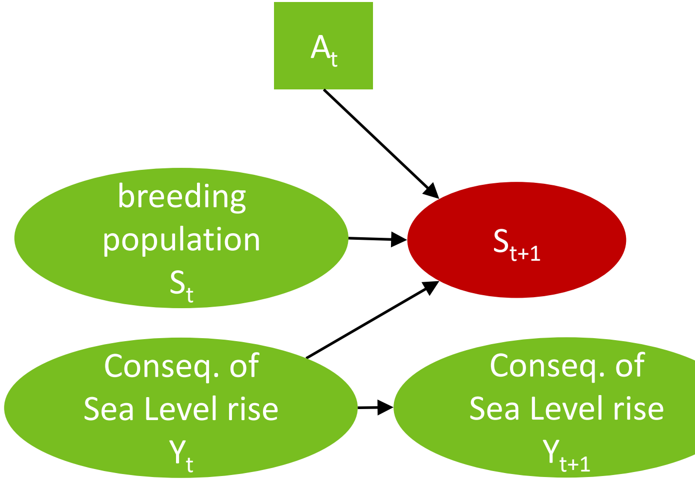

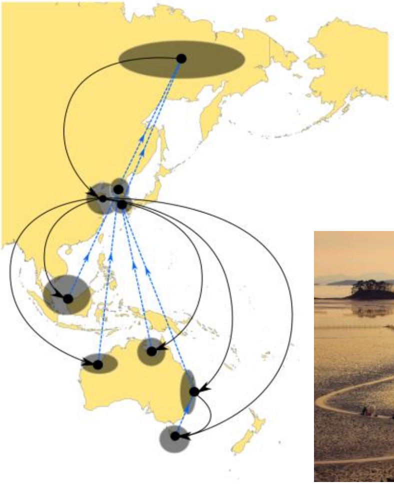

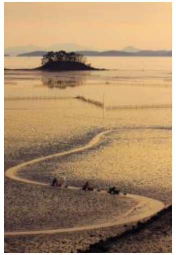

An application

Climate change, conservation ([Nicol et al. 2015](#))

Maximize migratory shorebirds populations across space and subject to seal level rise

Three models representing alternative responses to management under SLR

S= *breeding population*

Y={0m, 1m,2m}; //uncertain

Belief_t=[b(0m), b(1m), b(2m)]

b_t=[1/3, 1/3, 1/3];

b_{t+1}=[0.4, 0.2, 0.2];

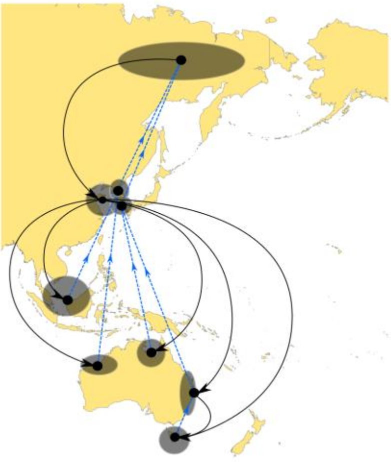

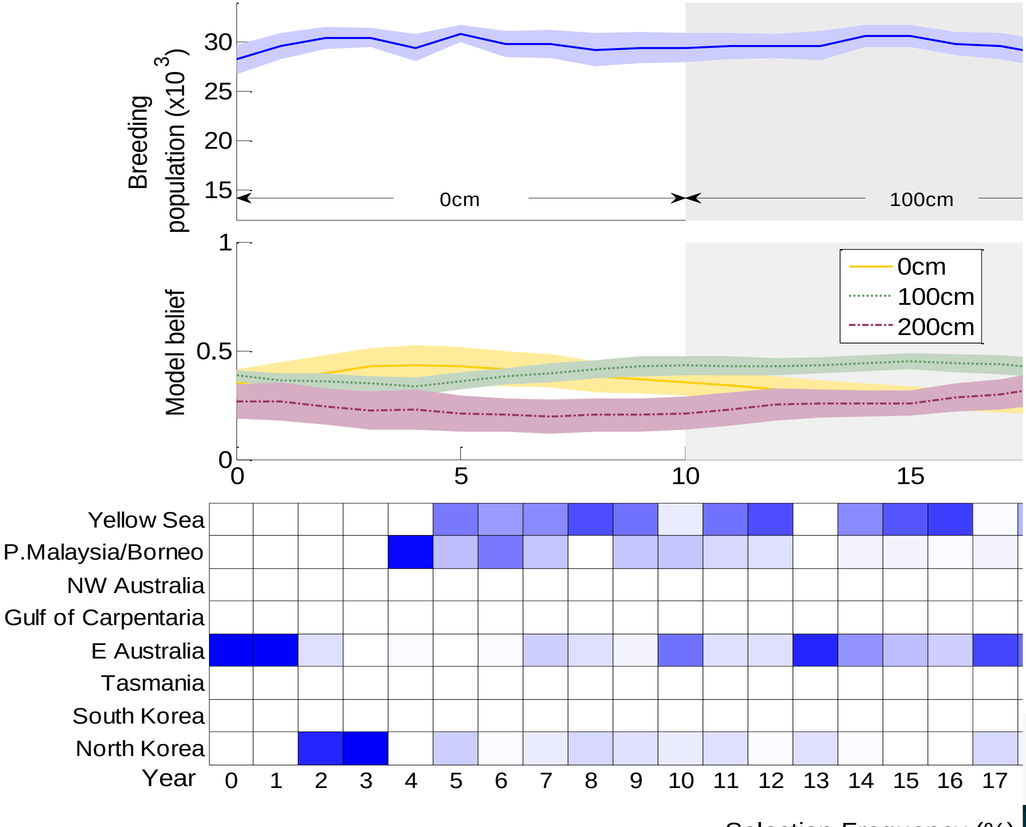

Bayes' rule is the underlying mechanism for learning in all AM problems

P(B|A)=P(A|B)P(B)/P(A)

↓

$$b_{t+1}(y|s_t, a_t, s_{t+1}, b_t) = \frac{\mathrm{P}_y(s_{t+1}|s_t, a_t)\, b_t(y)}{\sum_{y \in Y} b_t(y)\, \mathrm{P}_y(s_{t+1}|s_t, a_t)}.$$

where $\mathrm{P}_y(s_{t+1}|s_t, a_t)$ is the state transition probability assuming that the true model is y. The discrete belief value $b_t(y)$ is interpreted as the probability that y best describes system dynamics of the available models.

Choice of sufficient statistics influences the AM optimisation approach

For problems with hidden variables that can take **finite** values:

- **belief states** are widely used sufficient statistics.

Belief states are probability distributions over finite quantities and can be updated using Bayes' theorem.

Does the uncertain variable (parameter or model) take **finite** number of values?

no

yes

Augment state space with the **sufficient statistic belief states** over the finite values of the uncertain variable

Solve **passive adaptive management** problem using MDP

Solve **active adaptive management** problem using discretized belief MDP or POMDP

Finding sufficient statistics?

For problems with **hidden variables** that can take **infinite** values:

- sufficient statistics that take **finite** values facilitate the use of fast and accurate solution methods.

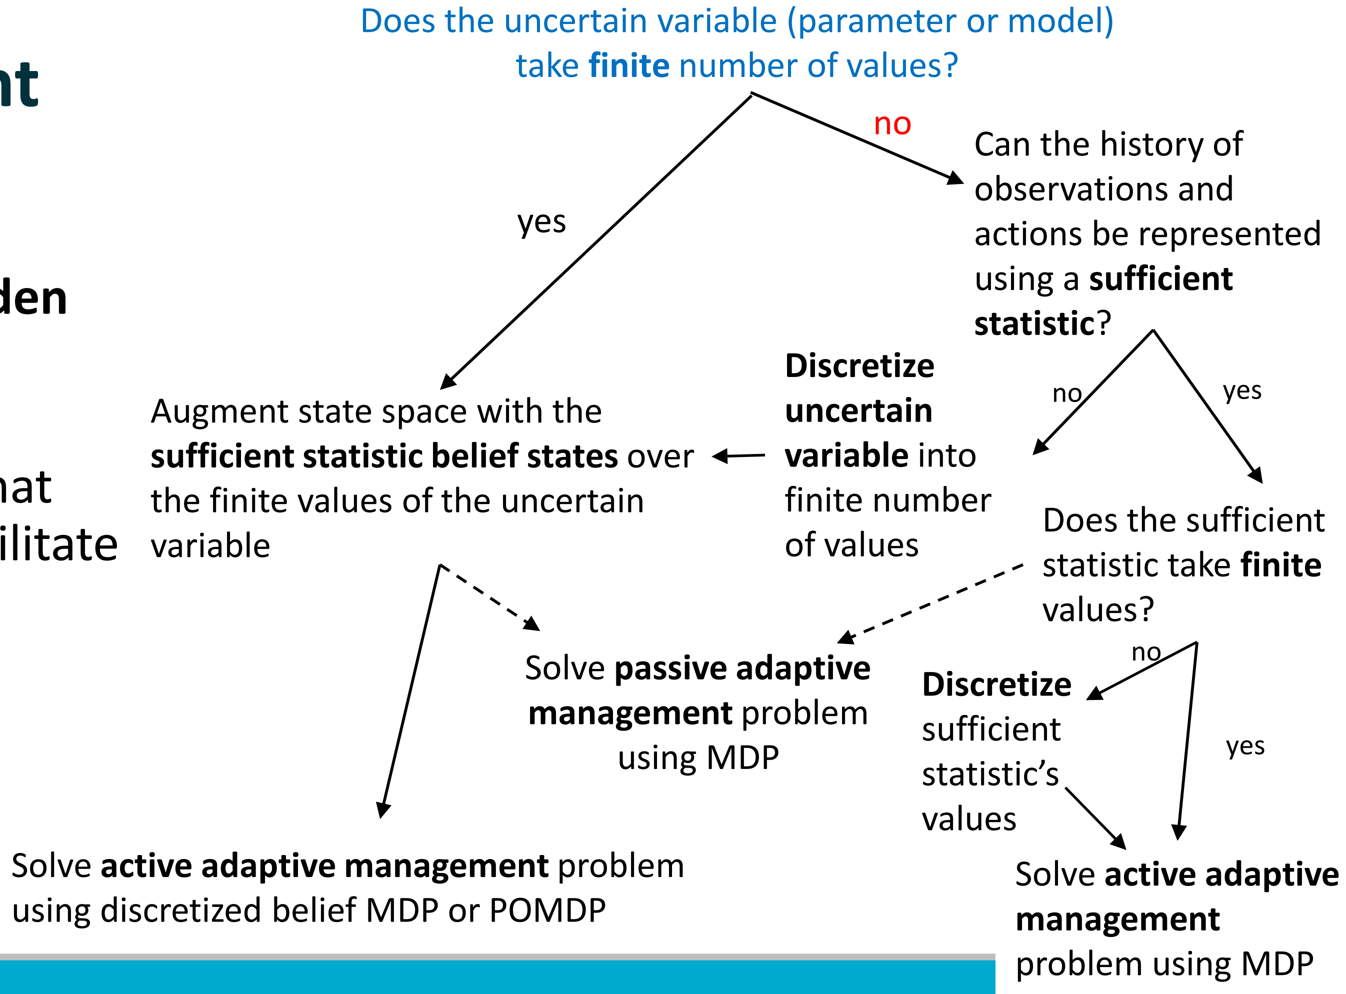

Bayes rule for infinite values

- The distribution $b_t(\theta)$ represents the values of parameter θ at time t as a probability density function: “belief in θ”
- Observing the system response to management actions between times **t and t + 1** provides information that can be used to update this belief.
- Bayes’ theorem provides a means of updating distribution $b_t(\theta)$ as the system is managed (a_t) in a given configuration (s_t) and data are gathered (s_{t+1}):

$$b_{t+1}(\theta|s_t, a_t, s_{t+1}, b_t) = \frac{P_\theta(s_{t+1}|s_t, a_t)b_t(\theta)}{\int_\theta \; P_\theta(s_{t+1}|s_t, a_t, \theta)b_t(\theta)d\theta},$$

$P_\theta(s_{t+1}|s_t, a_t)$ is the state transition probability assuming that the true parameter value is θ. **Useful sufficient statistics for $\boldsymbol{b_t(\theta)}$ can be found when $\boldsymbol{b_t(\theta)}$ is a conjugate prior for $\boldsymbol{P_\theta(s_{t+1}|s_t, a_t)}$.**

Beta distribution with binomial updating is a must try …

Domain	Objective	Uncertain parameter
Forestry, Conservation (McCarthy and Possingham 2007; Moore and McCarthy 2010)	Maximize the expected number of successes over a specified number of time periods	**Probability of success** of management action defined as Beta distribution with binomial updating
Fisheries (Hauser and Possingham 2008)	Maximum long-term fish stock harvest	**Recovery rate** after stock collapse, modeled as a Beta distribution with binomial updating
Conservation, translocation (Rout et al. 2009)	Translocation of threatened species, choosing between introducing to two sites	**Mortality rate** at one site represented as Beta distribution with binomial updating

- Uncertain parameter p in [0,1] e.g. management success, rate.
- Uncertainty surrounding p can be represented as a beta distribution.
- Given a Beta(∝, β) prior for p, the posterior is a beta distribution with new parameters ∝ + R (number of successes) and β+N-R (number of failures).
- Consequently, ∝ and β can be used as sufficient statistics.
- The transition probabilities are derived for all possible value of ∝ and β.
- The optimal policy matches an action to a population size and values of ∝ and β.

Examples of discrete conjugate distributions and sufficient statistics with potential for use in adaptive management with parameter uncertainty (Chades et al 2017).

Data updating process	Uncertain parameter	Conjugate Prior distribution and sufficient statistics	Posterior distribution and update of sufficient statistics	Posterior predictive
Binomial $x \sim \text{Bin}(n,p)$	p	$p \sim \text{Beta}(\alpha, \beta)$ α β	$p' \sim \text{Beta}(\alpha+x, \beta+n-x)$ $\alpha \rightarrow \alpha+x$ $\beta \rightarrow \beta+n-x$	Beta-Binomial $x \sim \text{BetaBin}(n, \alpha, \beta)$
Negative Binomial $x \sim \text{NB}(r, p)$	p	$p \sim \text{Beta}(\alpha, \beta)$ α β	$p' \sim \text{Beta}(\alpha+x, \beta+r)$ $\alpha \rightarrow \alpha+x$ $\beta \rightarrow \beta+r$	$Pr(X = x) = \binom{x+r-1}{x} \frac{B(\alpha+x,\beta+r)}{B(\alpha,\beta)}$
$x \sim \text{Poisson}(\lambda)$	λ	$\lambda \sim \text{Gamma}(k,\theta)$ k θ	$\lambda' \sim \text{Gamma}\left(k + x, \frac{\theta}{1+\theta}\right)$ $k \rightarrow k + x$ $\theta \rightarrow \frac{\theta}{1+\theta}$	Negative Binomial $x \sim NB\left(k, \frac{\theta}{1+\theta}\right)$
$x \sim \text{Geometric}(p)$	p	$p \sim \text{Beta}(\alpha, \beta)$ α β	$p' \sim \text{Beta}(\alpha+1, \beta+x)$ $\alpha \rightarrow \alpha+1$ $\beta \rightarrow \beta+x$	$\Pr(X = x) = \frac{B(\alpha+1,\beta+x)}{B(\alpha,\beta)}$

Some continuous conjugate distributions and sufficient statistics with potential for use in adaptive management with parameter uncertainty (Chades et al 2017).

Data updating process	Uncertain parameter	Conjugate Prior distribution and sufficient statistics	Posterior distribution and update of sufficient statistics	Posterior predictive
Normal x ~ N(μ, σ^2)	μ	m ~ N(μ_0, σ_0^2)	$\mu' \sim N\left(\frac{\mu_0\sigma^2 + x\sigma_0^2}{\sigma^2+\sigma_0^2}, \frac{\sigma^2\sigma_0^2}{\sigma^2+\sigma_0^2}\right)$	$x \sim N(\mu_0, \sigma^2 + \sigma_0^2)$
		μ_0	$\mu_0 \rightarrow \frac{\mu_0\sigma^2 + x\sigma_0^2}{\sigma^2+\sigma_0^2}$	
		σ_0^2	$\sigma_0^2 \rightarrow \frac{\sigma^2\sigma_0^2}{\sigma^2+\sigma_0^2}$	
Normal x ~ N(μ, σ^2)	σ^2	$\sigma^2 \sim \text{InvGamma}(\alpha, \beta)$	$(\sigma^2)' \sim \text{InvGamma}\left(\alpha + 1, \beta + \frac{(x-\mu)^2}{2}\right)$	$f(x) = \frac{1}{\sqrt{2\pi}} \frac{\alpha}{\beta} \left[\frac{2\beta}{(x-\mu)^2 + 2\beta}\right]^{\alpha+1}$
		α	$\alpha \rightarrow \alpha + 1$	
		β	$\beta \rightarrow \beta + \frac{(x-\mu)^2}{2}$	
Exponential x ~ exp(l)	l	l ~ Gamma(a, b)	$\lambda' \sim \text{Gamma}(\alpha + n, \beta + \sum_{i=1}^n x_i)$	$f(x) = \frac{\Gamma(\alpha+n)}{\Gamma(\alpha)} \frac{\beta^\alpha}{(\beta + \sum_{i=1}^n x_i)^{\alpha+n}}$
		α	$\alpha \rightarrow \alpha + n$	
		β	$\beta \rightarrow \beta + \sum_{i=1}^n x_i$	

CSIRO

Conclusion

- Learn uncertain quantities using sufficient statistics and applying Bayes' theorem;

- Find the optimal adaptive management strategy by augmenting the state space with sufficient statistics and stochastic dynamic programming;

Bayes theorem

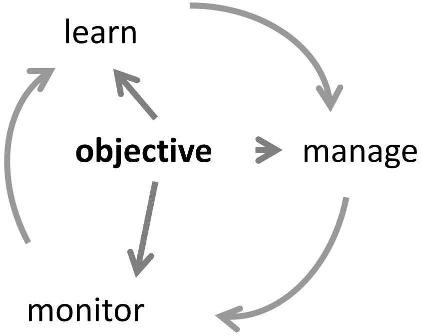

R. Bellman

Stochastic dynamic programming

Session 5:

- Solving AM under model uncertainty

Session 6:

- Solving AM under parameter uncertainty

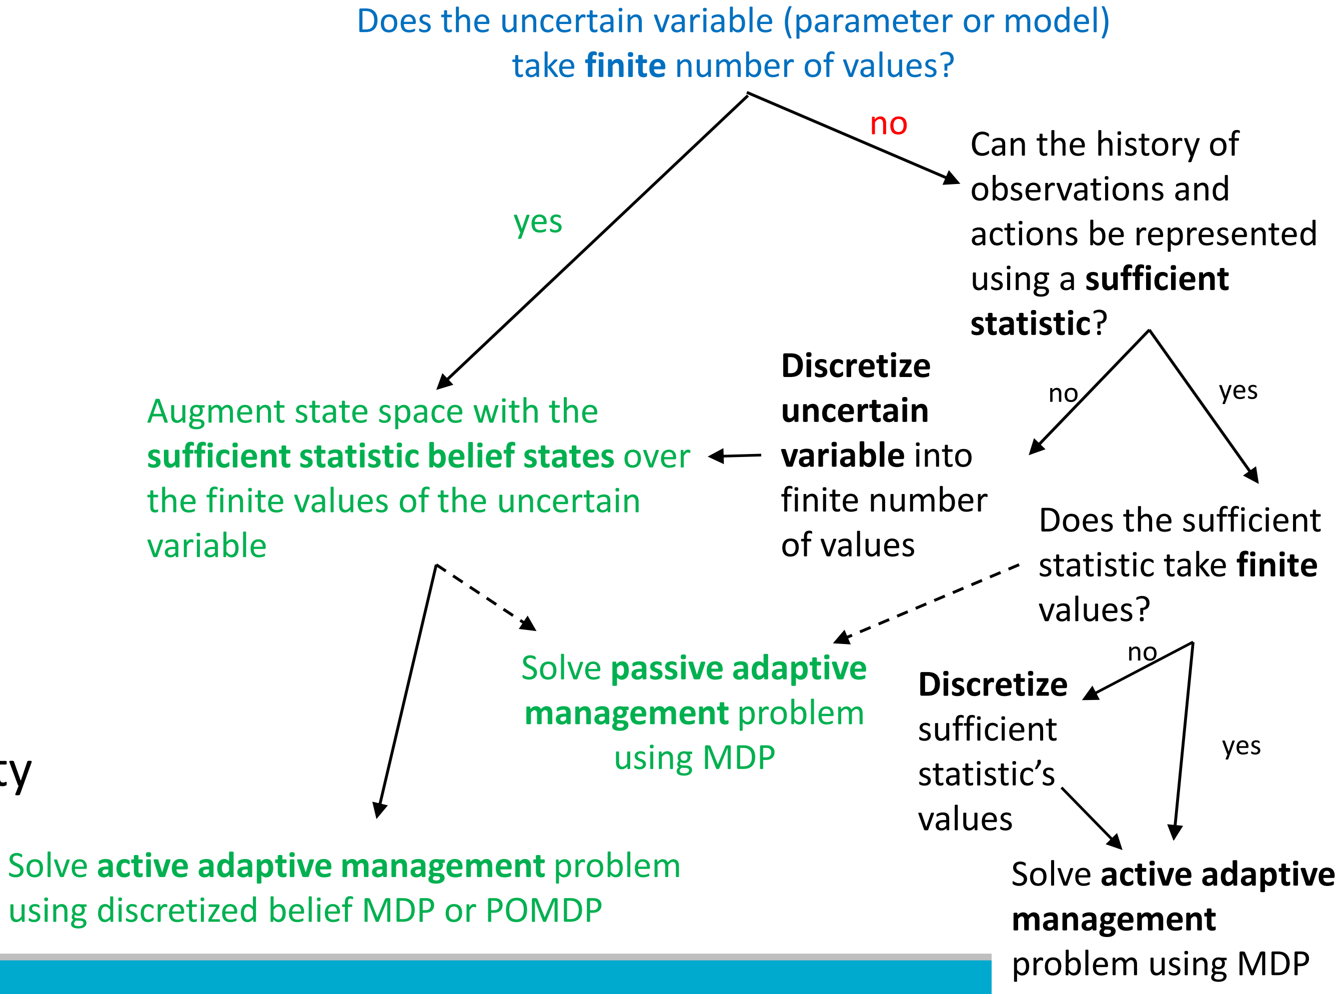

Additional material and references:

Theor Ecol (2017) 10:1–20
DOI 10.1007/s12080-016-0313-0

CrossMark

REVIEW PAPER

Optimization methods to solve adaptive management problems

Iadine Chadès[1] · Sam Nicol[1] · Tracy M. Rout[2] · Martin Péron[1,3] · Yann Dujardin[1] ·
Jean-Baptiste Pichancourt[1] · Alan Hastings[4] · Cindy E. Hauser[2]

CSIRO