

Coordination in irrigation systems:
An analysis of the Lansing-Kremer model of Bali

Marco A. Janssen

School of Informatics, Indiana University

&

Center for the study of Institutions, Population, and Environmental Change

Indiana University

408 North Indiana Avenue

Bloomington IN 47408, USA

maajanss@indiana.edu

Phone: 812 855 5178

Abstract

Farmers within irrigation systems, such as those in Bali, solve complex coordination problems to allocate water and control pests. Lansing and Kremer's 1993 study of Balinese water temples showed that this coordination problem can be solved by assuming simple local rules for how individual communities make their decisions. Using the original Lansing-Kremer model, I analyzed the robustness of their insights and found that the ability of agents to self-organize is sensitive to pest dynamics and assumptions of agent decision making.

Keywords: irrigation, coordination, networks, synchronization, agent-based model

1. Introduction

The question of whether irrigation systems require centralized authority to solve complex coordination problems has held the interest of scholars for a long time, but empirical analysis has not provided a clear answer (Hunt, 1988). Wittfogel (1957) argued that a central control was inevitable for larger irrigation systems and hypothesized that some states have emerged because of the use of irrigation. On the other hand, there are various examples of complex irrigation systems and drainage systems that have evolved without central coordination. The drainage systems and the water boards in the Netherlands are an interesting example of this. Irrigation on the Indonesian island of Bali has been a source of debate on the origin of state between those who favour an important role of the state, and those who argue that the state was not essential for the coordination of irrigation systems (Geertz, 1980; Lansing, 1991). I will not delve into debate on the role of the origin of state in Bali (see Hauser-Schäublin, 2003 for a discussion) but will show the consequences of different interpretations of the role of the state in recent history.

The Bali irrigation system existed for more than a thousand years, and the Bali *subak*¹ system of coordination was almost destroyed within a decade of national intervention to maximize rice production (Lansing, 1991). Due to insights of anthropologists and ecologists in the functioning of the system, a collapse was prevented and the system largely recovered, although it is still under threat of external disturbances.

The irrigators have to solve a complex coordination problem. On one hand, control of pests is most effective when all rice fields have the same schedule for planting rice. On the other hand, the terraces are hydrologically interdependent, and to balance the

¹ A *subak* is an organized group of farmers in a little village.

need for coordinated fallow periods and use of water, a complex calendar system that states what actions should be done on each specific date for each *subak* has developed.

These actions are related to offerings to temples, starting with the little temples at the rice terrace level, to the temple at the village level, to the regional level, and then up to the temple of the high priest Jero Gde, the human representative of the Goddess of the Temple of the Crater Lake. Crater Lake feeds the groundwater system which is the main source of water for irrigating in the entire watershed. These offerings are collected as a counter gift for the use of water that belonged to the gods.

The function and power of the water temples were invisible to the planners involved in promoting the Green Revolution during the 1960s. They regarded agriculture as a purely technical process. Farmers were forced to switch to the miracle rice varieties, which were predicted to lead to three harvests a year, instead of the two of the traditional varieties. Farmers were stimulated by government programs that subsidized the use of fertilizers and pesticides. After the government incentive program was started, the farmers continued performing their rituals, but they no longer coincided with the timing of rice-farming activities. Soon after the introduction of the miracle rice, a plague of plant-hoppers devastated the rice production. A new variety was introduced, but then a new pest plague hit the farmers. Furthermore, there were problems of water shortage.

During the 1980s, an increasing number of farmers wanted to switch back to the old system, but the engineers interpreted this as religious conservatism and resistance to change. It was Lansing (1991) who unraveled the function of the water temples and was able to convince the financiers of the Green Revolution project on Bali that the irrigation was best coordinated at the level of the *subaks* with their water temples.

Anthropologist Steve Lansing and ecologist Jim Kremer built a computer model of an artificial ecosystem and showed that for different levels of coordination, from farmer level up to the level of the watershed, the temple level was the level of scale where decisions could be made to maximize the production of rice (Lansing and Kremer, 1993).

In this paper I analyze Lansing and Kremer's original model in depth in order to understand why the temple level would be the best level for coordination. In their original analysis, they provided some illustrative simulations (Lansing and Kremer, 1993, p. 106), but they performed no rigorous analysis to provide sufficient insight into understanding the tradeoffs. One of the problems is that social coordination is not the same as synchronizing cropping plans. Therefore, I explore different possibilities for governing the irrigation network at the *subak* and watershed levels. Suppose we look at coordination at the watershed level: if a central planner were to optimize the cropping plans of all individual *subaks*, total rice production would at least be at the level of optimizing cropping plans synchronized at the temple level. The more degrees of freedom to tailor coordination among *subaks*, the higher the harvest might be. However, a higher aggregated harvest might be at the cost of the harvest within individual *subaks*. In my optimization experiments I wanted to know the tradeoffs of more detailed coordination and inequality of harvest among the self-supporting *subaks*. I also explored different decision rules at the *subak* level to determine if they might lead to high-level performance at the watershed level.

It is important to understand under which conditions it is possible for *subaks* to make decisions on cropping plans that lead to high-level performance of the irrigation

system. This paper aims to contribute to this endeavour. In section 2 I discuss the original Lansing-Kremer model. Section 3 discusses the potential total harvest when all *subaks* cooperate with different assumptions on rainfall and pest dynamics. We will see that there is a trade-off between total harvest and inequality. In section 4, I explore in detail with a simplistic two-node model the tradeoffs between pest dynamics and water supply. Bottom-up solutions for different decision algorithms are explained in section 5. Section 6 concludes.

2. The Lansing-Kremer model

The Lansing-Kremer model of the Bali irrigation system describes the water flows and rice terrace ecology along two rivers in south-central Bali (Lansing, 1991; Lansing and Kremer, 1993).² Low, middle, and high estimates are given for seasonal rainfall patterns at various elevations. Rainfall and the water from the volcano lake provide the water for the 172 *subaks*. Twelve dams allocate the water to the *subaks*. 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 there is supply, all *subaks* receive the same reduction of water supply, and the fraction of demand that is met is linearly assumed to be a measure of water stress of the crops in these *subaks*. The time step of the model is one month, and 49 cropping plans specify what crop is growing in a *subak* each month, e.g., triple cropping a high-yield rice variety or planting two traditional varieties with six- and four-month maturation times and one-month fallow periods between them.

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

Water demand of a *subak* depends on which crop variety is planted and the area of the *subak*. Each rice variety has to grow for a number of months. After this period the harvest is calculated by multiplying the rice variety's specific potential yield times the accumulated water stress. If a rice variety takes three months to grow and had water shortages of 0%, 10%, and 50% during each month, respectively, the water stress is $(1+9/10+5/10)/3$ which is equal to 0.8, and thus the harvest is 20% lower than the potential yield.

The harvest can also be lowered by damage from pest outbreaks. Each *subak* has a pest density p which changes by migration of pests and local growth. The direction and magnitude of the migration of pests depends on the gradient in concentrations between a *subak* and each of its neighbors. If a *subak* has four neighbors, the rate of change in pest level can be described as³

$$p_j(t+1) = g(x_j) \cdot (p_j(t) + \frac{1}{2}(p_{n1,j}(t) + p_{n2,j}(t) + p_{n3,j}(t) + p_{n4,j}(t) - 4p_j(t))) + \frac{1}{2}(p_{n1,j}(t) + p_{n2,j}(t) + p_{n3,j}(t) + p_{n4,j}(t) - 4 \cdot p_j(t)) \quad (1)$$

with $g()$ as the growth rate of pest p_j on *subak* j , depending on whether rice is growing in the field or not. P_{ni} refers to neighbors i of *subak* j . When rice is in the field, $g()$ is between 2 and 2.4; when the field is fallow, it is 0.1.

³ Note that one would expect a diffusion equation like $p_j(t+1) = g(x_j) \cdot p_j(t) + (p_{n1,j}(t) + p_{n2,j}(t) + p_{n3,j}(t) + p_{n4,j}(t) - 4 \cdot p_j(t))$, as well as a dynamic growth model of rice, but probably due to the limited computational power of PCs in the late 1980s, a shortcut was used to calculate pest dynamics with a monthly time step, and a fixed potential rice production at the end of the growth period was used. Since I analyzed the results of the original Bali irrigation model, I also used this shortcut as defined in equation (1).

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.

Lansing and Kremer distinguish six levels of social coordination, which are analyzed in separate model experiments. The assumption is that within a group of *subaks* the planting and harvesting occur at the same times, which means that they synchronize. The first level that is considered is one group of all 172 *subaks*. The second level of synchronization is two groups, the highlands and the lowlands. The third level is seven groups as pairs of temples. The fourth level distinguishes the 14 Masceti temples. The fifth level distinguishes the 28 groups at the Ulun Ski temples, and the sixth level considers each *subak* as a separate group.

3. Cooperative solutions at different levels of synchronization

The original code of the Lansing-Kremer model was reimplemented in Java⁴ and used to perform a number of optimization experiments to investigate the potential harvest level of the system. Since I ignored transaction costs, such a solution is not very realistic, but it provides us a benchmark of potentials within the system. Furthermore, I was interested in the sensitivity of the solutions for different assumptions on rainfall and the growth and dispersal rate of pests.

The original model included 49 cropping plans, which is reduced to 21 by not including plans with vegetables (since the objective function is rice production). Given

21 cropping plans, and 12 starting months of the cropping plan, I searched for the plan that maximizes rice production for each level of coordination. I used a period of 10 years and based the maximum on yield for the last five years to avoid initialization problems.⁵ The six levels of coordination were the same as used by Lansing and Kremer: 1, 2, 7, 14, 28, and 172 groups. The optimization was performed by a heuristic local search routine (hill climbing), which draws a group randomly and optimizes the cropping plan and starting month, given the existing values of the other groups, and updates the solution with the best local solution. Due to the character of the local search routine, the optimization was performed with multiple starting points. Nevertheless, due to the nonconvexity of the solution space, a global optimum cannot be guaranteed, except for the first two levels, in which all possible solutions are investigated. The optimization criterion was the total rice harvest of all 172 *subaks* over the last five years of the 10-year simulation period.

The results show that with an increasing number of smaller groups, there is a higher amount of total rice harvest (Figure 1). This is to be expected, since a solution of a small number of large groups is one of the possible solutions when the flexibility of smaller levels of synchronization exists. Thus increasing the number of groups should lead to the same or higher harvest because there is more flexibility to tailor the cropping plans. Interestingly, there is also an increasing inequality between annual harvest levels of *subaks*. Some *subaks* must give up production in favor of more productive *subaks*. Since *subaks* are self-supporting this inequality signals a potential source of conflict. One

⁴ Some small errors in the network of dams and in the network of pest diffusion were found in the original code and have been corrected. These errors do not change the results of Lansing and Kremer's (1993) model experiments in a qualitative way.

cannot assume that *subaks* will reduce their harvest significantly in favor of the production of nearby *subaks*. In the analysis below we will study the case where *subaks* have decision rules to change their crop plan given the information they have on production of their neighbors, the available water, and pests in the neighborhood. In that case, local interactions reduce inequality.

[Figure 1]

I analyzed the consequences of different deterministic levels of rainfall and stochastic rainfall variation. For stochastic rainfall I used the average over 10 simulations. Figure 2 shows that the harvest and inequality levels are more dependent on the level of synchronization than on variation of rainfall. I allowed all three crops, but all solutions from the optimization experiment favor crop number 3, a rice variety that has high yield but is sensitive to pests. I have not explored the sensitivity of the solutions to adding additional rice variations by varying the potential yield and sensitivity to pests. That should not affect the qualitative nature of the results. In scenarios with low rainfall there are quite a number of *subaks* with two harvests per year, skipping a cropping during the dry season. With higher levels of rainfall an average of three crops per year is dominant.

[Figure 2]

⁵ This is quite different from the original Lansing-Kremer simulations, which run the model for only one year.

Variation in the growth rate of pests has an important effect on optimal harvest levels (Figure 3). The benefit of synchronization is only derived for the medium growth rate of pests. When the pest growth rate is low, pests do not matter, and *subaks* only care about water coordination. Given a medium rainfall almost all *subaks* plant three crops per year. There is some loss due to water shortage, but the *subaks* do not have to let the system rest for a long period to reduce the pest population. When pests have a high growth rate, the *subaks* switch to two harvests per year to allow sufficient time for them to die out.

[Figure 3]

Different pest growth rates also have an important effect on inequality (Figure 3). Namely, there is not an increasing level of inequality like in the other experiments and I will show later that this affects the best level of synchronization. Finally, I analyzed the consequences of different dispersal rates, given medium rainfall and medium pest growth rate. When the pest spreads quickly, the harvest is severely affected (Figure 4) and the *subaks* switch to two harvests per year, again to allow time for the pests to die out. If the pest spreads less quickly, the benefit is marginal.

[Figure 4]

The results of these optimization experiments provide some interesting findings on the impacts of pest dynamics in both growth rate and dispersal rate. I wanted to

understand the synchronization problem related to pests in more detail and used the simplest possible irrigation network, namely consisting of two nodes.

4. A two-node irrigation model

The simplest possible model to study synchronization and coordination in irrigation systems is to distinguish upstream and downstream nodes. Suppose water supply is first available to the upstream node n_u and the leftover water is available for the downstream node n_d . In line with Lansing and Kremer's work, we consider a second problem for coordination: pests. In a two-node irrigation model pest biomass is defined as

$$p_u(t+1) = g(x_u) \cdot (p_u(t) + 0.5 \cdot (p_d(t) - p_u(t))) + 0.5 \cdot d \cdot (p_d(t) - p_u(t)) \quad (2a)$$

$$p_d(t+1) = g(x_d) \cdot (p_d(t) + 0.5 \cdot (p_u(t) - p_d(t))) + 0.5 \cdot d \cdot (p_u(t) - p_d(t)), \quad (2b)$$

where $p_u (\geq 0)$ and $p_d (\geq 0)$ are the levels of the pest upstream and downstream, respectively. The growth rate g is dependent on whether rice is in the fields, g_r , or whether there is no rice in the fields, 0.1. The diffusion rate d affects how fast pests move between nodes.

I considered two periods in a year, and a node therefore could have one out of three strategies for planting a crop: no crop, crop in the first period, and crop in the second period. If no crop is in the field, the growth rate of the pests is equal to 0.1, and when rice is in the field the growth rate of the pests will vary. I considered only one type of crop, which had a yield of one. In the first analysis I assumed that sufficient water was available. Therefore the coordination is solely based on controlling pest outbreaks. Figure

5 shows the total harvest of both nodes for different growth and dispersal rates of pests.. This total harvest is the maximum harvest if both nodes were to cooperate, and I analyzed all nine options for each combination of growth rate and dispersal rate. The harvest is calculated for 100 years to ensure convergence of the pest population, and the harvest in the last, converged year is depicted in Figure 5.

When the growth rate is less than 10, both nodes can plant a crop. Since the growth rate of pests during the fallow period is 0.1, and the average growth rate for a whole year is smaller than one when the pest growth rate during the cropping season is less than 10 ($g \cdot 0.1 \leq 1$ if $g \leq 10$). A higher growth rate leads to an explosion of pests, and the nodes need to coordinate. For high growth-rate levels, one crop can be planted depending on the dispersal rate. If the dispersal rate is low, the pest density remains too high in the node such that the pest population explodes and damages the crop. If the dispersal rate is high, the pest density in the other node increases to such a high level that pest outbreaks damage crop production. Only if dispersal is within boundaries, such that the pest population spreads evenly among the nodes, and remains below a certain level, the pest population can be controlled.

[Figure 5]

We now increase the complexity of the two-node model by assuming 12 different periods (months) within a year, leading to a larger variety of cropping strategies. A cropping pattern determines when three-month crops are planted. Using one month of fallow after each harvest allows a maximum of three crops within a year. When two crops are

planted, there are three variations with the maximum number of months of fallow.

Together with a one-crop option there are five different types of cropping patterns that can be started in one of the 12 months. The cropping plans from the total of 3600 possible plans, $(5*12)*(5*12)$, maximize the total harvest of the two nodes. The harvest H in a node is determined by

$$H = (1 - \text{MIN}(1, P)) \cdot WS, \quad (3)$$

where WS is the water scarcity during the three months in which the crop grows. Given the rainfall above the upper node, that node extracts one unit of water, and the remaining water is available for the lower node. For example, during a three-month period there are 1.5 water units available for the upper node, after extracting the water there is only 0.5 units of water available for the lower node. In this case, WS is 1 for the upper node and 0.5 for the lower node.

Figures 6, 7, and 8 depict the harvests for rainfall of 2, 1.5, and 1 units of water per month (no seasonal fluctuations). With 1 unit, only one node will have a sufficient amount of water. With 2 units, there is no water constraint. If nodes synchronize, a fallow month reduces the pest biomass to 10% of the original value. If the growth of pests in the three months when the rice is in the field is less than 1000%, the pest biomass does not grow beyond the initial values. Thus if the growth rate is smaller than 2.14, or $\sqrt[3]{10}$, pests cannot grow exponentially and a maximum number of crops is possible (Figure 6). Beyond this growth rate, we see a drop in the maximum harvest. Like the previous model, a high level of crop harvest is possible if dispersal does not lead to high pest

density. Due to the different temporal structures of cropping, a different pattern is derived.

[Figure 6]

Figures 7 and 8 depict the total harvest when there is a constraint to the availability of water. Due to the water shortage, it is not possible to grow the maximum amount of crops. Higher dispersal and growth rates of pests reduce the total harvest level. The distinct jumps are caused by the discrete nature of the cropping pattern. The results confirm the huge differences found for growth rates 2, 2.2, and 2.4 in the original Lansing-Kremer model (with a dispersal rate of 0.3 and almost no water shortages). The parameter range in the Lansing-Kremer model is 0.18 to 0.45 for the dispersal rate, and 2 to 2.4 for the growth rate.

Lansing and Miller (in press) also analyzed a two-node version of the coordination problem. They analyzed cooperation between a downstream *subak* and an upstream *subak* and showed that cooperation is a rational strategy when the pest is a key problem. The downstream *subak* traded pest control for water allocation by the upstream *subak*. This is in line with empirical observations that upstream *subaks* are relatively more concerned about pests than water, the opposite of concerns of downstream *subaks* (Lansing and Miller, in press). However, their model is static and does not consider various cropping patterns, in contrast to my analysis. Including dynamics leads to a more specific understanding of how pest growth rates and dispersal rates affect coordination.

[Figure 7]

[Figure 8]

Imitation and the emergence of temple groups

The two-node model provides a deeper understanding in the coordination for water and against pests. So far, I have assumed that the nodes cooperated to derive the maximum total harvest or that there was a central control that forced a cooperative solution.

Although this is interesting from an analytical perspective, it does not provide insights into how *subaks* make their decisions in a more decentralized way without perfect control and information. Lansing and Kremer (1993) performed exercises where they allowed *subaks* to imitate the cropping pattern from the neighbor with the highest production. I used a similar approach but instead of defining a limited set of neighbors, I assumed that *subaks* had access to the cropping patterns of all other *subaks*. I defined networks for water and pest relationships and calculated the minimum number of connections it takes for each node to connect the other nodes in the network. The network for water includes both dams and *subaks* as nodes. I assumed that *subaks* that are more distant were 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 shown in equation (4) of when to imitate a cropping pattern. A cropping pattern is considered to be imitated if

$$H_i < \frac{H_j}{1 + \text{MIN}\{\gamma_p \cdot \chi_p^2, \gamma_w \cdot \chi_w^2\}}, \quad (4)$$

with parameters γ_p and γ_w and the number of connections separating two *subaks* via pest relationships χ_p and water relations χ_w . From all the *subaks* meeting this condition, the *subak* with the highest harvest per hectare will be imitated.

Starting with randomly distributed cropping patterns, *subaks* update their cropping patterns each year. A *subak* i compares the derived harvest per hectare with each other *subak* j , but only updates the cropping pattern when the condition in equation (4) 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 as worked out by Lansing and Kremer, who assumed a fixed set of neighbors. We also assume that there is opportunity for innovation. When a *subak* i performs worse than the average harvest per hectare within the watershed, it is assumed that there is a probability ρ that the cropping pattern will be changed to a random configuration.

By analyzing the performance of the system for different values of γ for water and pest transaction costs, we can analyze the different types of bottom-up coordination patterns that perform the best for different pest dynamics (Figure 9). Using a ρ equal to 0.04, we derive for the default model high harvest levels when γ_p and γ_w are positive, and γ_p is less than 0.5. This means that in the default case the best aggregated solution is derived when not all *subaks* are copying the best performing *subak* but focus on their own local area with a maximum of three connections through which pests in *subak* i can disperse. The difference between the effects of γ_p and γ_w relates to the fact that for coordination on pest outbreaks it is useful to synchronize the cropping patterns of a large

enough neighborhood and include all neighbors j . The inequality is low when cropping patterns are imitated within a large irrigation network (Figure 10).

[Figure 9]

[Figure 10]

When the growth rate of the pest is equal to 2, the results are very different. The results are not sensitive to the level of coordination when we use a default water availability that does not cause water scarcity (Figure 11). When a high pest growth rate is used, the best solutions are similar to those used with the default growth rate, except that when *subaks* look only at direct pest related *subaks*, it does not reduce harvests significantly.

[Figure 11]

There is a strong overlap between the 14 Masceti temples and *subaks* connected via pest relationships in the empirical dataset of the Lansing-Kremer model. This is consistent with my analysis that the pest dynamics in this particular watershed in Bali leads to synchronization of cropping plans among *subaks* who are connected via the spreading of pests.

Adaptive subaks

An alternative, plausible way *subaks* can make decisions on cropping patterns is to make decisions during the year whether to leave a field fallow or to plant a crop. I assumed that a *subak* made this decision based on the availability of water and dispersal of pests. A crop needs $150 \text{ m}^3/\text{day}$ of water per hectare. If water is expected to be above a certain threshold m_w , 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 .

In Figure 12, we see that if m_p is very low *subaks* never plant crops, leading to a low performance of the system. When m_p is large, crops are planted too early and pests are not controlled effectively. We also see that a larger value of m_w leads to a lower performance, since crops are not planted frequently when a high surplus of water is demanded before a *subak* starts planting a new crop.

Inequality among the *subaks* increases with a higher tolerance of pests (Figure 13). This leads to destruction of harvest for *subaks* that are prone to dispersal of pests from other *subaks*. A low growth rate of pests does not lead to a different relative performance of the best adaptive strategy when pest growth rates are at the default value (Figure 14). This suggests that adaptive strategies are also not sensitive to pest dynamics.

[Figure 12]

[Figure 13]

[Figure 14]

The effect of changed network structures

For both the imitative *subaks* and the adaptive *subaks*, I explored the consequences if the neighbors with whom they are connected by pest dispersal change. 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 the experiments of the imitative *subaks* I used a threshold of 0.4 for both pests and water, which led to the maximum harvest per hectare in the default case (Figure 9). For the adaptive *subaks*, I used 0.05 as a minimum for water and 0.02 as a maximum for pests, a combination that maximized the default case as well (Figure 12). In Figures 15 and 16 we see that the average harvest per hectare is sensitive to varying the probabilities. For the imitative *subaks*, the harvest decreases when pest-related links are removed. With fewer pest-related links, the *subaks* have no need for synchronization of crops. In fact, now water is the only reason for coordination, and synchronization with neighbors might be the worse thing to do. Therefore, a high probability of removing links has the lowest performance. The performance of imitative *subaks* is not sensitive to the probability of adding links.

The results for the adaptive *subaks* are very different (Figure 16). Adaptive *subaks* are not sensitive to removing existing pest-related connections. They are, however, sensitive to adding additional connections. With more links there is more need

for synchronization of cropping patterns, which is less necessary with individually adaptive *subaks*. Thus, although the adaptive and imitative *subaks* led to similar results for the original Bali irrigation network, they may lead to very different performance in a somewhat perturbed network.

[Figure 15]

[Figure 16]

5. Discussion

The Lansing and Kremer (1993) model is seminal and shows that simple bottom-up interactions of *subaks* can lead to good performance of a very complex large-scale irrigation system. But how much does it contribute to a more general understanding of the evolution of complex irrigation systems? In this discussion a number of problems of the original study are mentioned and some ideas for future work are sketched.

One of Lansing and Kremer's key topics is the level of social coordination. They analyzed different levels of synchronization that led to effective control of pests. However, synchronization is not the same as coordination. If water availability is the key problem, it is wise to coordinate with others not to use water at the same moment. Thus coordination might be formulated in very different cognitive and social processes. The Lansing-Kremer model is elegant in its simplicity of the agents, but if we want to explain coordination in irrigation systems in a more general fashion, we may need more comprehensive agents. For example, agents who learn to make a decision of when to

plant a crop based on water availability, pest biomass in the neighborhood, observed behavior of their neighboring agents, and expectations on water availability in the coming months. We may equip agents with neural networks to facilitate their learning in dynamic environments. When networks change, the agents may learn to adapt their rules about when to plant. We represent a group of farmers as an agent at the *subak* level. We may also consider “temple level” agents who collect information from various *subak* agents and give advice to the *subak* about when to plant crops. The *subak* agents may decide to follow the advice or may decide not to.

We expect that environmental conditions like connectiveness, pest dynamics, and rainfall variability may affect whether *subak* agents and temple agents are able to learn to cooperate, such that coordination processes may emerge. An important aspect of cooperation that is not included in the original Lansing-Kremer model is the maintenance of the canals. *Subaks* contribute to the maintenance of the irrigation system at the temple level. Lansing (in press) discusses conflict resolution between *subaks* when some do not contribute to maintenance. An increasing degree of sanctions could lead to cutting off a *subak*'s access to water. Maintenance of irrigation systems can be seen as a contribution to public goods. In future work, I may explore how different degrees of perceived costs and benefits of the public goods game affect the ability of a group of *subaks* to maintain the irrigation system.

If we have more adaptive agents that are able to coordinate activities in a complex dynamic irrigation system, we may explore the problem of the evolution of the irrigation network. Falvo (2000) argues that the Lansing-Kremer model cannot explain how the temple rituals have evolved, since pests were not an issue 1000 years ago. Falvo argues

that the irrigation system started in the coastal region and grew uphill. New *subaks* upstream were eager to cooperate with the large downstream population to avoid retaliation. The question is whether we are able to “grow” irrigation systems at different geographical landscapes that mimic observed regularities of various irrigation systems around the world.

6. Conclusions

The analysis in this paper confirms Lansing and Kremer’s (1993) insight that individual-level decision making can lead to a pattern of interactions that further lead to high performance of a whole irrigation system. The original Lansing-Kremer model is an important stepping stone toward understanding social coordination processes in complex dynamic irrigation systems. However, more comprehensive models of agent decision making, at various levels, need to be developed in order to generalize the insights, since the results are sensitive to specific assumptions of the ecology and the decision making of that area.

Acknowledgments

Thanks to Steve Lansing and Jim Kremer for providing the original source code of the Bali irrigation model and for their detailed comments on earlier versions of this paper. Marty Anderies provided helpful feedback on earlier versions of the paper, and Shade Shutters is acknowledged for his discovery of some critical errors in my original recoding of the Bali-irrigation model. The audience of seminars at Indiana University, University of Alaska Fairbanks, and Santa Fe Institute provided useful feedback. Finally, I gratefully

acknowledge support from the Center for the Study of Institutions, Population, and Environmental Change at Indiana University through National Science Foundation grants SBR9521918 and SES0083511. Indiana University Computing Systems kindly allowed me to run a portion of the simulations on the UITS Research SP System.

References

- Falvo, D.J., 2000. On modeling Balinese water temple networks as complex adaptive systems. *Human Ecology*, 28(4):641–649.
- Geertz, C., 1980. *Negara: The Theatre State in Nineteenth-Century Bali*. Princeton University Press, Princeton, NJ.
- Hauser-Schäublin, B., 2003. The precolonial Balinese state reconsidered: A critical evaluation of theory construction on the relationship between irrigation, the state, and ritual. *Current Anthropology*, 44(2):153–170.
- Hunt, R.C., 1988. Size and structure of authority in canal irrigation systems. *Journal of Anthropological Research*, 44(4):335–355
- Lansing, J.S., 1991. *Priests and Programmers: Technologies of Power in the Engineered Landscape of Bali*. Princeton University Press, Princeton, NJ.
- , in press. Tyrants, sorcerers and democrats. Chapter four in *Perfect Order*. Princeton University Press, Princeton, NJ.
- Lansing, J.S., and Kremer, J.N., 1993. Emergent properties of Balinese water temples. *American Anthropologist*, 95(1):97–114.
- Lansing, J.S., and Miller, J.H., in press. Cooperation in Balinese rice farming. *Current Anthropology*.
- Wittfogel, K.A., 1957. *Oriental Despotism: A Comparative Study of Total Power*. Yale University Press, New Haven, CT.

Figure captions:

Figure 1: Results of optimizations for 6 different synchronization levels from watershed (1) to individual *subaks* (172). Inequality is measured as harvest per ha. Harvest is the average harvest per ha per year.

Figure 2: Harvest (left) and inequality (right) levels for different rainfall scenarios for different coordination levels. The rainfall variability is the average for 10 runs, where for each year there is a 25% chance on a low rainfall year, 50% on a medium rainfall year and 25% on a high rainfall year. The synchronization level varies from the watershed (1) to the individual level.

Figure 3: Harvest (left) and inequality (right) levels for different growth rates of pests.

Figure 4: Harvest (left) and inequality (right) levels for different dispersal rates of pests.

Figure 5: Total harvest in two nodes for the different growth rates of pests and dispersal rates of pests between nodes. Black refers to two crops, dark gray to one crop, and light gray to no crops.

Figure 6: Total harvest for the different growth rates of pests and dispersal rates of pests between nodes when rainfall is 2.

Figure 7: Total harvest for the different growth rates of pests and dispersal rates of pests between nodes when rainfall is 1.5.

Figure 8: Total harvest for the different growth rates of pests and dispersal rates of pests between nodes when rainfall is 1.

Figure 9: Average harvest per *subak* per hectare when *subaks* imitate neighbors using different parameters. The irrigation system functions according to the default parameter settings.

Figure 10: Inequality of harvests per *subak* per ha when *subaks* imitate neighbors using different parameters. The irrigation system functions according to the default parameter settings.

Figure 11: Average harvest per *subak* per ha when *subaks* imitate neighbors using different parameters. A low pest growth rate of 2 is used.

Figure 12: Total harvest per *subak* per hectare when *subaks* adapt to methods of other *subaks* using different parameters. The irrigation system functions according to the default parameter settings.

Figure 13: Inequality among the *subaks* with adaptive *subaks* using different parameters. The irrigation system functions according to the default parameter settings.

Figure 14: Average harvest per *subak* per ha when *subaks* with adaptive *subaks* using different parameters. The irrigation system functions according to low growth rate of pests.

Figure 15: Average harvest per *subak* per hectare for different degrees of perturbation of pest relations between *subaks* when *subaks* are conditional imitators.

Figure 16: Average harvest per *subak* per ha for different degrees of perturbation of pest relations between *subaks* when *subaks* are conditional adaptive.




Figure 1

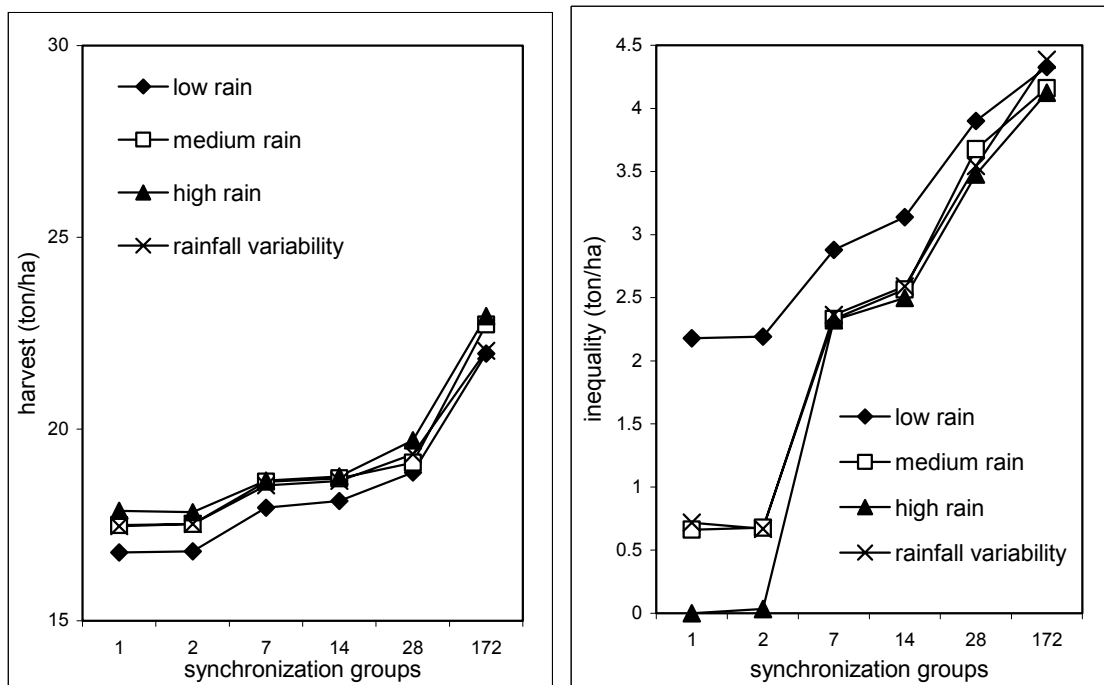


Figure 2

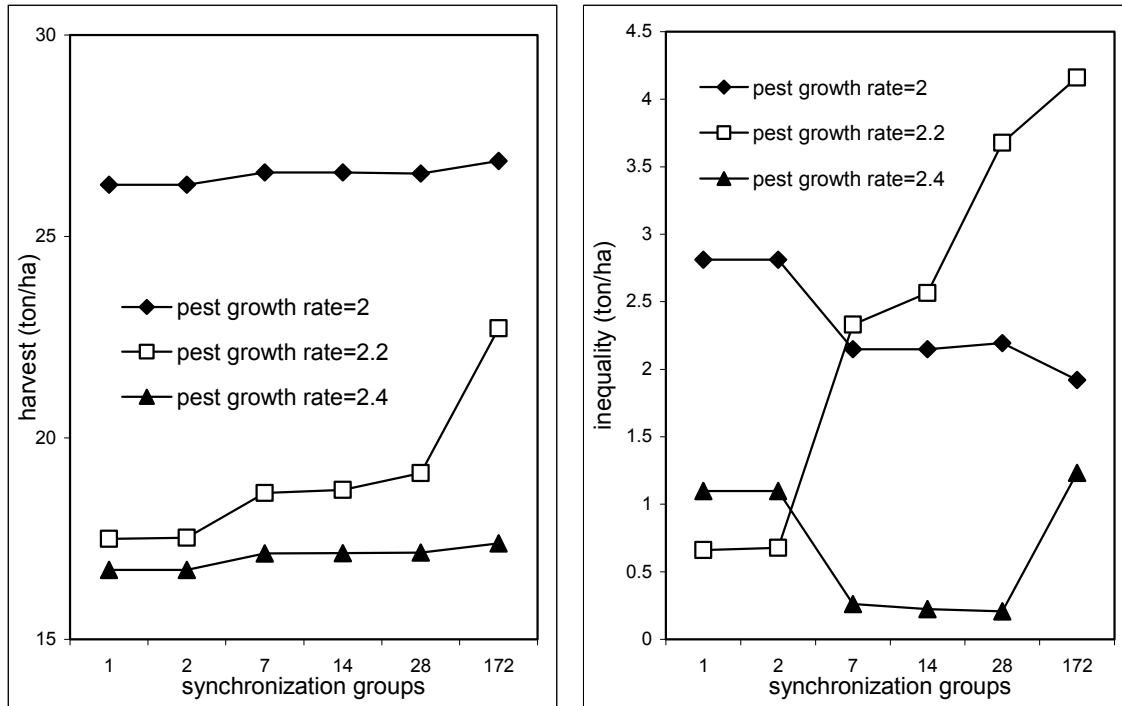


Figure 3

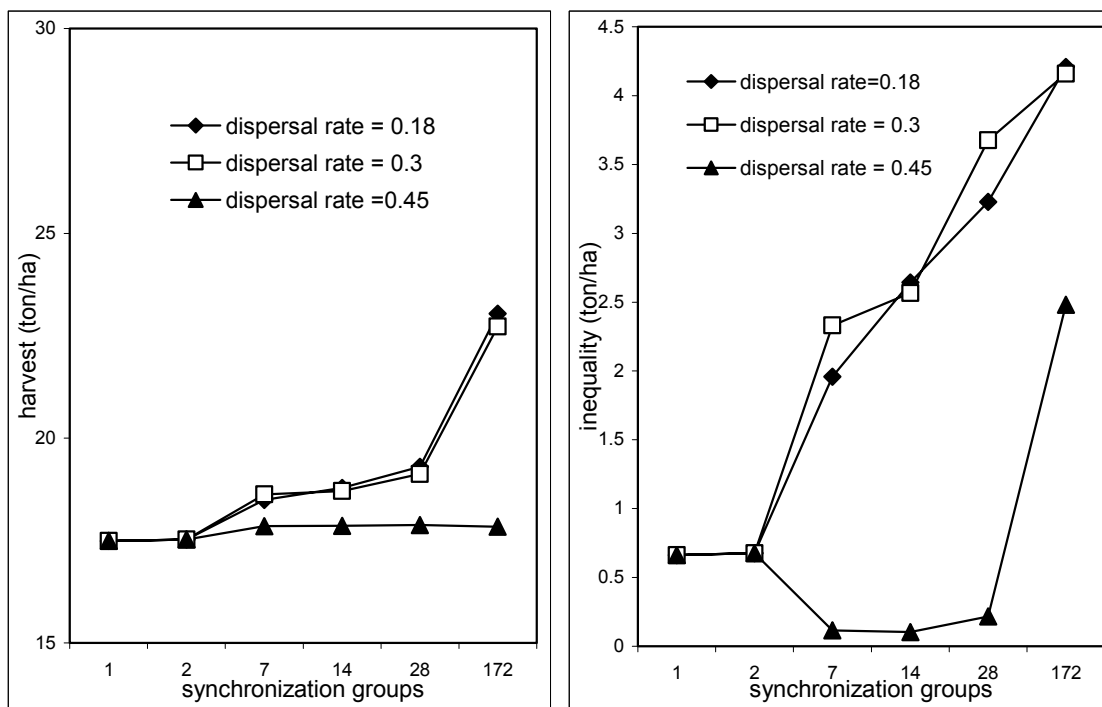


Figure 4




Figure 5

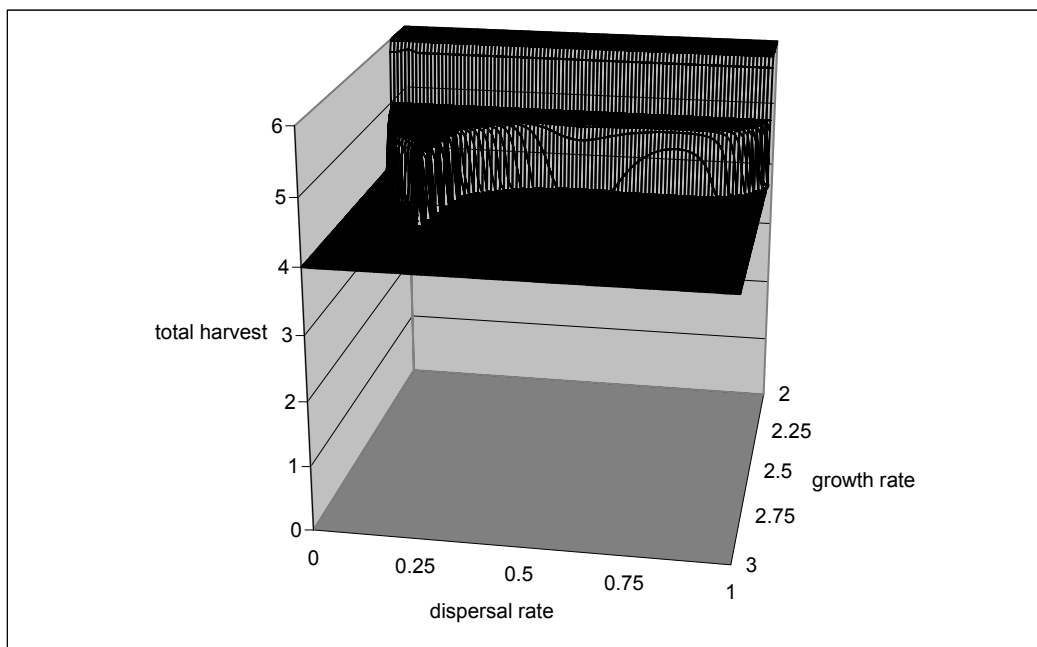


Figure 6




Figure 7




Figure 8

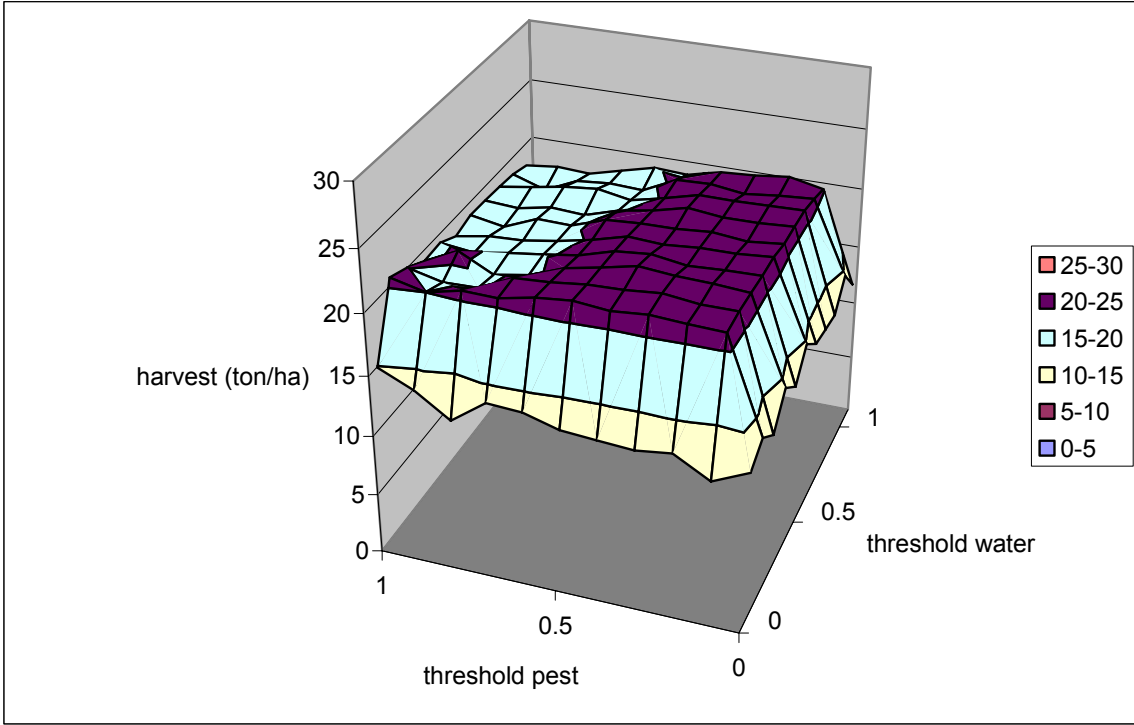


Figure 9




Figure 10

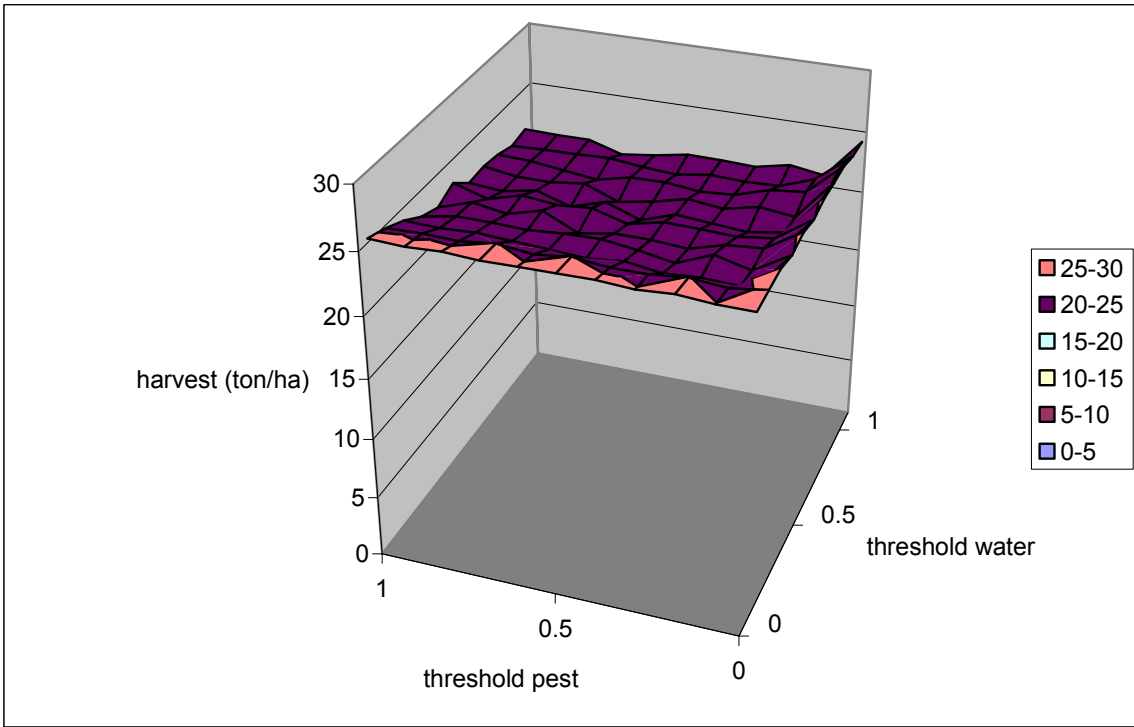


Figure 11




Figure 12

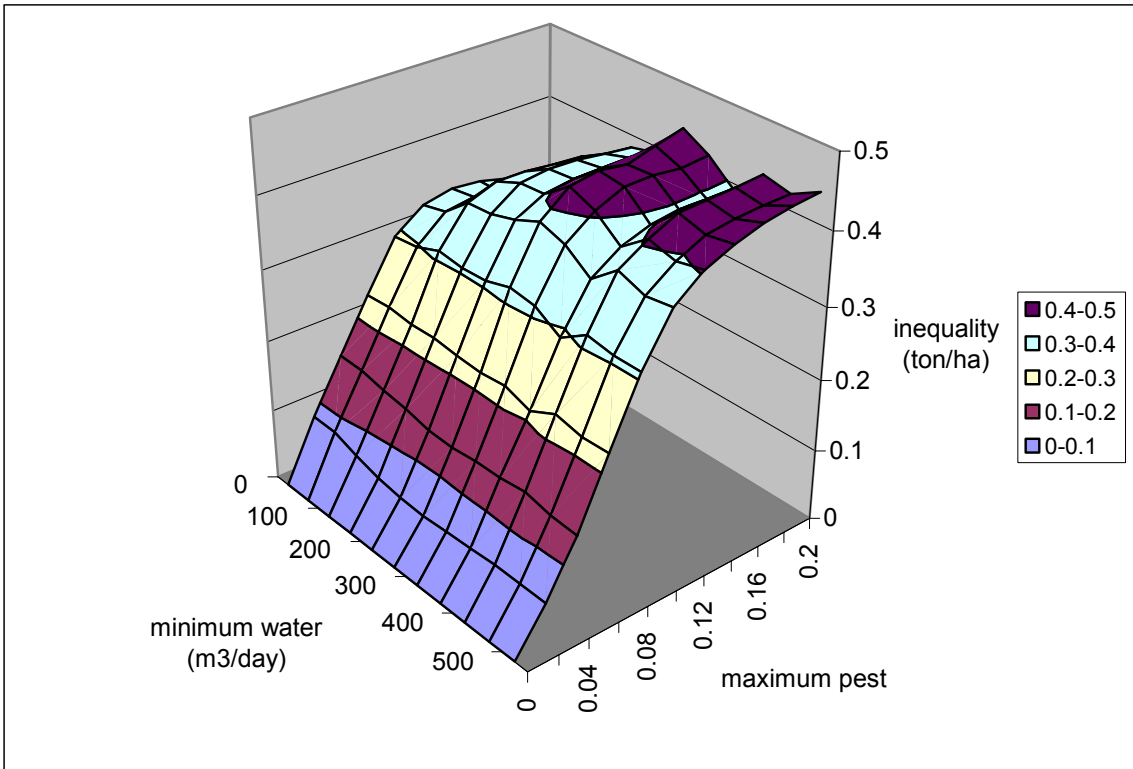


Figure 13

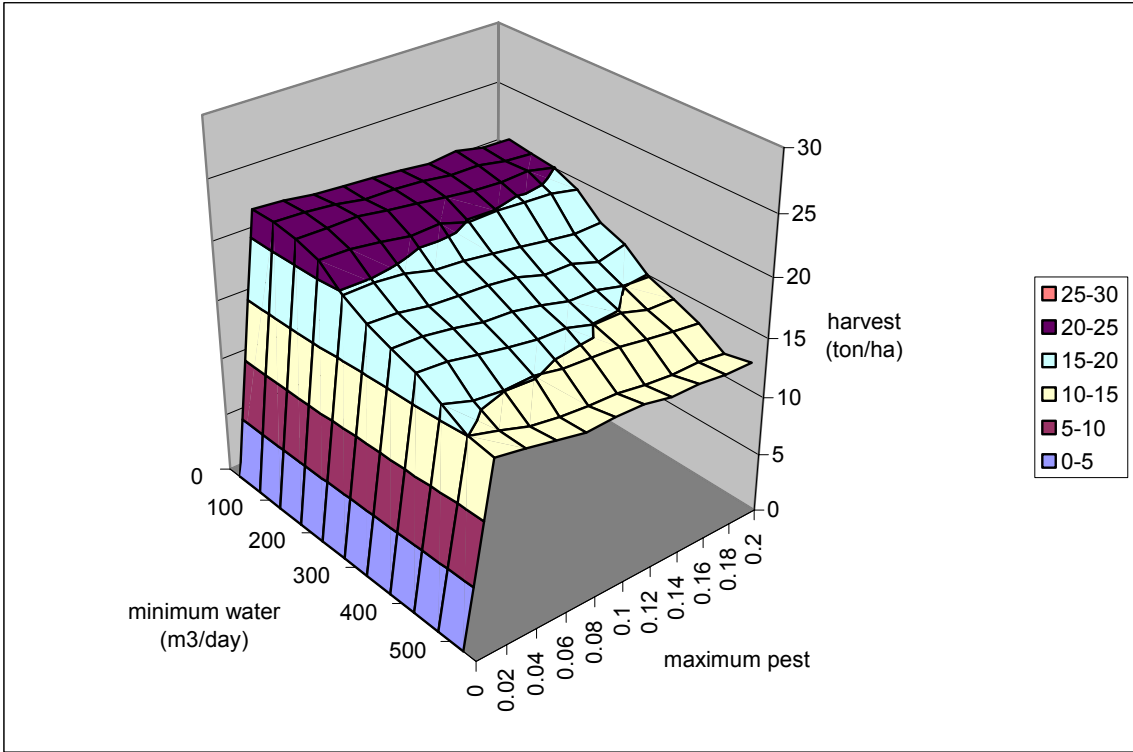


Figure 14

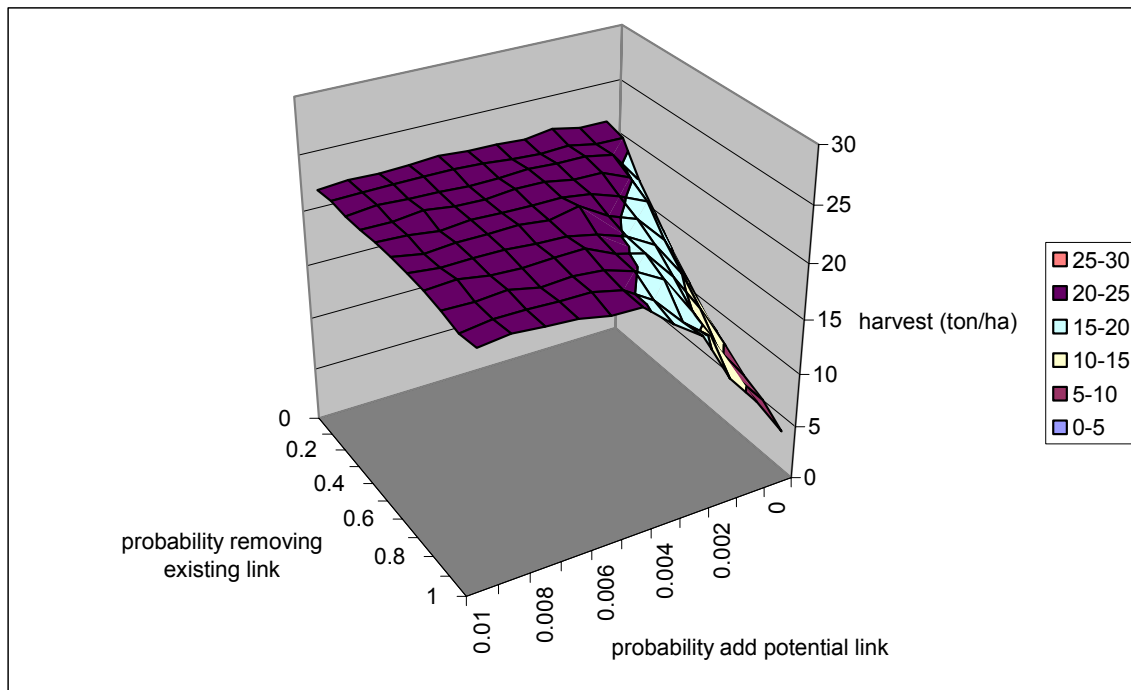


Figure 15

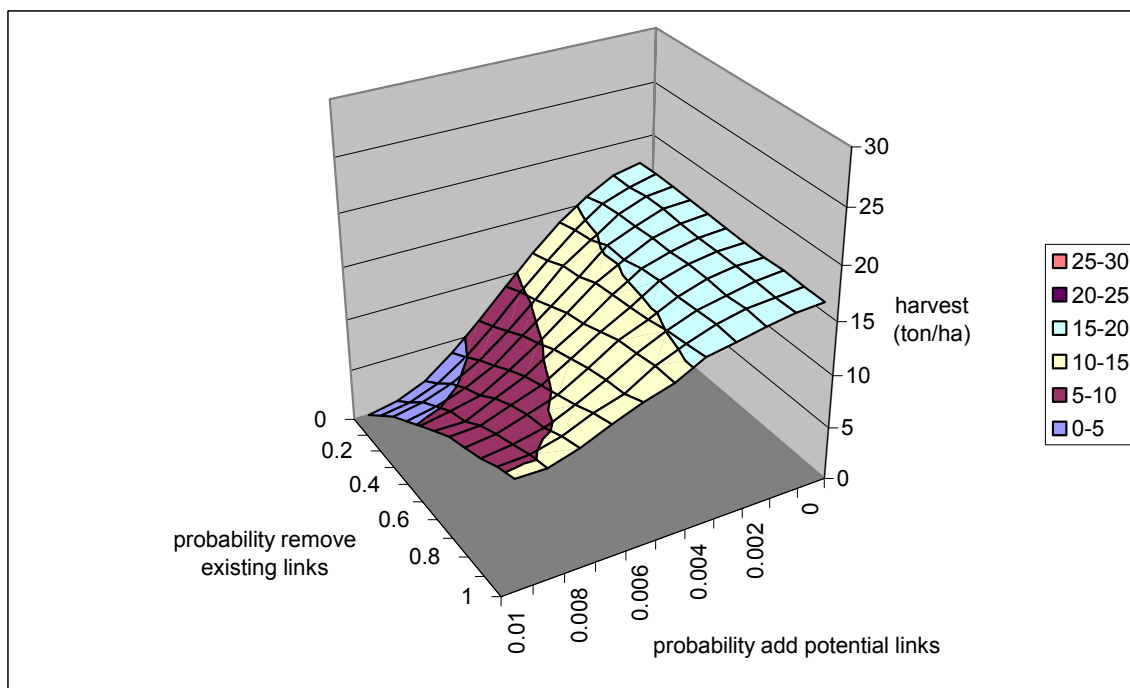


Figure 16