Case Study: Native Prairie Adaptive Management in the USFWS Refuge System

Dynamic Decision Making

Case Study Module D

Module Developed by:
Jill Gannon
USGS Northern Prairie Wildlife Research Center

<u>NPAM Framework Components</u>

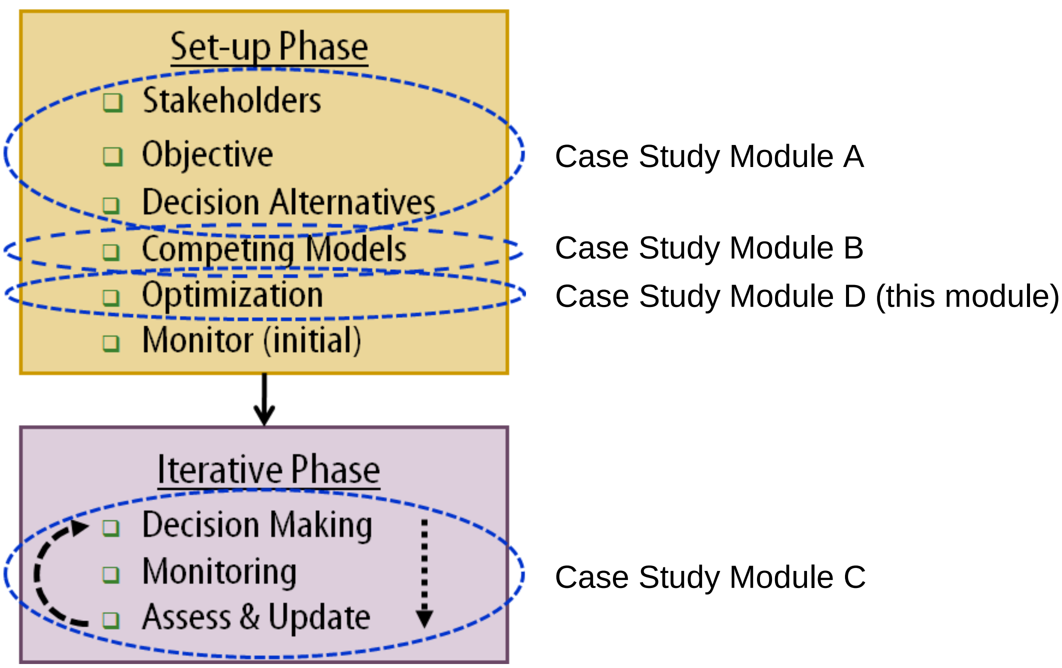

<u>Objectives of this Module</u>

- Introduce elements used in optimization process
  - Management actions and partial controllability
  - Utility and Objective function
- View the output of optimization – the decision table
- Compare optimal decision policies
  - Policies for a known model
  - Policies under model uncertainty

Nature of the Management Decision

- An annual decision must be made per unit
    - Decision is conditional on the state of the unit (vegetation state and defoliation state)
- Today's decision affects the future state of the unit
    - By extension, the decision sets a path of opportunities and/or constraints on all future management decisions
- Management interest is in persistent, abundant native cover, over a long time frame
- Must acknowledge uncertainty about system response to management
- Must account for stochastic effects
- Managers have
    - A clear management objective
    - A clear set of alternative management decisions (actions)
    - A system of monitoring
- ➢ All facets of this decision problem point to solution through adaptive stochastic dynamic optimization

System State Structure

Vegetation State Structure

Native Cover	Dominant Invasive: SB	CO	KB	RM
60 – 100%	1	2	3	4
45 – 60%	5	6	7	8
30 – 45%	9	10	11	12
0 – 30%	13	14	15	16

Defoliation State Structure

Years Since Defoliation	Defoliation Level: Low	Med	High
5+	1		
2 – 4	2	3	4
1	5	6	7

- Combined, there are 16 x 7 = 112 possible discrete states that a unit can be in at any one time

Decision Alternatives

- Menu of management action alternatives
    - Rest
    - Graze
    - Burn
    - Burn / Graze
- Each management year a manager chooses one of these actions to apply to a unit

Degree of Management Control

- Action carried out is not always the intended action
  - Unfavorable conditions, lack of resources, access to burn crew or grazing contractors, etc., influence choice of action and make it impossible to fully control what is implemented
- ***Partial controllability***: An irreducible form of uncertainty that we must explicitly take into account in the decision framework
- Because it can affect which management action is optimal, we must formally recognize the probability of implementing each management action

Partial Controllability

No Previous Defoliation

Recommended Action	Implemented Action			
	Rest	Graze	Burn	Burn/Graze
Rest	0.92	0.06	0.01	0.01
Graze	0.10	0.81	0.06	0.04
Burn	0.15	0.18	0.61	0.06
Burn/Graze	0.12	0.25	0.13	0.49

Previous Defoliation - 75% Reduction of Burn and Burn/Graze

Recommended Action	Implemented Action			
	Rest	Graze	Burn	Burn/Graze
Rest	0.92	0.07	0.00	0.00
Graze	0.13	0.85	0.01	0.01
Burn	0.38	0.45	0.15	0.01
Burn/Graze	0.29	0.56	0.03	0.12

Decision Utility

- Describes what the manager wants from the system through management

  ** Increase cover of native prairie while minimizing cost **

- Combines both aspects of the objective
  - Expresses the balance between the ***value*** of having native prairie with the ***cost*** of achieving it

- Subjective expression of values – separate from beliefs about the behavior of the system (model set)

Decision Utility

Native Cover *after* Mgmt Action

Native Cover *before* Mgmt Action		60 - 100%	45 - 60%	30 - 45%	0 - 30%
	60 - 100%	0.92	0.61	0.35	0.18
	45 - 60%	0.95	0.75	0.51	0.29
	30 - 45%	0.96	0.79	0.59	0.37
	0 - 30%	0.97	0.82	0.64	0.44

Native Cover *after* Mgmt Action

Native Cover *before* Mgmt Action		60 - 100%	45 - 60%	30 - 45%	0 - 30%
	60 - 100%	0.92	0.61	0.35	0.18
	45 - 60%	0.95	0.75	0.51	0.29
	30 - 45%	0.96	0.79	0.59	0.37
	0 - 30%	0.97	0.82	0.64	0.44

Value of remaining in a state

Native Cover *after* Mgmt Action (columns); Native Cover *before* Mgmt Action (rows)

	60 - 100%	45 - 60%	30 - 45%	0 - 30%
60 - 100%	0.92	0.61	0.35	0.18
45 - 60%	0.95	0.75	0.51	0.29
30 - 45%	0.96	0.79	0.59	0.37
0 - 30%	0.97	0.82	0.64	0.44

Value of improvement in state

Native Cover *after* Mgmt Action (columns); Native Cover *before* Mgmt Action (rows)

Value of degradation in state

	60 - 100%	45 - 60%	30 - 45%	0 - 30%
60 - 100%	0.92	0.61	0.35	0.18
45 - 60%	0.95	0.75	0.51	0.29
30 - 45%	0.96	0.79	0.59	0.37
0 - 30%	0.97	0.82	0.64	0.44

Rest

	60 - 100%	45 - 60%	30 - 45%	0 - 30%
60 - 100%	0.92	0.61	0.35	0.18
45 - 60%	0.95	0.75	0.51	0.29
30 - 45%	0.96	0.79	0.59	0.37
0 - 30%	0.97	0.82	0.64	0.44

Graze

	60 - 100%	45 - 60%	30 - 45%	0 - 30%
60 - 100%	0.87	0.55	0.29	0.13
45 - 60%	0.90	0.69	0.43	0.21
30 - 45%	0.91	0.72	0.52	0.28
0 - 30%	0.92	0.75	0.56	0.36

Cost

Burn

	60 - 100%	45 - 60%	30 - 45%	0 - 30%
60 - 100%	0.83	0.50	0.25	0.10
45 - 60%	0.86	0.64	0.37	0.16
30 - 45%	0.88	0.67	0.46	0.22
0 - 30%	0.89	0.70	0.50	0.29

Burn / Graze

	60 - 100%	45 - 60%	30 - 45%	0 - 30%
60 - 100%	0.81	0.47	0.22	0.07
45 - 60%	0.83	0.61	0.33	0.12
30 - 45%	0.85	0.64	0.42	0.17
0 - 30%	0.87	0.67	0.46	0.25

Decision Utility – Recap

- Decision utility is a subjective expression that describes how managers feel about the system response to the actions they apply
- Utility value is dependent on:
  - Where you came from → Native cover level at time *t*
  - Where you end up → Native cover level at time *t*+1
  - How you got there → Management action taken
- Utilities were elicited from cooperators

Decision Utility & Objective Function

- Utility provides an annual measure of what manager receives from the system for what he/she invests
- Objective stated in mathematical form →Objective Function

$$\max_{a_t \in A} \sum_{t=0}^{T} E(u_{t+1} \mid x_t, x_{t+1}, a_{t+1})$$

  - The objective function is a long-term sum of annual utilities

$$U_T = u_{t+1} + u_{t+2} + u_{t+3} + \ldots + u_T$$

- The endpoint, T, is sufficiently large to represent an indefinite time horizon → the goal is sustained conservation of native prairie

Decision Utility & Objective Function

- Formulated utility to focus on improving and maintaining native cover with each time step
  - Seeking a recurrent reward each year
- Formulated the objective function to maximize cumulative utilities over the long-term
  - Seeking to sustain native cover by accruing recurrent rewards into the indefinite future
- Formulating the objective function in terms of a long-term accumulated sum prevents myopia in decision making
  - Tomorrow's decision always must be taken into account

Utility, Objective Function, & Discounting

- ***Discounting***: Returns in the future have less value relative to the same return today
- When the aim is long-term conservation of a resource, one might chose not to discount
    - For practical reasons of computational stability we incorporated discounting
- With discounting, the objective function is a long-term sum of annual ***discounted*** utilities

$$U_T = u_{t+1} \cdot d^0 + u_{t+2} \cdot d^1 + u_{t+3} \cdot d^2 + \ldots + u_T \cdot d^{T-1}$$

- Where '*d*' is the discount factor, exponentiated by time

Utility, Objective Function, & Discounting

- With no discounting, $d = 1$
- The larger the discount factor (i.e., the closer it is to 1), the less the future is discounted
- We chose a discount factor of 0.993
    - Utility gained 100 years in the future is discounted by half of what you would receive for the same event had it occurred today ($0.993^0 = 1$, whereas $0.993^{100} = 0.5$)
- Choice of '*d*' is sufficiently large so as to not overly undervalue the future, which is important when the aim is long-term conservation of the resource.

Decision Utility and Objective Function

- Objective function – the quantitative description of the management objective – is used in the optimization procedure to find the sequence of decisions through time that maximizes cumulative utility
- Because utility directly reflects the objective, maximizing cumulative utility identifies the sequence of decisions through time that will achieve the objective

Optimization: Finding the Best Management Action

- Adaptive stochastic dynamic programming
    - Method of optimization – integrates models and utility
    - Identifies sequence of decisions through time that maximizes cumulative utility and achieves the management objective
    - Accounts for:
        - Future dynamics of system state and knowledge state
        - Current and future expected returns (utility)
        - Degree of management control (partial controllability)

<u>Optimization – Decision Policies</u>

- Produces an optimal decision table that identifies the best decision for every combination of:
  - System State x Knowledge State

vegetation states (16) x
defoliation states (7)
= 112 system states

	M1	M2	M3	M4	sum
weight					1.0

- Model weights
  - Individual model weights are between 0 and 1, discretized by 0.08333
  - Four weights sum to 1.0
- 455 possible model set weight combinations

system states (112) x knowledge states (455) = 50,960 possible states

Veg state	Defol state	Model Weights M1	M2	M3	M4	Optimal Mgmt
1	1	0.25	0.25	0.25	0.25	Burn

1 of the 50,960 rows in the full decision table

- One decision policy per knowledge state
  - A complete policy includes the optimal management decision for each of the 112 system states

<u>Optimization – Decision Policies</u>

- Decision policies vary depending on the weights on each competing model
- Weights create differential influence of each model on the decision policy
  - Models with greater weight have greater influence
  - Models with lesser weight have lesser influence
- The policy that is used each year depends on the current knowledge state
  - Knowledge state is updated each year as go through the annual decision making and monitoring cycle

Recall Case Study
Module C -
Monitoring & Learning

<u>Decision Policy – Anatomy</u>

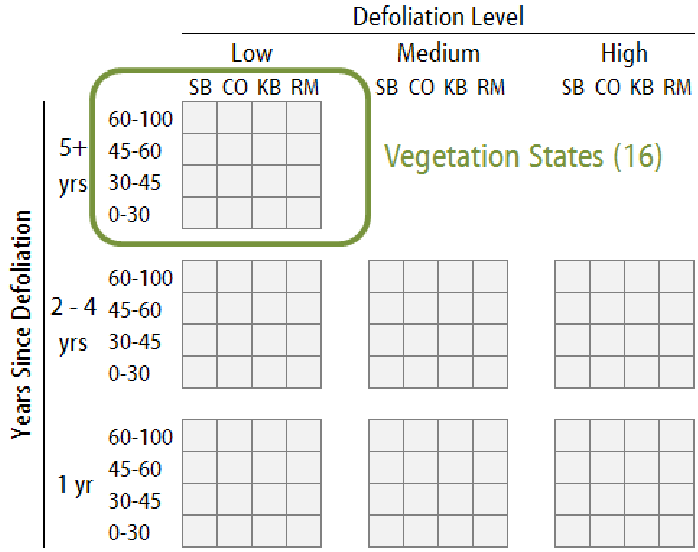

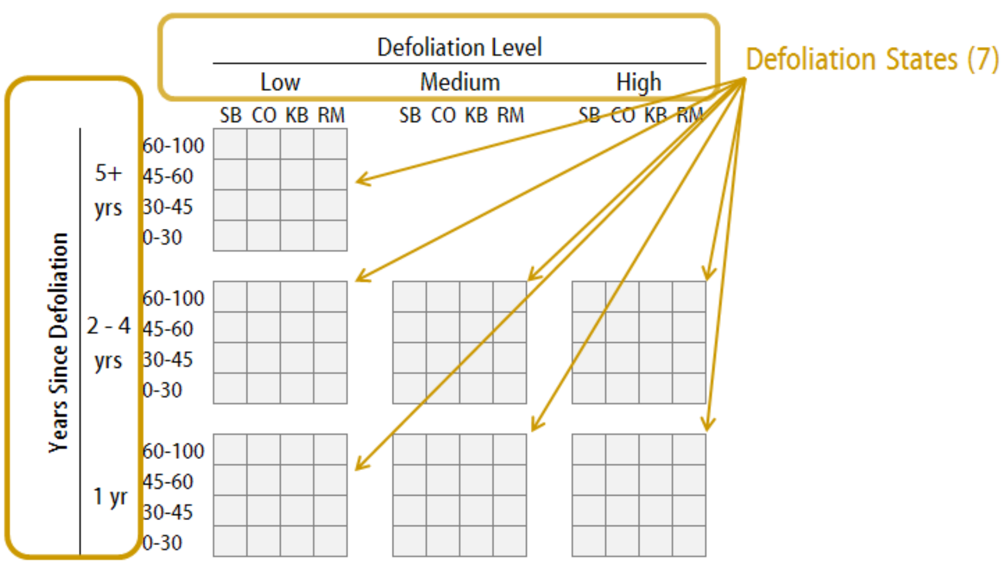

Vegetation States (16) x Defoliation States (7) = 112 states

Decision Policies – Examples

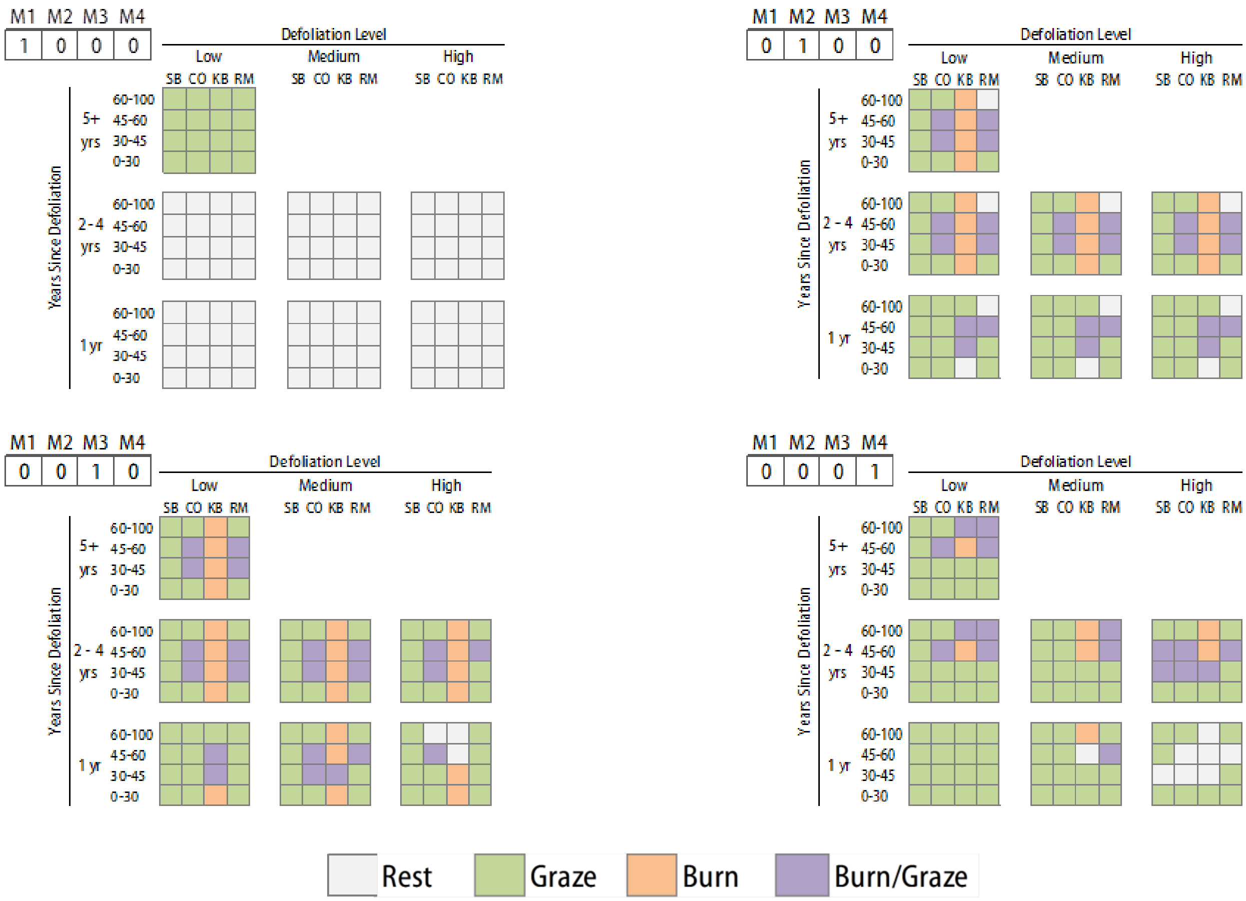

Knowledge State

	Model 1	Model 2	Model 3	Model 4
Weight	0.25	0.25	0.25	0.25

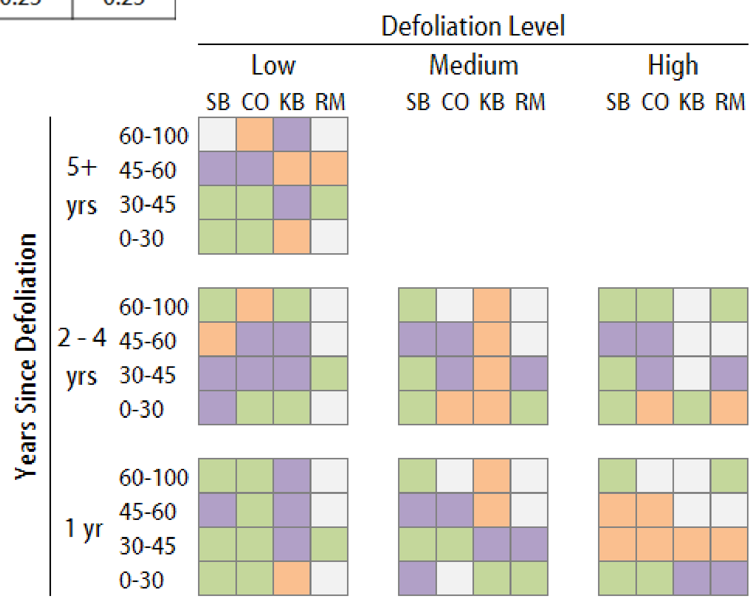

Decision Policies – Recap

- Different knowledge states result in different policies
- At any point in time, the current policy depends on the current knowledge state, and is weighted by the current belief in each competing model
- The knowledge state, and thus decision policy, may change each year as we learn and update our knowledge, or our confidence in each competing model

Summary – Dynamic Decision Making

- Nature of the problem is that an annual decision must be made, but under uncertainty
  - Components of the decision are in place – Decision Alternatives & Objective Function
- The action carried out is not always the intended action (partial controllability)
  - We must formally take this into account because it affects the management decision that is optimal to recommend
- Utility is an annual measure of what is received from the system (native cover) for what is invested (cost of management action)
- Expressing the objective function as a long-term accumulated sum of recurrent rewards prevents myopic decision making
- Optimization
  - Integrates the competing models and utility
  - Accounts for future dynamics of system state and knowledge state, current and future expected returns, and degree of management control
  - Identifies sequence of decisions through time that maximizes cumulative utility and achieves the management objective
  - Provides a decision table with individual decision policies per knowledge state
- The decision policy for each year is pulled from the full decision table and depends on the current understanding of system behavior (as represented by the confidence in each competing model), or our knowledge state

Literature Cited

Gannon, J.J., T.L. Shaffer, C.T. Moore. 2013. Native Prairie Adaptive Management: A Multi Region Adaptive Approach to Invasive Plant Management on Fish and Wildlife Service Owned Native Prairies: U.S. Geological Survey Open File Report 2013-1279, 184 p. with appendixes, http://dx.doi.org/10.3133/ofr20131279

Next Up…
Case Study Module E – Seeing the Whole Picture