Big Bend National Park, Texas

Chapter 4: Uncertainty and Learning

In this chapter we discuss some important technical issues that arise in the implementation of adaptive management, in particular the treatment of uncertainty in resource management and the influence of long-term (and uncertain) environmental trends. We also address attributes of models and management alternatives that promote learning.

4.1. Components of uncertainty

Here we revisit the components of uncertainty that can affect natural resources in the context of thematic areas explored in this guide. We focus on the uncertainty factors highlighted in Chapter 2, including environmental variation, partial controllability, partial observability, and structural uncertainty. These uncertainties influence natural resources management in different ways and at different points in a resource system (Figure 4.1). Taken separately or in combination, they can limit understanding of resource functions and restrict our ability to identify useful management strategies. The difficulties they introduce vary with the particular ecological situation, but as a rule their potential impacts increase with the scale and complexity of the resource system.

Environmental variation. Environmental conditions can be viewed as external factors that influence, but are not influenced by, resource conditions and dynamics. Here we consider environmental conditions in terms of the physical environment, as expressed in precipitation patterns, temperature regimes, ambient light conditions and other measures, as well as extremes in these conditions. Environmental conditions directly and indirectly influence the ecological and physical processes that determine resource dynamics. Because they vary randomly over time, future conditions cannot be predicted with certainty.

Environmental fluctuations may be thought of as lacking a discernable pattern of change in central tendency or range of variation. Alternatively, they may be seen in terms of directional trends, such as a long-term decrease in average precipitation or an increase in the range of ambient temperatures. The latter framework is especially relevant to climate change, which is characterized in terms of directional environmental change over an extended period.

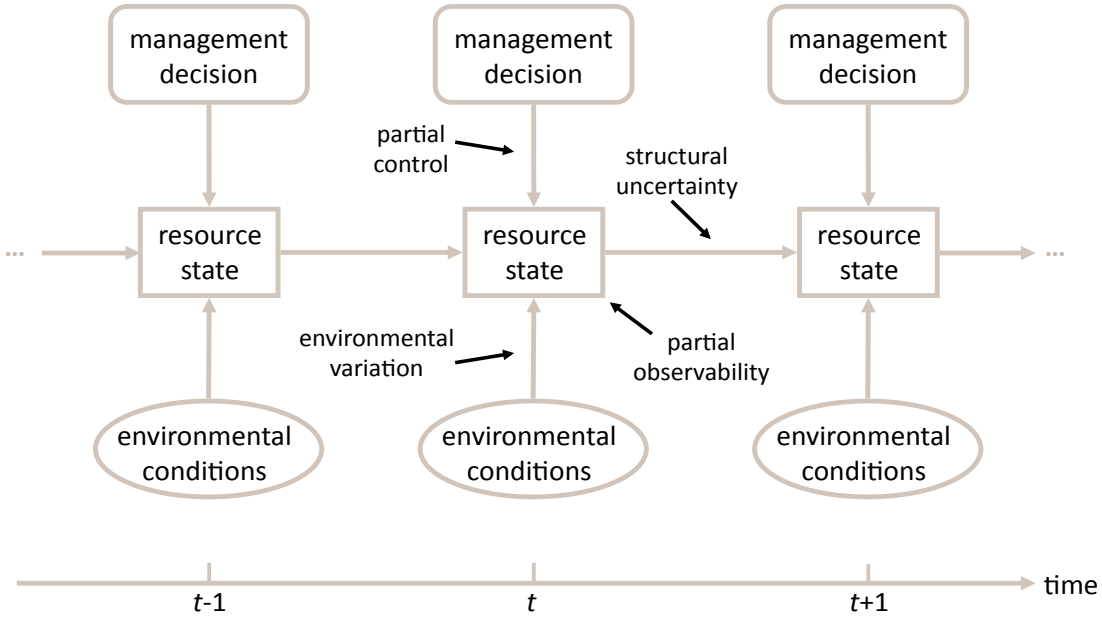

Figure 4.1. Uncertainty factors in natural resource management. Partial control limits the influence of management actions. Environmental variation affects resource system status and dynamics. Partial observability limits the recognition of system status. Structural uncertainty limits the ability to characterize system change.

Fluctuations in the environment can interact with land-use and land-cover changes that occur during the same time that the climate is changing. Urbanization, deforestation, industrial agriculture, manufacturing, mining, transportation, and other activities have increased worldwide, with potentially profound impacts on resource systems. Because climate change and human development have occurred simultaneously, their impacts are difficult to separate. However, there is little doubt that in combination they are altering natural resource systems and causing long-term changes in resource dynamics.

It often is useful to include unrecognized landscape heterogeneity and unpredictable human impacts on the landscape as a part of "environmental variation." In combination these factors can greatly influence resource responses to management, depending on the scale and unpredictability of the change. For example, management strategies needed for irregular, large-amplitude environmental fluctuations may differ from those needed for more predictable fluctuations of smaller amplitude. Though environmental variation is assumed to be uncontrolled, it can be tracked through monitoring and incorporated into forward-looking management strategies.

Partial controllability. Partial controllability refers to the difference between the results intended by a given management decision and the results that actually occur. Stated formally, it describes a random association between intended and realized management actions. Unintended outcomes are often a result of management decisions implemented by indirect means. For example, hunting permits may be used as an indirect means to attain a chosen rate of waterfowl harvest, as in our example of adaptive harvest management (Section 4.4); or forestry regulations may be used to limit logging-related impacts on wildlife. The net effect is that the intended outcome of a management decision is only partially accomplished by the action taken. One way to account for this is to characterize an anticipated action probabilistically, with a distribution that assigns probabilities of occurrence over a range of potential outcomes.

A somewhat different version of partial controllability can arise if there is a delay between identifying an action and implementing it. In this case partial controllability is induced not by an imprecise or indirect linkage to a control mechanism, but rather by unforeseen circumstances that restrict or prevent the implementation of the action. One example is an unanticipated loss of funds for a management intervention. In such a case there is a point between the identification and implementation of an action when the manager recognizes that the chosen action cannot be carried out.

In actual operations, partial controllability differs from environmental variation in terms of the nature and timing of its effect. Thus, partial controllability occurs at specific points in the resource system where management alters resource conditions and states, with decisions and actions linked at each point in time (Figure 4.1). On the other hand, environmental variation is expressed through fluctuations and trends in environmental conditions over time. Fluctuating environmental conditions influence ecological processes in ways that are uncontrolled, uncertain, and often unrecognized.

Notwithstanding these differences, environmental variation and partial controllability are sometimes combined in models of resource dynamics, mainly because of similarities in the way they are characterized. Like environmental variation, partial controllability

increases with geographic scale and ecological complexity: the larger and more complex the resource system, the less certain we can be that management decisions will have the intended outcome. For example, regulations for hunting ungulates may not result in the intended harvest rates if the animals occur in wide-ranging groups (perhaps based on age or sex) with different likelihoods of being seen by hunters.

Partial controllability is likely to cause significant uncertainty in managing projects in the thematic areas emphasized in this guide. For example, indirect mechanisms like tax incentives, permit systems, and carbon trading arrangements may be used to prevent atmospheric carbon dioxide from rising above a certain level, but the net effects of these emissions controls may differ from the outcome that is intended. Climate change adaptation also may involve indirect control mechanisms, such as land transfers, outreach and communication efforts, and regulatory mechanisms. Similarly, partial controllability is likely to be an important uncertainty factor in water management, energy development, and large-scale human activities on the landscape, each of which involves many of the same kinds of management controls that can be used for climate change. The importance of partial controllability varies with scale, and it tends to be less important in localized, smaller-scale projects for which random variation is limited and control can be exercised more directly.

Partial observability. Partial observability expresses our inability to observe completely the resource system that is being managed. Natural resources are almost always partially observed. For example, only a part of the area where a fish population occurs can be monitored, and a sampling strategy needs to allow inferences over the whole area on the basis of the observation of only a part of it. Observability is further complicated by the fact that individuals (e.g., plants and animals) often escape detection, even in areas that are intensively monitored. In combination, incomplete geographic coverage and incomplete detectability mean that observations collected in the field are associated with – but not the same as – actual system states.

Partial observability obscures the resource status on which effective management depends. This reduces management effectiveness, even if environmental variation is minimal and management actions are precisely controlled. For example, decision makers without accu-

rate information can fail to recognize the need to protect a resource, or overlook opportunities for sustainable resource exploitation (Moore and Kendall 2004). These problems become more pronounced under highly variable environmental conditions.

Partial observability is commonly measured by sampling variation, which occurs when field data are collected and analyzed. Unlike environmental variation, over which we have little if any control, the accuracy with which resources are observed can be controlled by designing field sampling efforts efficiently. For example, we can reduce uncertainty about resource status with more intensive sampling, optimal geographic design of the sampling effort, and the use of standard survey principles like randomization, replication, and controls. Nonetheless, partial observability can rarely be eliminated, no matter what the design and sampling intensity.

There are several ways of dealing with partial observability in decision making. One is to estimate resource status with field data, and then treat the estimate as if it accurately represents resource conditions. Another is to state the uncertainty about resource status explicitly, with probabilities for possible resource states, and incorporate these probabilities directly into the decision-making process (Williams 2009). The first approach is far more common in natural resource management. Of course, the most direct way to address partial observability is to reduce it as much as is practicable with well-designed monitoring.

Like the other forms of uncertainty, partial observability increases with geographic scale and ecological complexity. For example, wildlife population abundance is more difficult to estimate if populations consist of widely dispersed age or size groups that are not equally detectable. As a general rule, the larger and more complex the resource system, the less certain we can be that the resource estimates on which management is based track the actual system state.

Structural uncertainty. Structural uncertainty is a result of a lack of understanding (or lack of agreement) about the processes that control resource dynamics. In virtually all cases there is some degree of uncertainty about the forms and functions – i.e., the structure – of natural processes. Structural uncertainty can limit our ability to manage resources effectively and efficiently, even if monitoring is exact, management actions are rigorously controlled, and environmental variation is minimal.

The differing views held by stakeholders about how natural processes work and how they respond to management are examples of structural uncertainty. These views can be framed as hypotheses about system processes and responses and then embedded in models, which in turn can be used to make testable predictions. Examples of uncertainty about resource form and function include hypothesized associations between different attributes of the resource, or relationships between controls and resource elements, or connections between environmental conditions and resource states, or parameterizations of these relationships. The hypothesized forms and parameterizations can be incorporated in different models, and structural uncertainty then is expressed in terms of uncertainty about which model (and its embedded hypothesis) best represents resource dynamics.

In adaptive decision making, structural uncertainty changes over time because it is based on evolving resource conditions and management actions. These changes are quantified with measures of confidence in the ability of the models to predict resource dynamics. A common mathematical approach is Bayesian updating,

which combines confidence values and monitoring data at each point in time to generate new confidence values for the next point in time (Lee 1989). Confidence increases in models that make accurate forecasts of resource conditions, and confidence declines in models that do not make accurate forecasts. Of course, changes in confidence differ from a change in the hypotheses themselves, which occurs through the process of double-loop learning (see Section 2.7).

Structural uncertainty, like the other forms of uncertainty, has a tendency to obscure the effects of management and reduce effectiveness. However, it differs from environmental variation and partial controllability in its point of influence (Figure 4.1) and the manner in which it is treated. Structural uncertainty can be reduced by applying management strategies to affect the measures of confidence in models. In contrast, environmental variation (and in some cases partial controllability) are effectively uncontrolled.

4.2. Systemic resource changes over time

Adaptive management is usually framed in terms of an (often unstated) assumption that the features and processes of a resource system are stable over the management time frame, so that uncontrolled fluctuations change little in overall direction or range of variation. A generic model for adaptive management assumes that at any given time, resource change is influenced by the state of the resource, environmental conditions, and the management action taken at that time (Figure 2.1). Randomness in environmental conditions induces random resource changes, and directionality in these conditions over time means that uncontrolled resource dynamics also tend to exhibit directionality over time. Conversely, random and non-directional environmental fluctuations tend to preserve dynamic stationarity in resource behaviors. Approaches to system analysis and control, including the framework typically used in adaptive decision making, have traditionally rested on the assumption that system features and patterns of fluctuation are stable over time.

It is becoming increasingly clear that for a great many resource systems, the ecological structures and processes controlling resource dynamics are changing in ways not fully expressed by the management framework depicted in Figure 2.1. Of particular importance is that environmental conditions, and the ecological processes influenced by them, are exhibiting directional patterns of change. An obvious example is climate change, in which the environment is seen as evolving directionally in terms of temperature, precipitation, and other variables.

An important challenge for an adaptive approach is to include directional trends in the environment. Such an extended framework is especially relevant to climate change, as expressed in terms of directional environmental change like a long-term decrease in average precipitation or an increase in the range of ambient temperatures. Of course, directional change can be important over shorter periods as well; many anthropogenic forces exhibit large-scale directional change on shorter time scales than climate change. In either case, directional change has the potential to induce directionality in resource behaviors, i.e., to generate non-stationary resource dynamics.

Non-stationary dynamics become especially challenging for a forward-looking, learning-based approach like adaptive management. Learning about resource processes and management impacts proceeds through an iterative process of decision making, follow-up monitoring, and assessment of impacts. The cycle of learning becomes more difficult when the subjects of investigation – the ecological processes that determine resource change – are themselves evolving.

One way to address this problem is to track and even model the environmental drivers of change, and to use trends in environmental conditions to account for changes in patterns of resource change over time.

Another way is to look for limited periods during which resource processes are largely stable so that learning-based management can be effective. A third approach is to develop environmental scenarios with different patterns of directional change, and try to design acceptable management strategies that account for uncertainties among the scenarios. Adaptive decision making then can be used to address uncertainty about which scenario is appropriate (and therefore which strategy should be chosen).

Non-stationarity is a newly recognized and serious challenge to adaptive decision making, one for which we need new approaches that go beyond the standard ways of framing and conducting learning-based management. At a minimum it is necessary to look for directional trends in environmental conditions and systematic changes in resource structures and functions, and consider ways to accommodate them.

Finally, it is worth re-emphasizing that systemic change in resource dynamics can also be caused by large-scale, effectively permanent human interventions on the landscape. At a certain scale the human footprint on the landscape can be thought of as part of the external environment, and long-term growth of that footprint can easily cause changes in physical and ecological processes. Because it is the result of human actions, the footprint presumably is partially controllable. However, long-term changes, driven in large part by the growth of human populations, economic growth, technological change, and demands for natural resources and space, are unlikely to stabilize in the foreseeable future. Like the directional change in environmental conditions caused by climate change, long-term patterns of increasing resource use and disturbance, and the directional trends they cause in resource dynamics, will need to be taken into account in adaptive decision making.

Models play a key role in adaptive management by incorporating different hypotheses about how a resource system works and how it responds to management. Agreements, disagreements, and uncertainties about resource behaviors can be highlighted with models and used to guide investigations through basic research and learning-oriented management interventions.

4.3. Models, management alternatives, and learning

In an adaptive management application, both models and management alternatives are identified and agreed upon by managers and other stakeholders. The models, which embed different hypotheses about how the resource system works, represent uncertainty (or disagreement) about ecological processes and the influence of management on them. The management alternatives express a range of potential actions that can be taken at each decision point.

Models play a key role in adaptive management by incorporating different hypotheses about how a resource system works and how it responds to management. Agreements, disagreements, and uncertainties about resource behaviors can be highlighted with models and used to guide investigations through basic research and learning-oriented management interventions.

The management alternatives also play a key role. The identification of informative and effective strategies depends upon differences in predicted responses to management actions. For optimal management, distinctly different predictions should be produced for the alternative actions, so as to facilitate the identification of an optimal action. To promote learning, distinctly different predictions should be produced by the alternative models.

These conditions suggest two ways that adaptive decision making can be compromised. One way is that the models representing uncertainty about system structure produce similar predictions for the alternative actions. Under these circumstances there is little practical value in resolving uncertainty about how the system works, since the models describing system performance all perform equally well. The other way adaptive decision making can be compromised is that the available actions produce essentially indistinguishable results for each model. In this situation there is little value in attempting to discriminate among management alternatives because they all produce similar results. The latter situation often occurs when the potential management alternatives differ only marginally.

This implies that the models and management alternatives should be considered in combination. For a given set of actions, the various models under consideration should predict distinctly different outcomes, so that learning through management becomes possible. Similarly, the various alternatives should produce distinctly different predictions, so that the best action can be clearly seen. Adaptive decision making works best when (*i*) there is substantial variation in the hypothesized forms and functions for the resource system, and (*ii*) management alternatives differ substantively in their predicted resource responses.

A special case of adaptive decision making treats the management alternatives themselves as hypotheses (Williams 2011*a*). Each alternative is seen as a hypothesis about the effectiveness of the action, much as hypotheses in experiments are expressed in terms of responses to experimental treatments (Graybill 1976). The emphasis here is restricted to system responses to management, rather than improved understanding of the ecological processes behind those responses.

As an example, consider the alternatives of clear cutting, thinning, and selective cutting as hypotheses about the best way to manage a forest stand. A choice of one of the alternatives sets up an "experiment," which provides evidence that either does or does not support the intervention as an effective management action. If the forest's response contributes to meeting the management objectives, the intervention is a viable candidate for continued use. A response differing from what was expected or desired suggests that the intervention should be rejected in favor of another. The problem, of course, is that there is always uncertainty about system responses to management interventions, and predictions about the responses must somehow account for those uncertainties. Without a mechanism for learning based on the comparison of alternative predictions against observed evidence, this "experimental" approach can easily become a form of trial and error management.

There are at least two ways to strengthen the inferences of such "experimentation." A traditional way is to use randomization, replication, and controls, when possible, in the spirit of experimental design. Thus, we might use simultaneous interventions on different management units in different places. This makes it possible to compare the effect of one intervention on a group of management units against a different intervention on other units. Our example of post-wildfire management after the Biscuit Fire in Oregon describes such a management study by the Forest Service. Standard statistical treatments can be used for the comparison, with results that can contribute to improved management.

If the interventions are carried out sequentially, one can compare monitoring data against predictions for each alternative to update confidence in the alternatives over

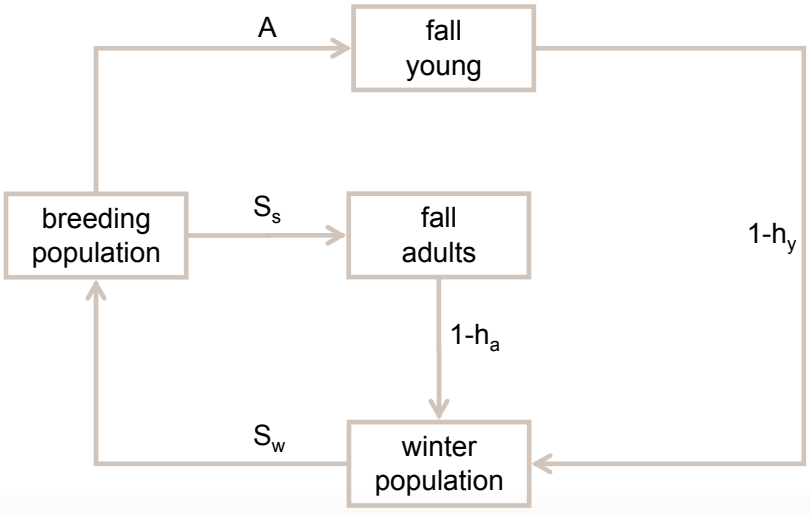

Figure 4.2. Conceptual model of annual cycle of mallard population dynamics. Model includes survival rates for spring-summer (S_s) and fall-winter (S_w), along with harvest rates for young (h_y) and adults (h_a) and age ratio (A) for reproduction/recruitment.

time. In this case, one must describe a predicted response to a given intervention under each of the alternative models. For example, hypotheses (models) can be formulated based on the relative responses to clear cutting, selective cutting, and thinning. One common method is to identify confidence values for the intervention models and update these values at each decision point by comparing predicted responses with post-decision monitoring data (Williams et al. 2002). In this way the confidence values can evolve over time, increasing for alternatives that are supported by the data and decreasing for alternatives that are not. A change in the confidence values then becomes a measure of learning over time, leading gradually to recognition of the best intervention.

Learning through experimentation typically involves the use of classical hypothesis testing, in which interventions are considered experimental "treatments" and analysis-of-variance methods are used to recognize statistically significant treatment effects. When interventions are implemented sequentially, a popular alternative for learning is to update the credibility of different hypotheses over time on the basis of post-decision monitoring data.

Adaptive decision making works best when (i) there is substantial variation in the hypothesized forms and functions for the resource system, and (ii) management alternatives differ substantively in their predicted resource responses.

4.4. Example: Uncertainty and learning in waterfowl management

An example that highlights many of the points in this chapter is the framework for adaptive harvest management of waterfowl. Adaptive harvest management was begun in 1995 as a process for setting annual regulations for the sport hunting of waterfowl in North America (Williams and Johnson 1995, Williams et al. 2002). It uses a simple model to represent associations among fall harvest, seasonal survivorship, and spring reproduction (Figure 4.2). Contrasting hypotheses about the impact of harvest on annual survivorship are easily incorporated into different versions of the model by describing different functional relations between harvest rates and post-harvest survival. In addition, contrasting hypotheses about the importance of density dependence in recruitment are incorporated by describing recruitment in terms of spring population size. In combination, these hypotheses define different models, each with its own predictions about harvest impacts and each with its own measure of confidence that evolves over time. The models and their measures of confidence characterize structural uncertainty, which is reduced as harvest actions are taken and post-harvest monitoring data are used to update the confidence measures. Learning is expressed through the updating of these measures and is folded into the annual process of setting hunting regulations.

The forms of uncertainty we have described in this chapter enter naturally into this problem. For example, harvest rates that are targeted through the use of regulations result in partial controllability. Environmental variation affects recruitment through water conditions on the breeding grounds, as measured by the abundance of ponds. The change in pond numbers each year is based on the number in the current year and the amount of precipitation the next winter and following spring. Precipitation amounts are assumed to be random and independent from year to year, with no long-term trend in the average amount or severity of precipitation events. Finally, one of the most comprehensive monitoring programs for wildlife in the world (Martin et al. 1979, Smith et al. 1982) is used to estimate the status of waterfowl populations and the parameters that control waterfowl population dynamics.

The assumption of dynamic stability underlies the approach currently used to identify optimal harvest regulations in the presence of the various sources of uncertainty. Thus, harvest strategies are assessed in the context of a dynamic but stable resource system. It is straightforward to incorporate non-stationarity in the waterfowl harvest problem simply by including directionality in the amount of precipitation over time. Long-term directionality in annual precipitation induces systemic changes in the average pond conditions, which in turn induce long-term patterns of change in waterfowl populations and harvests. Under these circumstances the structures and processes of the resource system change through time, even in the absence of harvest. These changes should be taken into account as we evaluate forward-looking harvest strategies (Nichols et al. 2010).

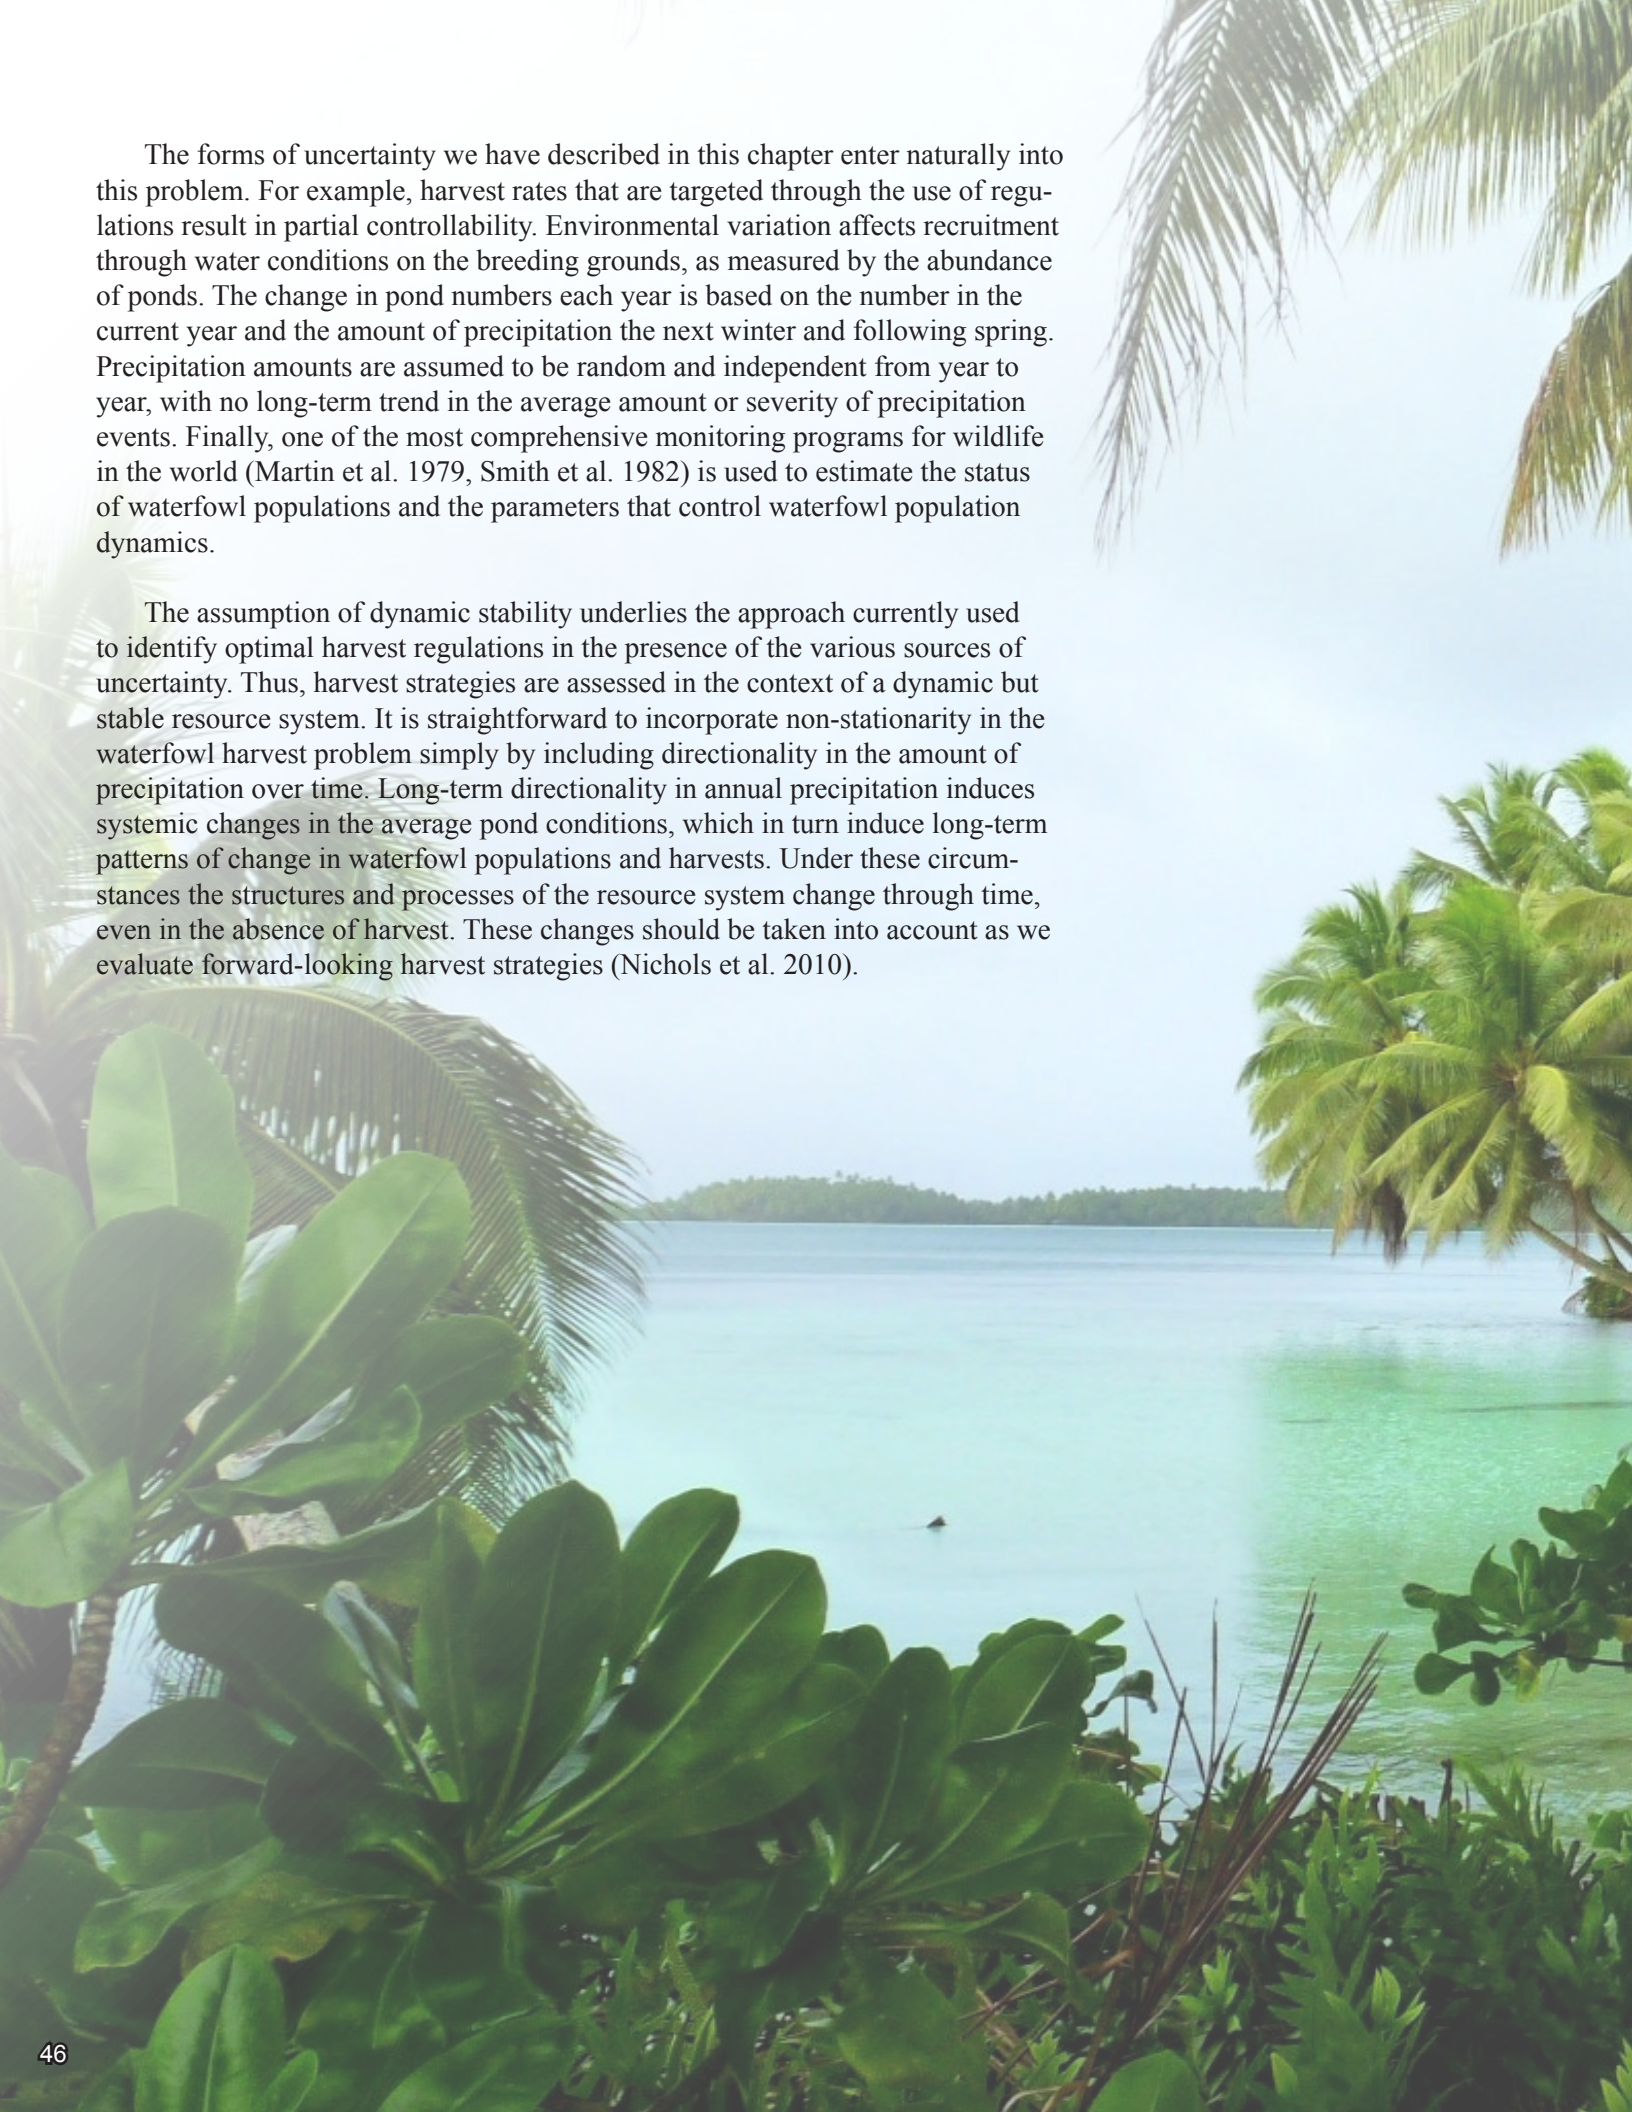

Palmyra Atoll National Wildlife Refuge, Pacific Ocean