CHAPTER 8

Stylized Models to Analyze Robustness of Irrigation Systems

Marco A. Janssen and John M. Anderies

Irrigation systems are a prevalent feature in human history. They are archetypal examples of human efforts to control spatial and temporal environmental heterogeneity. In this chapter, we report on some preliminary results to understand the emergence, robustness, and collapse of irrigation systems, using a suite of stylized models. As a starting point, we perceive irrigation systems as similar to control systems in engineering. The "irrigation society" may usefully be thought of as a controller engineered to meet specific performance criteria defined by sustained, relatively constant food output, given both spatially and temporally variable water input. That is, the well-functioning irrigation society control system will get the right amount of water to the right place at the right time.

This, of course, is a nontrivial task. Flying a large jet on autopilot is a nontrivial task as well, yet modern feedback-control systems can do this quite well. Control engineers have devised highly effective control systems that perform well under a variety of conditions. Key issues faced by feedback control designers are noisy feedback from the system (either because the system has some random components or because the system's state cannot be perfectly known) and uncertainty about how the system actually works (model uncertainty). These conditions would certainly be characteristic of an irrigation system. How does one design a feedback (closed-loop) control system under such circumstances? This question is the subject of the field of robust control. It is in

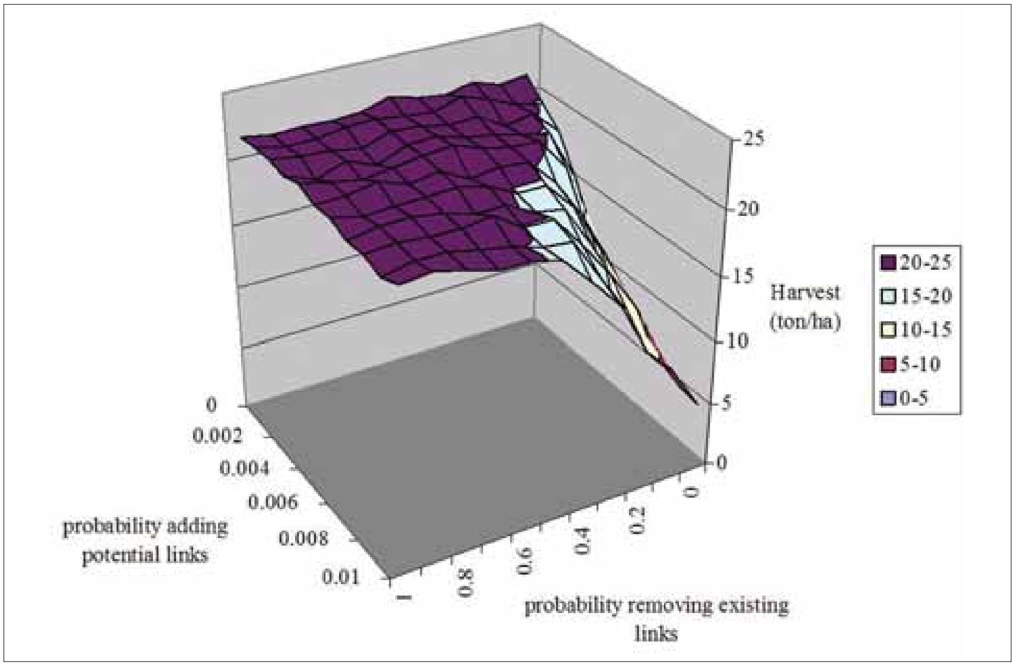

Figure 8.3 Average harvest per subak (per hectare) for various degrees of perturbation of links among subaks that disperse pests and are conditional imitators.

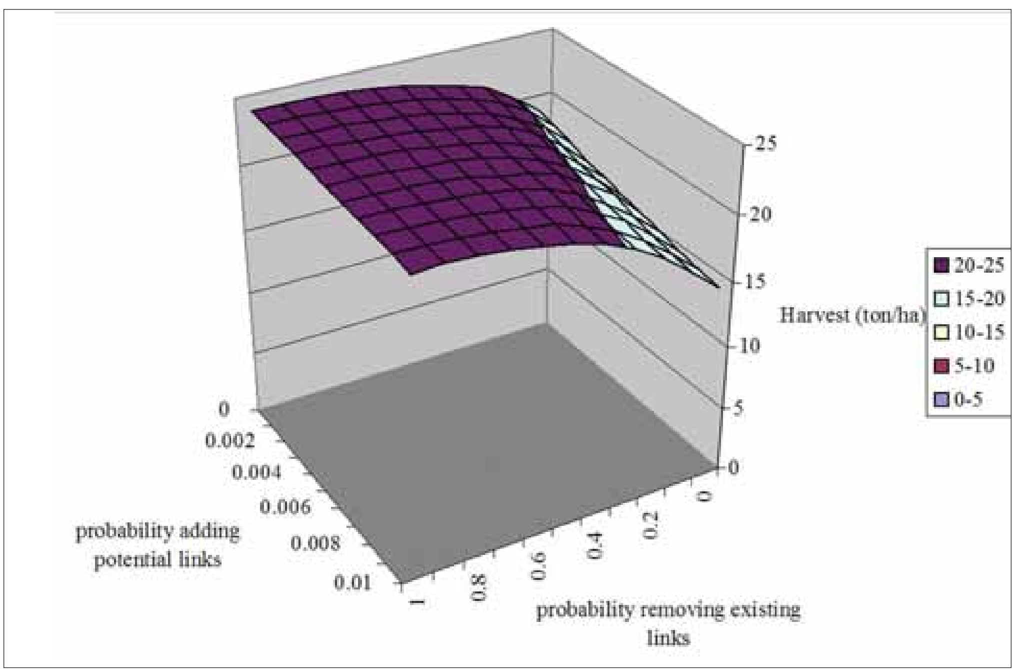

Figure 8.4. Average harvest per subak (per ha) for various degrees of perturbation of links among subaks that disperse pests and are conditionally adaptive.

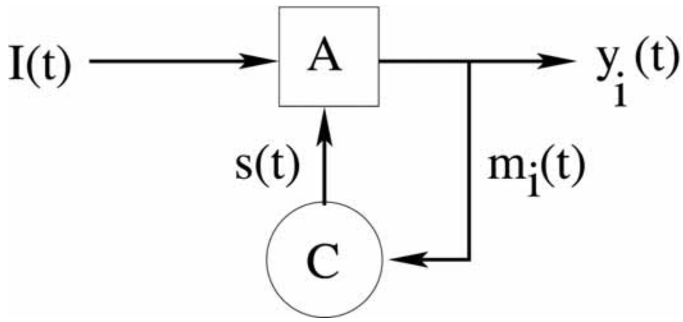

Figure 8.1 Schematic of a simple control system. Given outputs $y_i(t)$ *at each location,* i*, and time,* t*, the controller, C, receives measurements,* $m_i(t)$*, and generates a signal,* s(t)*, which it sends to the actuator, A, to adjust flows, given* I(t).

this context that we are interested in the robustness of irrigation systems.

Although there is a nice analogy between irrigation societies and engineered control systems, they are different in several respects. These differences are what make the former so interesting.[1] To illustrate, consider the following task: given a known distribution of incoming water volume and an area of productive land, design the infrastructure to deliver water in a way that maximizes some performance index, most likely involving the amount of food produced, conditioned on some aspect of the cost of producing it. In the reasonably controlled environment characteristic of most engineering problems, this is a straightforward task. The infrastructure would comprise a series of pipes (perhaps open, as in canals), a series of valves (or head gates), and a set of flow monitors that are at key delivery points and feed back to the valves to control flow. The design problem is to produce the infrastructure that accomplishes the objective with the lowest capital cost (fewest values, minimal pipe size and length, and so forth) and lowest operation costs (least friction loss, simplest configuration, and so forth) and is not "too" sensitive to noise in measurement or in the water input data.

Two important issues for irrigation systems are not mentioned in the preceding paragraph. The first is the matter of scale. The description above assumes that the system to be controlled is small enough in scale that a single entity coordinates the signals arriving from flow monitors and controls the valves accordingly, as shown in figure 8.1.

Now consider that the scale of the problem is such that getting the measurements, mi(t), from each location, i, to the controller and subsequently propagating the signal, s(t), to the actuators is difficult, simply because of the distances between entities. This is the problem faced by many irrigation systems. Next, consider the situation in which the controller, C, is not a single entity with control over the system, but rather a collection of individuals who must cooperate to generate accurate mi(t), agree on how to transform them into s(t), and cooperate again to operate the actuator, based on the signals s(t). This additional complexity emerges because the scale of the task at hand is too large for any single agent to achieve. So challenging is this task, it has been argued that its solution required the evolution of complex societies (Wittfogel 1957). That the "controller" of irrigation societies comprises a collection of independent agents who must coordinate their activities is what makes this problem so interesting.

When societies construct irrigation systems and become adapted to regular variability, based on past experience, this can lead to the development (or emergence of) specialized institutions. Such adaptations may generate highly optimized complex systems that are robust to a particular range of variability but may be fragile to changes of that variability (Carlson and Doyle 2002). We hypothesize that many long-term irrigation systems developed a highly optimized tolerance (HOT) to a particular set of conditions but then became vulnerable to qualitative shifts in those conditions (a regime change). A system becomes "robust yet fragile" because it can never be robust to each possible disturbance; Csete and Doyle (2002) refer to this as "conservation of fragility." As a result, whatever adaptation is made to the irrigation system, it will always be fragile to certain disturbances. Thus, becoming more robust to disturbances of type A might increase the fragility of the system to disturbances of type B.

In this chapter, we extend some ideas concerning robustness from control theory, based on systems like that in figure 8.1, to explore the robustness characteristics of irrigations systems. We do not apply results from control theory strictly but rather use them as a motivation to explore the robustness of systems, based on trading off one type of robustness for another. Certain characteristics of this trade-off can be characterized precisely (Bode 1945). For the systems we are interested in, we do not precisely characterize this trade-off but rather explore the vulnerability trade-offs between various types of organizational structures in irrigation systems. Compared with most other chapters in this volume, we use more stylized models for which we can more thoroughly explore the parameter space that characterizes the underlying assumptions on which the model is built. We think that the use of styled models complements more comprehensive models because the simpler models can be used to understand the consequences of specific assumptions and to identify the main drivers of observed phenomena.

Two areas in which vulnerabilities may enter the system shown in figure 8.1 are explored in this chapter. First is that the controller is composed of individual agents interacting in a social network to achieve a common, larger-scale goal. Differences in goals, information-processing abilities, and the usual problems associated with collective action can influence the ability of an artificial social-ecological network to function (in this case, generate accurate $m_i(t)$, produce $s(t)$, and enable A). We use a model based on the Bali irrigated rice agriculture and temple system to explore this question. The second area is related to the way in which the investment infrastructure required to generate the feedback controller pictured in figure 8.1 affects the range of options open to deal with external shocks. To explore this issue, we use a simple model motivated by the Hohokam experience in the American Southwest.

Network Density, Interacting Agents, and Irrigation Performance

In this section, we illustrate the consequences of various assumptions about decision rules for controllers—in our case, communities deciding when to plant rice—on the performance of a rice irrigation system. We show that different assumptions about decision making have different sensitivities to a change in the density of pest-related connections among the communities. We use the Lansing-Kremer model of a Bali irrigation system, reimplemented in Java (Janssen in press), as an arena for our analysis (Lansing and Kremer 1993). The Lansing-Kremer model describes water flows and rice terrace ecology along two rivers in south-central Bali (Lansing 1991; Lansing and Kremer 1993).[2] Rainfall and the water from the volcanic lake provide water, allocated by twelve dams, for 172 *subaks* (collections of farmers). The runoff between dams is formulated as the difference between supply (runoff of dams from higher elevation and rainfall) and demand from the subaks related to each dam. When subaks ask for more water than is available, all subaks receive the same reduction of water supply, and the fraction of demand that is met depends linearly on a measure of water stress on the crops in these subaks. The time step of the model is one month.

A subak's water demand depends on rice variety and area planted. Each rice variety has a maturation time after which it yields a harvest per hectare, calculated by multiplying the rice variety's specific potential yield by the accumulated water stress. If a rice variety takes three months to grow and had water shortages of 0, 10, and 50 percent during each month, respectively, then the water stress is

$$\left(\frac{1+9/10+5/10}{3}\right)$$

which is equal to .8. Therefore, the harvest is 20 percent lower than the maximum potential yield.

The harvest can also be reduced by pest outbreaks. Each subak has a pest density, p, that changes because of migration and local growth. The direction and magnitude of pest migration depend on the gradient in concentrations between a subak and each of its neighbors.

For diffusion of pests, up to four adjacent neighbors are defined for each of the 172 subaks. Furthermore, for each subak, the source dam that provides the water is given, as well as the return dam for water that is not used. The source dam and return dam can be the same. Finally, seasonal rainfall patterns are known for each subak.

Local Control and Macro-level Patterns

Lansing and Kremer (1993) performed exercises in which they allowed subaks to imitate the cropping pattern from the neighbor with the highest production. We looked at two alternative models of agent decision making and illustrate here the implications of various strategies and changes in network structure on system performance. We can analyze the decision-making process of subaks from a control per-

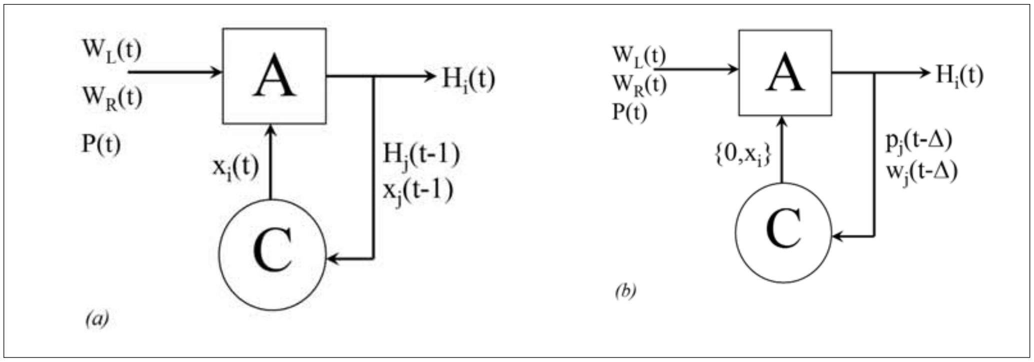

Figure 8.2. Schematic of subak decision making from a control system perspective: (a) the imitative subak and (b) the adaptive subak.

spective (figure 8.2). The input to the system, a subak, is water from the lake, $W_L(t)$; water from rainfall, $W_R(t)$; and pests from neighboring subaks, $P(t)$. An imitative subak derives information about cropping patterns, $x(t–1)$, and harvest, $H(t–1)$, of other neighboring subaks from the preceding year. The controller, C, makes a decision on the cropping pattern, x, for time, t. The adaptive subak updates each month and derives information on pests, p, and water availability, w, from the preceding month, $(t - \Delta t)$, of the neighboring subaks. In case a subak had no crop on its land, the controller decides whether to start planting rice, $\{0, x\}$. We describe the decision rules in more detail below.

Imitative Subaks.

The first type of agent is similar to the Lansing-Kremer model, but instead of defining a limited set of neighbors, we assume that subaks had access to the cropping patterns of all other subaks. We defined networks for water and pest dispersal and calculated the minimum number of connections it takes for each node to connect to the other nodes in the network. The network for water includes both dams and subaks as nodes. We assume that subaks that are more distant are less likely to be imitated. Differences of harvest between distant subaks are less influential on changing cropping patterns than differences between closely connected subaks. The harvest of a distant subak has to be significantly higher before that cropping pattern is imitated. This results in the decision rule of when to imitate a cropping pattern shown in equation 1. A cropping pattern in subak j will be considered for imitation by subak i if

$$H_i < \frac{Hj}{1+MIN\{\gamma_p \cdot \chi_p^2, \gamma_w \cdot \chi_w^2\}}$$

with sensitivity parameters, γ_p and γ_w, and the number of connections separating two subaks via pest dispersal links, γ_p , and water flow links, γ_w. From all the subaks meeting this condition, the subak with the highest harvest per hectare will be imitated. If no subak meets this condition, then the existing cropping pattern of the subak itself is repeated.

Starting with randomly distributed cropping patterns, subaks update their cropping patterns each year. Subak *i* compares the derived harvest per hectare with every other subak *j* but updates the cropping pattern only when the condition in equation 1 is met. This means that subaks take care of adjusting inequalities with their neighbors but generally do not change their cropping patterns when distant subaks perform better. This is a more general, but similar, implementation of imitating neighbors than worked out by Lansing and Kremer, who assume a fixed set of neighbors. We also assume that there is opportunity for experimentation. When a subak *i* performs worse than the average harvest per hectare within the watershed, it is assumed that there is a probability γ that the cropping pattern will be changed to a random configuration.

We assume that agents are able to learn and adapt their strategies by changing the values of γ_p and γ_W. Each year, agents update their values by comparing these with their direct neighbors. Neighbors with high harvest will have more impact on changing the values by using the equation

$$\gamma_i = \sum_{j \in N} \frac{\gamma_j \cdot e^{\mu \cdot H_j}}{\sum_{k \in N} e^{\mu \cdot H_k}}$$

where γ is a parameter that values the sensitivity of updating γ to value differences between the harvest values among the neighbors.

Adaptive Subaks.

An alternative, plausible way subaks can make decisions on cropping patterns is to make decisions during the year about whether to leave a field fallow or to plant a crop. We assume that a subak makes this decision based on the availability of water and dispersal of pests. A crop needs 150 cubic m/day of water per hectare. If water availability is expected to be above a certain threshold, m_W, then the subak may expect to have sufficient water to make planting crops worthwhile if the pest biomass per hectare among the neighbors and within the subak is, on average, below m_p. Similar to imitative subaks, adaptive subaks update their threshold values m_p and m_W in line with their best-performing neighbors.

The Effect of Changed Network Structures

For both the imitative subaks and the adaptive subaks, we explored the consequences if the neighbors with whom they are connected by pest dispersal change. In the default network are 161 connections among the subaks that disperse pests. We define probability p_e as the probability that an existing connection is deleted. Furthermore, we define probability p_n as the probability that a new connection is created. New connections are created between subaks who share a source and/or a return dam and are assumed to be geographically in the same neighborhood.

For each experiment, we run the irrigation for one hundred years one hundred times and calculate average harvest per hectare during the last fifty years of the simulation. When pe and pn are equal to zero, the harvest for the original network is calculated. The resulting harvest is similar to the value we derive when optimizing the thresholds γ_W and γ_p or m_p and m_W, indicating that the agents are learning to find high-performance threshold values. In figures 8.3 and 8.4, we see that the average harvest per hectare is sensitive to varying the probabilities. For the imitative subaks, the harvest decreases when links are added that can disperse pests. With more of those links, the subaks will synchronize on a larger scale, leading to water shortages. Decreasing the number of links that disperse pests does not affect the average performance of the subaks. The results for the adaptive subaks are different (figure 8.4). Adaptive subaks are not as sensitive to adding links that disperse pests. They are able to find suitable threshold values m_W and m_p to adjust for high connectivity; however, they are sensitive to adding additional connections. When the number of connections declines, the performance of adaptive subaks increases significantly. Note that the effects of adding and removing links is linear only to adding or only to removing links, which suggests that the effects are sensitive to density and not to topology of the network.

When we calculate the average harvest per hectare as a function of the original density of the pest-related network within the Bali irrigation network, we see that the imitative subak leads to the highest performance of the original network but is outperformed by the adaptive subaks (figure 8.5). There is no systematic understanding of how the real subaks make their decisions. Their decision-making processes are almost certainly much more comprehensive than those of our simple agents. Nevertheless, with regard to highly optimized tolerance, we argue that imitative agents are only suitable for a very particular density of pest networks. The results of Lansing-Kremer (1993) show that the use of imitative subaks is a simple solution to generate cropping patterns as observed. Our results show that this might be an artifact of the particular ecological system.

Irrigation Infrastructure, Short-term Robustness, and Long-term Fragility

In the preceding section, our focus is on the nature of the controller as represented by a network of subaks. The measure of the controller's effectiveness is harvest per hectare of a subak. In this section, we look at the problem of irrigation from a slightly different perspective. In many situations, the system being controlled actually comprises more than one subsystem. The actual control of the system then involves decisions about how to combine varying outputs from subsystems to meet an overall objective. In this case, social-ecological systems face two issues: not only the issue associated with controlling a system via a network of individuals, but also the problem of shifting vul-

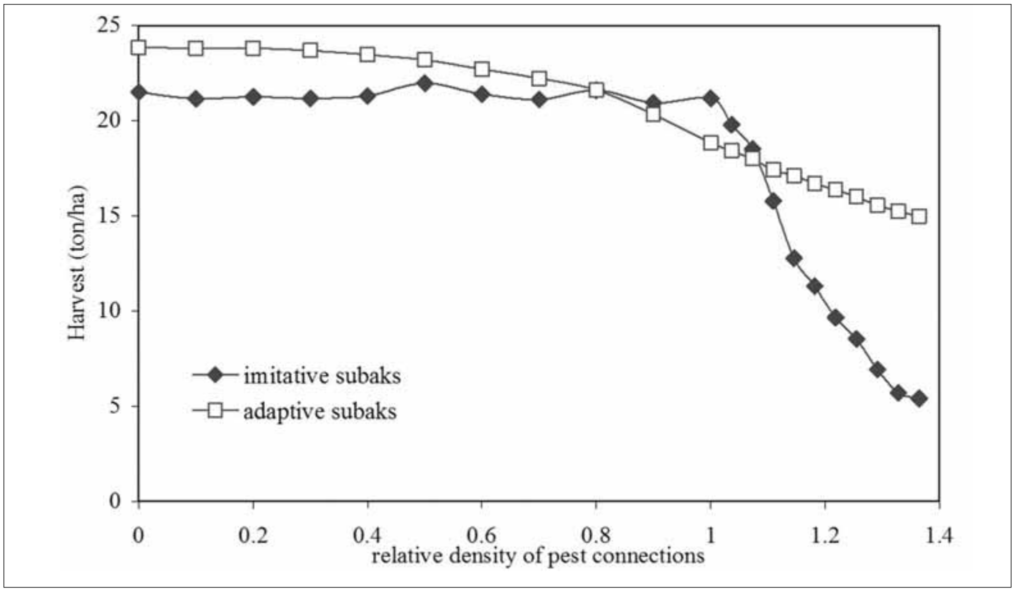

Figure 8.5. Average harvest per subak (per ha) as a function of the relation to the original pest network density (original density is equal to 1).

nerabilities associated with shifting emphasis from one resource type to another. Here we focus on this latter issue.

A common example of balancing a multiple resource system composed of elements with different risk characteristics is portfolio management in finance. In this case, productive resources are financial investments such as equities, options, and bonds. The problem is to design an optimal portfolio (proportion of each type of investment owned) at each time over a planning horizon, based on past performance of each investment (except in the case of options). Another example is a society that relies on a combination of irrigated agriculture and hunting and gathering of wild resources. In terms of risk, irrigated agriculture might be thought of as a bond. The yield of the irrigation-system capital stock (land, canals, and institutions) is agricultural output. The yield is less variable than the alternative of hunting and gathering, but when measured per unit of labor input, the yield is comparatively low. The high-yield activity of hunting and gathering is more subject to environmental fluctuations and is thus more variable. What is the best allocation of labor to irrigated agriculture and to hunting and gathering?

Obviously, societies typically do not make optimal investment decisions, and the time frames of interest are typically much longer than those in decision analysis problems. Nonetheless, ideas from portfolio management can provide important insights if applied in a broader context. For example, thinking in terms of asset management, we can ask the question of how the risks change as the asset mix between wild (hunting and gathering) and highly managed (irrigated agriculture) changes over time. Our focus is how the changing risk may feed back into the system and influence the coevolution of the social-ecological system.

Motivation: The Hohokam Cultural Sequence

The model developed below is based on the archaeological record from the American Southwest, specifically the Hohokam cultural sequence. The model is not intended to capture specific characteristics of the Hohokam experience but rather attempts to capture key structuring processes most likely at play. As such, the model could be equally well applied to any situation characterized by the use of wild and irrigated resources.

The context for the Hohokam cultural sequence is extremely rich and detailed, filling the pages of books (for example, Abbott 2003) and articles (for example, Bayman 2001). Only a very rough outline of major events is presented here, based on these sources and personal communication with archaeologists specializing in the American Southwest. The Hohokam occupied Arizona and northern Mexico from around AD 1 to 1450. Table 8.1 summarizes the key periods and events. This chronology tells of the expansion of irrigated agriculture and the development of pottery production, along with a regional trading system in exotic items associated with ball court settlements. Early in the sequence (the Pioneer period), houses were built in pits, settlements were small and dispersed, and populations probably relied most heavily on wild resources. During the Colonial period that followed, the archaeological record reflects a time of continued expansion, but with elements of the solidification of Hohokam culture as a regional system. Ball courts appear, possibly signifying the increased formalization of a regional trading network. More stable settlement patterns emerge in the form of courtyard groups. This period was probably associated with the continued heavy reliance on wild resources at a larger spatial scale, as evidenced by the regional trading networks. However, reliance on irrigated agriculture was probably also intensifying, and there may have been extensive trade between the wild and irrigated resource sectors.

During the subsequent Sedentary period, major aspects of Hohokam culture remain stable yet expand in scale, as evidenced by what was perhaps the mass production of pottery for an expanding, regional trading network associated with the ball court system. The signatures of this period give a sense of success; the culture enjoyed material abundance and ideological expansion. The areal extent of this cultural signature, which probably reached its maximum during the Sedentary period, covered approximately a third of present-day Arizona. Within the core area in the Phoenix Basin, an impressive irrigation infrastructure developed. However, this period of relative stability and success gave way to the period termed the "Classic," which saw "unprecedented changes in patterns of settlement, technology, material culture, and ideology" (Bayman 2001:281).

The abandonment of the ball court system and a contraction of the regional system marked this change. Platform mounds replaced ball courts, and aboveground residential areas with compound walls replaced the open-courtyard pithouse settlements. Community centers became more nucleated. Later in the Classic period, the platform mounds were walled off. Finally, toward the end of the Classic, just before the

Table 8.1 Summary of the Hohokam Cultural Sequence

Period	*General Characteristics*
Classic (AD 1150–1450)	Aboveground residential areas with compound walls emerge. Hohokam interaction outside Gila-Salt river valleys declines as the overall regional system shrinks. Rectangular platform mounds with compound walls dominate villages. Ball court system is abandoned. Platform mounds have similar spacing to ball courts, but community centers become more nucleated. Highly stylized crafts associated with ancestor worship disap pear. Population declines, and Hohokam culture collapses around AD 1450.
Sedentary (AD 900–1150)	Expansion from the colonial period continues. Mass production of pottery. Use of ball courts continues. Maximum extent of regional system reached.
Colonial (AD 750–900)	Period of expansion. First ball courts appear, and increased trade in exotic items is evident. Artistic florescence, accompanied by elaborate cremation rituals. Colonial Courtyard groups with shared ovens emerge. Ball court system expands, related to regional exchange networks.
Pioneer (AD 1–750)	Irrigation system begins to develop; first canals built. Bow and arrow begins to be used in the Southwest. Irrigation systems continue to expand, and large canal systems appear on the north and south sides of the Salt River.
1500 BC–AD 1	Hunter-gatherers with limited agriculture. Small pithouse settlements and seasonally occupied hamlets were typical.

Hohokam cultural collapse, massive structures now called "great houses" were built. This historical progression can be seen as a shift in resource portfolio from one dominated by an extensive, highly productive (per unit of labor) but highly variable resource to one dominated by a less productive (labor-intensive) resource. As the resource portfolio shifted, the system became exposed to different risks. In terms of the schematic shown in figure 8.1, the system can be pictured as a control system with two subsystems and two controllers (figure 8.6). The objective of the controller is to produce a sufficient food supply with minimum effort, given environmental variation.

The top subsystem consists of uncultivated land and a small amount of land in riparian areas cultivated with flood irrigation—that is, with little built infrastructure. On the left, the various vulnerabilities associated with each resource base are listed. Food output from wild resources is vulnerable to local drought. Food output from riparian agriculture is less susceptible to drought but more susceptible to flooding. Food output based on a portfolio composed of these two resources would be robust to short-term fluctuations in rainfall. This robustness is due to their being "orthogonal" in risk space—that is, they are robust against types of shocks that typically do not occur simultaneously. The bottom system—irrigated agriculture—is very robust against short-term fluctuations but is vulnerable to infrequent, large floods. Furthermore, the irrigation system becomes vulnerable to new types of internal shocks associated with the collective action required to maintain infrastructure.

The shift to irrigated agriculture from a portfolio dominated by wild resources and small-scale riparian agriculture is related not only to managing environmental risks but also to population growth. Success of the top system in figure 8.6 may generate population growth and, with it, a rise in food demand. As food demand surpasses the maximum possible return from wild resources and riparian agriculture, intensification of production is required, and movement to irrigated agriculture becomes necessary. The Hohokam cultural sequence can be mapped onto this asset reallocation process. The pre-Pioneer period is characterized mainly by wild resource use. The Pioneer period probably saw a gradual increase in riparian agriculture. Toward the end of this period, some irrigated agriculture began to augment the wild resource/riparian agriculture portfolio. The Colonial period, with the expansion of the ball court system, probably saw an increase in the areal extent of the wild resource base, with concentration and intensification of riparian and irrigated agriculture in the most suitable areas. The ball court system may have played an important role in mediating the exchange of outputs from the various resource types as they became more spatially distinct.

Throughout this period, food production from both agriculture and wild resources may have increased, with a slow shift in dominance toward irrigation. It seems that in the Sedentary period the maximal extent of food production from wild resources may have been reached, so the use of both types of resources was no longer expanding. Thus, intensification of irrigated agriculture increased at the expense of wild resource use. Crossing this threshold may have induced some of the social changes seen in the archaeological record—that is, the shift from the extensive ball court system to the intensive platform mound system in the Colonial period.

This shift in the resource portfolio, driven by population growth, changed the risks faced by the Hohokam. It is now clear why this situation is different from optimal portfolio management. The Hohokam did not optimally manage output from different resource classes. Rather, at each point in time they responded to a wide range of pressures that forced them into a certain development trajectory. Along this trajectory, vulnerabilities changed. An interesting question to explore is whether and how these shifting vulnerabilities may have played a role in the Hohokam cultural collapse. The model in the next section focuses on this question.

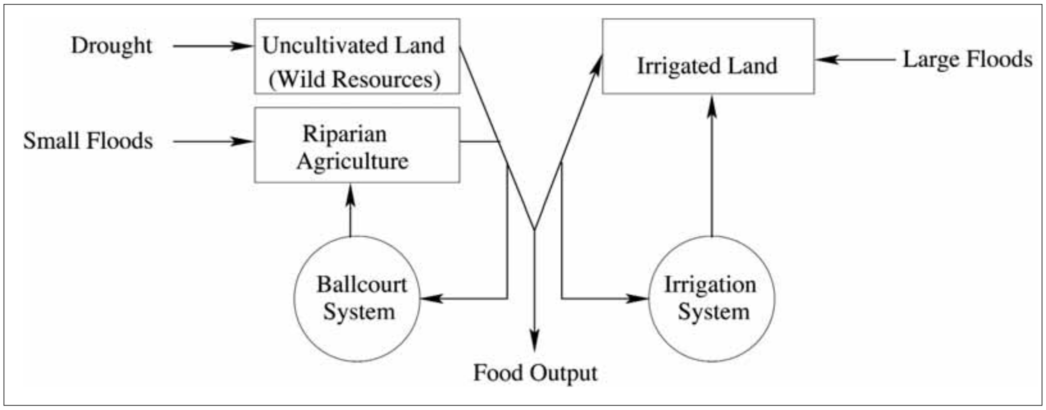

Figure 8.6. Control system consisting of three resource types (see text for explanation).

A Simple Model of Extensive and Intensive Food Production

This model, a formal representation of the system shown in figure 8.6, is motivated by several simple, representative, agent bioeconomic models of renewable resource use (Anderies 1998, 2003; Brander and Taylor 1998; Janssen and Scheffer 2004). We do not make a sharp distinction between riparian and irrigated agriculture but instead represent them on a continuum from low to high physical-infrastructure systems. Thus, we have two resource types: extensive (type 1) and intensive (type 2). The resources naturally self-regenerate and are degraded by exploitation. For the extensive resource, this regeneration is in terms of biomass that can be harvested, as in fishery models (Clark 1990). For the intensive resource, the regeneration is in terms of soil fertility—that is, soil microorganisms generate nutrients for plant uptake over time (Anderies 1998, 2003).

The society utilizes these resources to produce two types of output: protein-rich (type 1) and carbohydrate-rich (type 2). Both types of output can be produced from both resource types. The wild resource base, however, has potentially higher productivity of type 1 resources, whereas irrigated agriculture has potentially higher productivity of type 2 resources. To model the society's choices regarding the mix of resources to produce, we assume that it attempts to meet its basic needs with as little labor as possible. This is equivalent to assuming that individuals prefer leisure to additional food when their basic needs are being met. Note that "basic needs" are not equivalent to subsistence needs: what an individual defines as basic needs may be well above subsistence, including food for gifts, rituals, and so forth.

The simplest mathematical representation for this system comprises a description of the resource bases' dynamics over time, as well as the allocation of labor to gathering wild resources and conducting irrigated agriculture. The former is given by two differential equations:

$$\frac{dx_1}{dt}=r_1(R)(1-x_1)-\alpha_{11}Y_{11}-\alpha_{12}Y_{12}$$

$$\frac{dx_2}{dt}=r_2(S)(1-x_2)-\alpha_{21}Y_{21}-\alpha_{22}Y_{22}$$

where $r_1(R)$ is the intrinsic regeneration rate of the wild resource stock, x_1, which depends on rainfall, R. When left alone, the wild resource will increase to its maximum level (which has been rescaled to 1). Y_{ij} is the quantity of output of type j from resource i, and α_{ij} is the impact of producing output type j from resource type i. The regeneration rate of irrigated agricultural soil productivity, x_2, depends on the stream flow, S. For this analysis, aimed at characterizing the possible dynamics, we assume "average" conditions—that is, R and S are constant, so r_1 and r_2 are constant.

For clarity and simplicity, we assume a linear production structure. Note that lower- and uppercase letters represent per capita and total quantities, respectively ($Y_{11} = hy_{11}$, where h is the total population size). Thus, $y_{1j} = A_{1j}x_1 l_{1j}$, where l_1j is the labor devoted to harvesting output of type j from resource 1 by a single representative agent. This is a simple mass-action production function commonly used in natural resource economics (Clark 1990). Similarly, $y_{2j} = A_{2j}x_2 l_2 K$. We do not make a distinction between labor directed at producing different types of output from resource 2—these come in a fixed ratio determined by A_{2j}. Output from irrigated agricultural activity depends not only on labor but also on capital, K. We refrain from defining this stock of capital precisely at this moment, but it certainly includes irrigation infrastructure. Given these definitions, each individual (all identical) chooses l_{ij} to meet her needs and minimize her total labor. This problem can be stated as follows:

$$\begin{aligned}
&\text{Min } l_{11}+ l_{12}+ l_2\\
&\text{Subject to}\\
&y_{11} + y_2 \geq y_{1\text{min}}\\
&y_{12} + y_{22} \geq y_{2\text{min}}\\
&l_{11} + l_{12} + l_2 \geq l_{\text{max}}
\end{aligned}$$

where y_{imin} is the minimum requirement for each resource type. Because the objective is to minimize labor in meeting basic requirements, equations 5 and 6 become equalities, and by rescaling, the right-hand sides can be replaced with 1s. The solution to this optimization problem yields the following rule: Define

$$l_2=\min\{1/A_{21}x_2K,\ 1/A_{22}x_2K\}.$$

Then, if

$$1-\frac{x_2K}{x_1}\left(\frac{A_{21}}{A_{11}}+\frac{A_{22}}{A_{12}}\right)>0 \qquad \text{(eq. 8)}$$

then $l_2 = 0$. Otherwise, . When l_2 is known, y_2 can be computed, and y_{11} and y_{12} can be computed from equations 5 and 6. When these are known, multiplying each y_{ij} by the total population size, *h*, yields Y_{ij}, which determines the ecological dynamics via equations 2 and 3.

Several points emerge from this simple model. First, equation 8 makes the conditions favoring a switch to irrigated agriculture clear: as the left side of equation 7 becomes negative, $l_2 = l_2 \neq 0$ so that society devotes some labor to irrigation. This becomes more likely as *K* increases (obviously), because the wild resource becomes degraded relative to irrigated agriculture (the ratio x_2/x_1 increases) or the productivity of labor in irrigation is higher than that of labor devoted to wild resources. Condition 8 is sensible and intuitive. What is not intuitive and requires further analysis is the effect of population and capital stocks on the system dynamics. Simple phase-plane analysis illustrates the point (figure 8.7).

Conceptually, the model provides a preliminary mathematical formalization of the relationships among capitalization, population, shifts in resource utilization (wild resource versus irrigated crops), and associated changes in the loci of critical vulnerability (rainfall versus stream flow versus capital stocks). Here we treat population and capital stocks as parameters and explore their effect on resource use decisions. When population density is zero, the dynamical system (equations 2–8) exhibits a single stable equilibrium such that, given any initial condition, the system will move to pristine conditions of full soil fertility and high levels of wild resource biomass (the basin of attraction is the entire light gray region in figure 8.7a). Because any initial condition will return to the steady state, the system is very robust to external perturbations (resets system to new initial condition) that may reduce wild resource biomass. Figure 8.7a characterizes the Hohokam in the pre-Pioneer and early Pioneer periods.

A second stable attractor emerges (the solid circle in the dark gray region of figure 8.7b) when a critical population size is reached. These two new attractors correspond to a situation in which either the wild resource base is in good condition and able to support the population (the light gray region) or the wild resource base has become degraded and the society has introduced irrigated agriculture to complement the wild resource. Equation 8 shows that the threshold between these basins of attraction depends on the amount of investment in agricultural infrastructure (for example, canals) and the relative productivities of the two resources. There is the interesting possibility that the society may move between states. In the steady state in the light gray region (figure 8.7b), wild resources dominate the society's resource portfolio, and irrigation plays a small role. For stable population levels, the only way the system could be moved to the equilibrium in the dark gray region is by a drought (for example, wild resource biomass is reduced from .6 to .3). Only in severe drought conditions

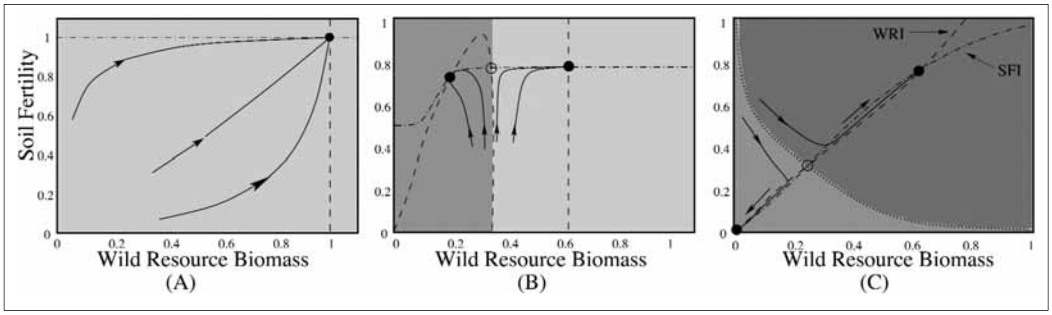

Figure 8.7. Phase plane analysis of the Hohokam subsistence agriculture model presented in equations 2–8. The curves with arrows show possible trajectories for various initial conditions. The arrows show the direction of flow over time. The population increases from (A) to (C).

does it pay this society to move into more intensive irrigation. Note that the decision to intensify irrigation locks this society into the degraded wild-resource state. Now a large perturbation (that is, a very wet phase that increases wild resource biomass from the equilibrium level of around .2 to .33) is required to push the system back into the alternate basin of attraction. In this way, the longer-term effects of social responses to immediate problems limit future options. This dynamic may characterize patterns of Hohokam development in the Colonial and early Sedentary periods. It is interesting to speculate whether an exceptionally dry period was associated with the shift away from the ball court system to the irrigation system.

As the population increases, the relative size of the wild resource basin shrinks (that is, the boundary moves to the right) (figure 8.7b, light gray), reducing the size of a perturbation (for example, drought) required to push the system from reliance on wild resources into the irrigation basin (figure 8.7b, darker gray). This illustrates how increasing population makes the system more vulnerable to ever-smaller drought events. Moreover, higher population decreases the likelihood that the system will return to the wild resource state after it is in the irrigation basin. At even higher population and capitalization levels, the wild resource basin vanishes (figure 8.7c), and two different basins emerge, including an undesirable one (lighter gray) in which the society cannot sustainably meet its minimum needs. The dynamics illustrated in figure 8.7c are fundamentally different than in figures 8.7a and 8.7b. In figures 8.7a and 8.7b, the stable equilibria are always positive—that is, neither resource base is destroyed, and the society can always meet its basic needs. In figure 8.7c, one equilibrium is positive (in the dark gray region), and the other is at zero. When the state enters the basin of attraction corresponding to the zero equilibrium, the resource base will inevitably be destroyed, and the human population reduced. Furthermore, the social response of additional intensification cannot correct this problem, and new vulnerabilities emerge, as described below.

Increasing population and capital stocks causes the boundary between basins to move away from the origin, decreasing the distance between the undesirable basin and the desirable attractor (the black dots in figure 8.7). In this situation, all else being equal, a small perturbation in the capital stock (for example, destruction of a canal headgate) causes the dashed curve (the wild-resource biomass isocline labeled *WRI* in figure 8.7c) to move up and leftward and the dash-dot curve (the soil fertility isocline labeled *SFI* in figure 8.7c) to move down and rightward. From figure 8.7c, it can be seen that this relative movement of these two curves will cause their intersection points (the open and solid circles) to move closer together. Eventually, the two curves will no longer intersect, and the desirable basin will vanish. This situation, depicted in figure 8.7c, may correspond to events during the Hohokam Classic period.

The Dynamic Interplay between Social Change and Shifting Vulnerabilities

The simple model and analysis presented here provide a caricature of the leapfrog process of shifting vulnerability and social change. A society's success (in feeding itself) may cause it to move from figure 8.7a to 8.7b. In so doing, it becomes more vulnerable to droughts. If a sufficiently large drought occurs, then the society must respond by intensification (move from light to dark gray regions in figure 8.7b). If society does reorganize around the new irrigation equilibrium, then it becomes more robust to drought (the system will always return to the equilibrium if perturbed to the left of the equilibrium point in the dark gray region of figure 8.7b). However, as the notion of "conservation of fragility" suggests, the society has traded off one kind of vulnerability for another. Worse yet, the movement from the high wild-resource, low intensive-irrigation equilibrium (in the light gray region of figure 8.7b) to the low wild-resource, high-intensive irrigation equilibrium (in the dark gray region of figure 8.7b) is difficult to reverse and becomes more so as population increases. Put another way, the society typically has no option of doing intensive irrigation for a while until things get better and it can return to an easier life style. Finally, again because of its success in coping with vulnerability, society moves from figure 8.7b to 8.7c. Now the situation is more unforgiving, and the system becomes vulnerable to any shock that reduces the capital stock (flood, social unrest, and the like). The process is one of social change in response to particular vulnerabilities, leading to new vulnerabilities, leading to further social change, and all the while leading to fewer and fewer options for society.

Discussion

Irrespective of the efforts to increase the robustness of an irrigation system, fragility always remains in the system. Because of the temporal and spatial heterogeneity of disturbances, adaptation to some disturbances may make the system more vulnerable to others. Our Hohokam analysis shows that an increasing population invests in irrigation systems to feed itself. However, this makes the system more vulnerable to large floods that destroy the infrastructure. Although irrigation systems lead to a higher production by controlling high-frequency variability such as small droughts and floods, they cannot rapidly recover from a low-frequency, high-impact perturbation. In the Bali case study, we saw that different assumptions on decision making of the subaks lead to different vulnerabilities. Although imitation of best-performing neighbors maximizes production in the original pest-related network, it does not do so when the density of the links changes. In some sense, the system performance is enhanced by particular behavior in a given network but becomes vulnerable to changes in the network.

The results reported in this study are initial results of a longer-term effort to develop computational models capable of simulating the evolution of irrigation systems. How do various assumptions affect the structure of the network? How do robust irrigation systems emerge? Do irrigation systems evolve to a HOT situation, being robust for high-frequency events but not for low-frequency events? In the longer term, we aim to "grow" irrigation systems and understand which assumptions about human behavior, institutions, ecology, and biophysical structure affect a system's development and, in turn, its robustness.

Acknowledgments

This research was funded by National Science Foundation grants BCS-0601320 and BCS-0527744.

Notes

1. Feedback control systems are used for real-time irrigation and drainage-canal control (Malaterre, Rogers, and Schuurmans 1998; Reddy and Jacquot 1999). However, we do not include these modern feedback systems in our analysis but focus solely on feedback control as a broad class of decisions irrigators had to make to control their irrigation systems.

2. For a visual representation of the Bali irrigation system, see Steve Lansing's homepage: http://www.ic.arizona.edu/~lansing/home.htm.