Journal of Environmental Informatics 21(2) 119-135 (2013)

Journal of Environmental Informatics

www.iseis.org/jei

Behavioural Adaptation towards Efficient Resource Sharing under the Lack of Communication

N. Barbalios[1,*], I. Ioannidou[2], P. Tzionas[3], and S. Paraskeuopoulos[1]

[1]Dept. of Special Education, University of Thessaly, Volos GR-38221, Greece
[2]Dept. of Early Education, Alexander Technological Educational Institute of Thessaloniki, Thessaloniki GR-54101, Greece
[3]Dept. of Automation, Alexander Technological Educational Institute of Thessaloniki, Thessaloniki GR-54101, Greece

Received 19 October 2011; revised 12 June 2012; accepted 24 September 2012; published online 27 June 2013

ABSTRACT. This paper introduces a novel multi-agent model for simulating water sharing scenarios under various irrigation policies, together with a novel self adaptive learning algorithm that achieves efficient resource allocation. The main contribution of this work lies in the fact that both the multi-agent model and the proposed learning algorithm operate under the lack of communication between the users of the resource, thus, no assumptions about the development of relations of trust between them are made. Moreover, the proposed learning algorithm uses only local information and operates in a decentralized manner, thus its implementation does not entail significant costs. The model was calibrated using data from a real world ecosystem and experimental results provided statistical and qualitative figures of merit for assessing typical irrigation policies. For all the irrigation policies examined, even if the users of the resource acted under profit maximization criteria, the proposed learning algorithm provided a means of achieving efficient resource allocation, despite the lack of communication. Thus, the proposed model and learning algorithm are valuable tools for assessing alternative irrigation policies and providing the best policy for any given scenario.

Keywords: tragedy of the commons, multi-agent simulation model, policy evaluation, natural resource sharing simulation

1. Introduction

Water sharing is a complex and difficult task that tends to become more and more irresolvable. Freshwater ecosystems gradually become scarce throughout the world and, at the same time, the earth's population is expected to increase by 50% in the next 40 years (Feek and Morry, 2003). In countries that already face water shortages, this increasing water crisis combined with inequitable access among the users of a resource, leads to confrontations, abuse and finally the depletion of the water resource, a condition referred to as 'the tragedy of the commons dilemma' (Hardin, 1968; Ostrom et al., 1994, 1999; Deadman, 1999).

Policies, consisting of specific set of rules, can be imposed in order to manage human activities so that harmful effects on natural resources would be prevented or reduced. Policies regarding water resource sharing usually describe rules under which users are entitled to exploit the water resource for irrigation purposes. Two main categories of policies can be found in the literature: centralized and decentralized. Centralized policies are enforced by central authorities (e.g. governments, institutions) and vary according to their complexity, implementation time, costs and the ecological and socio-economic impact entailed (Suleiman, 2005; Laycock, 2007; World Bank, 2007; Queensland Government, 2009; Queensland State, 2009). Empirical research findings however are not always in accordance with the recommendations of such policies (Ostrom, 2002). In more detail, centralized policies cannot be always implemented to their full extent and thus perform poorly. This is due to the costs involved, or to the absence of appropriate enforcement mechanisms, or even the lack of commitment by the users of the resource (Walker, 1989; Gurung, 2004). Further criticism on such policies is that they are based on erroneous assumptions, such as the belief that appropriate organization can be achieved only through centralized guidance (Ostrom, 2002).

Decentralized policies distribute decision making power over natural resource to local authorities. They are considered to outperform typical centralized policies, provided that control over the necessary financial and human resources is distributed and that local authorities are well organized (Smith, 1998). In the context of irrigation water sharing, irrigation schemes managed by the farmers themselves have in many cases outperformed expensive models proposed by governments (Shivakoti and Ostrom, 1992, 2002). Thus, a trend towards decentralized policies is reported in many countries (Rasmussen and Meinzen-Dick, 1995; Chemonics International Inc., 2004; Gurung, 2004). Decentralized policies however are not a panacea, since such local systems can also fail in many ways (Rasmussen and Meinzen-Dick, 1995; Ostrom,

* Corresponding author. Tel.: +30 2310 312003; fax: +30 2310 312003.
E-mail address: nbarba@autom.teithe.gr (N. Barbalios).

ISSN: 1726-2135 print/1684-8799 online
 doi:10.3808/jei.201300239

2002). Laboratory experiments have shown that when individuals are given the opportunity to exploit a resource, they start communicating with each other, they develop relations of trust, they exhibit self organization behaviours and they discover appropriate rules for efficient resource allocation (Ostrom, 1990; Ostrom et al., 1999; Ostrom, 2002). On the contrary, when such communication is lacking amongst the users of the resource, the tragedy of the commons is inevitably reached and the resource is destroyed (Ostrom, 1990; Ostrom et al., 1999; Ostrom, 2002).

It is thus clear that the interactions amongst the users of a resource and with their environment have a crucial impact on efficient resource allocation. Moreover, in such a decentralized management scheme, appropriate tools are required to stimulate joint learning and to integrate knowledge in order to establish shared understanding and coordination mechanisms, in the context of multiple resource users and their conflicting relationships (Gurung, 2004). To address these issues, multi-agent simulation (MAS) models were developed, that encapsulate the complex procedures of resource sharing and the human-environment interactions (Bousquet et al., 1998; Ginot and Page, 1998; Barreteau and Bousquet, 2000; Pauly et al., 2000; Becu et al., 2001). Multi-agent simulation is widely acknowledged as an appropriate modelling technique for simulating human-environment interactions.

A considerable review of MAS models and their applications to natural resource management can be found in Bousquet and Page (2004) and Matthews et al. (2007). In more detail, a MAS platform called CORMAS, specifically designed to simulate resource management by providing a framework to study various users with varying objectives, is presented in Bousquet et al. (1998). Utilizing CORMAS, the overexploited Kairouan water table in Tynisia is examined in Feuillette et al. (2003), where a model for the negotiation of management decisions about the water table is introduced. This is achieved by taking under consideration the global dynamics of the system and non-economic interactions between the farmers. In a similar context, CORMAS is used in Barnaud et al. (2008) to assist decision making over the allocation of rural credit in a small community of farmers in mountainous Northern Thailand. A multi-agent platform that outlines the consequences of water allocation rules is presented in Bars and Attonaty (2001), whereas in Schlüter and Pahl-Wostl (2007) a MAS model is presented for investigating the impact of the various organizational structures for water management to the water scarcity in the Amudarya river. The relation between water availability and water use is examined in Oel et al. (2010), where farmers modify land use and water use for irrigation based on empirical data gathered from surveys. A multi-layer MAS model is presented in Berger et al. (2007) that encapsulates collective action problems in water markets, whereas in Athanasiadis et al. (2005) a MAS model for water pricing is introduced that estimates water consumption taking into account economic and social parameters. Water management is also the case in Izquierdo et al. (2010), where the viability of irrigation in Senegal is examined, and a MAS model is used to examine the influence of existing social networks and the viability of irrigated systems. Although not utilizing agents, water management policies are also assessed in Tzionas et al. (2004), where a decision support system is developed to compare various policies proposed for the restoration of the water level and the rehabilitation of lake Koronia in Greece, based on their feasibility, environmental impact and costs.

Even though in MAS models the behaviour of the agents is usually empirically defined (Thebaud and Locatelli, 2001; Bousquet and Page, 2004; Schlüter and Pahl-Wostl, 2007; Oel et al., 2010), learning algorithms can be employed in the search of an optimal policy, i.e. a policy that achieves optimal resource allocation. In Alexandridis (2006), a MAS model for simulating the economics of land transition is presented where Bayesian learning is employed to adapt the agents' actions towards profit maximization. Janssen introduced a general learning method for MAS models that updates an agent's attraction to a specific strategy (Janssen and Ahn, 2006), whereas Yi-Chen et al. (2009) addresses optimal watershed management with agents objectives being optimized through a bargain scheme in which an agent announces his action to other agents.

The work presented in this paper introduces a novel MAS model that evaluates irrigation policies in cases where the users of the resource do not communicate with each other and thus they are not capable of self-organization or developing relationships of trust. Although humans have the tendency to communicate and trust each other, as demonstrated by the work of E. Ostrom in (Ostrom et al., 1999; Ostrom, 2002; Feek and Morry, 2003), a different approach is followed in this paper. We believe that as the economic crisis expands, farmers may be forced to take actions under extreme economic pressure, competing against each other, as water becomes scarce and could not sustain all the farmers of a community. When operating under pressure however, the self interest nature of individuals leads them to choices that improve immediate rewards, thus may lead them to self-lucrative behaviours (Sen et al., 1996). So, the assumptions of "good will" and "trust" existing amongst the farmers cannot be considered to hold for all cases of resource sharing. Thus, in the proposed work, it is chosen to investigate cases where although some communication between the farmers may exist, such information is not considered to affect their behaviour notably, and thus it is ignored. These assumptions distinguish the proposed model from all other MAS models proposed in the literature of natural resource sharing. Although MAS models applied in other disciplines (i.e. economics as argued in (Li, 2011)) may employ reduced inter-agent communication or partial knowledge of other agents utilities, the proposed MAS model does not assume any type of farmer communication (direct or indirect) at any part of its design; farmers are considered to act independently, consulting only their self interest and local knowledge. Additionally, since the lack of communication will most certainly lead to the tragedy of the commons (Ostrom, 2002), a novel multi-agent learning method that employs no communication amongst the farmers is introduced for the first time in this paper. The proposed learning algorithm operates as an additional layer of complexity to any given policy, and is capable of adapting the behaviour of the farmers towards efficient resource allocation, thus avoiding the tragedy of the commons. The main contribution of the work presented in this paper is:

- the introduction of a new MAS model for investigating water sharing schemes in the absence of communication amongst the users of a resource, capable of evaluating policies in scenarios where decisions are made with self-lucrative criteria and based on local information. As a result, the simulations that evaluate irrigation policies are not affected by unreliable assumptions, such as the assumption of "good will" or "trust". It should be also noted that the MAS model is directly related to (Tzionas et al., 2004), where various strategies were also evaluated, but only from an economical and environmental point of view, since the interaction with the farmers community was ignored.
- the introduction of a novel self-adaptive learning algorithm that achieves efficient resource allocation using exclusively local information, without employing any means of inter-agent communication. This is an advancement compared to typical multi-agent learning algorithms that mostly depend on communication or observation for coordinating the agents in order to reach their goals (Bųsoniu et al., 2008).

Extensive simulations were conducted, and experimental results have shown that the proposed model is capable of assessing irrigation policies based on their environmental and socio-economic impact. More specifically, the model was initially calibrated and its performance was verified based on empirical data derived from the ecosystem of lake Koronia. Representative irrigation policies were studied, and simulations were conducted in two-steps: initially with farmers employing only the policies under study, and subsequently with farmers adopting the learning algorithm introduced in this paper. The performance of each policy was evaluated using statistical and qualitative criteria, namely the condition of the lake at the end of the season and the survival rate of the farmer's community. These criteria were mapped on a new variable introduced in this paper, called community value, that provided a global means of assessing the performance of a policy. As it will be shown, the proposed MAS model and algorithm provide a means of adapting the farmers behaviour towards efficient resource allocation. Additionally, it will be shown that although no assumptions about trust or communication are necessary to be made, even if farmers act under profit maximization criteria, the introduction of the proposed learning algorithm will ensure both the sustainability of the resource and the maximization of their profit. In this sense, the proposed model and learning algorithm are valuable tools for assessing alternative irrigation policies and providing the best policy for any given simulation scenario.

2. Lake Koronia

Lake Koronia is presented in this section of the paper, since it will be used for the calibration of the proposed model, as it can serve as a representative example of the tragedy of the commons (Ioannidou et al., 2003; Tzionas et al., 2004).

Lake Koronia is one of the ten most important Ramsar protected wetlands, located 15 km northeast of the town of Thessaloniki in the region of Macedonia in Northern Greece, at a latitude of $40^{o}56'$ N and a longitude of $23^{o}15'$ E, and with a mean altitude of 75 m above sea level. During the 1970s, the lake was part of a sustainable environment, connected to the nearby lake Volvi that lies about 15 km to the downstream, it had a surface of 47 km^2 and a mean depth of 5 meters. In the last 20 years however, the lake suffered from a negative water balance, due to increased water diversion for irrigation purposes, that led to a dramatic water level decline, increased pollutant loads and reduced surface runoff. As a result, the area of the lake decreased dramatically to 30 km^2, the depth was decreased to 1 meter and the water quality deteriorated. These factors gave rise to the current hypertrophic conditions, which cannot support fish or any living organism. Particularly, the water level decline coupled with the water quality deterioration led to the death of a large number of fish and fish production was minimized in the summer of 1995 (Piesold et al., 1999; Ioannidou et al., 2003).

The aforementioned situation was additionally deteriorated by water-consuming irrigation schemes used by farmers of the area, coupled with water demanding agricultures, mainly consisted of corn. Another factor that further deteriorated the condition of the lake, was the failure of all central irrigation policies due to the lack of enforcement infrastructures (Tzionas et. al., 2004). A master plan was conducted, funded by the EC Directorate General XVI, Regional Policy and Cohesion Fund (Piesold et al., 1999), that suggested actions and measures to be taken in order to restore the water balance of the lake and ensure the viability of the overall ecosystem.

Regarding the restoration of the water resource, the lake water level and negative water balance was suggested to be restored through the transfer and artificial recharge of 45 million m^3 of water per year. Several restoration strategies were proposed to achieve that goal, including water diversion from nearby rivers, such as the Aksios river, Strymon river or Aliakmon river. Alternatively, the water required to restore the lake could be found either by diverting water from the Laggadiki and Scholari torrents, or by utilizing rainfall water from the village of Asvestochori. Restoration strategies also included water draining from the deep aquifer, or from the nearby lake Volvi. Finally, the master plan proposed the adoption of new water saving irrigation techniques and the limitation of water pumps and drills through a state regulatory policy.

Out of all the strategies that the master plan proposed, none was actually adopted, either due to controversial environmental impact, or due to implementation costs or even because some of these solutions cannot be considered sustainable unless accompanied by a drastic cut-off of agricultural activities across the lake (Kolokytha, 2010). Additionally, a fall in precipitation is observed that coupled with the high evapotranspiration will further diminish the natural inflows to the lake and will eventually lead to the reduction of the water reserves (Ministry of Rural Development and Floods, 2000). It is thus of great importance to examine the efficiency of irrigation policies using appropriate tools that not only capture the mathematical essence of the elements involved but also capture the socio-economical dynamics entailed in the complex water sharing procedure. The above findings support the motivation for developing novel water sharing models aiming at optimal resource allocation, such as the one proposed in this paper.

3. Proposed Model

In order to encompass the widest variety of irrigation schemes and deal with the imposed constraints, it was required for the proposed model and learning algorithm to be as modular, flexible and expandable as possible. Although several MAS platforms have been tested, such as Sesam (Klügl and Puppe, 1998), Cormas (Bousquet et al., 1998), RePast (Collier and North, 2011), NetLogo (http://ccl.northwestern.edu/netlogo/), MASON (Luke et al., 2004), none of them could meet these requirements adequately. Thus, the proposed model and learning algorithm were developed from scratch in the C++ programming language, which further to the flexibility and modularity also contributed to the encapsulation of novel parameters that were introduced in this paper.The model will be described using the ODD (Overview, Design Concepts and Details) standard protocol for describing agent based simulation models, presented in (Grimm et al., 2006).

3.1. Overview of the Model

Following the ODD protocol (Grimm et al., 2006) a general overview of the model is given, that provides a first insight about the model resolution, focus and complexity.

3.1.1. Purpose of the Model

This paper introduces a new MAS model that simulates the draining of a resource by a community of farmers, similar to the situation that arose in lake Koronia, and evaluates alternative irrigation policies. Additionally, a novel self-adaptive learning algorithm is introduced that is capable of achieving efficient resource allocation. The purpose of the model is two-fold: (a) it provides a simulation tool that reveals underlying behaviours of the farmers as well as the environmental and socio-economical impact of water management policies; the main assumption here is that communication amongst the users of the resource does not affect or alter their behaviour considerably; (b) it reveals that, provided the proposed learning algorithm is applied, behavioural self adaptation based solely on the maximization of profit (and not on communication skills or even bonds of trust developed amongst the farmers as suggested in (Ostrom, 2002; Ostrom et al., 1999; Ostrom, 1990)) can also lead to coordination towards the mutual benefit of the environment and the farmers community.

3.1.2. Variables and Scales

The model entails two main entities: a) the water resource, described by a hydrological model, that defines its water balance over a period of time, and b) the users of the resource that represent farmers draining water for irrigation purposes only. Users of the resource are considered only to drain water for irrigation purposes, since all other human related water consumptions from the resource (i.e. industrial needs, urban needs etc.) are considered insignificant (Kolokytha, 2010) and are incorporated within the hydrological model. The morphology of the ground is implicitly taken under consideration as well, by limiting the amount of water actually drained by the farmers to a certain degree when the resource water level reaches a certain threshold T. This can be realized by considering that below that certain threshold T, the resource degrades to a series of disconnected ponds, resulting to loss of water pressure, exposed or un-submerged water draining pipes etc. In general, the resource can be described by its initial water level X, the threshold level T, the amount of water provided under the threshold level W_T, and by the hydrological model employed.

The population of the users of the resource consists of farmers that drain water from the resource utilizing the same irrigation policy (i.e. behaviour) π_i. Farmers are characterized by their goal G_i, which is a variable defining the total amount of water they wish to drain throughout the season to cover their needs. These needs should ideally be derived from the type of cultivation and the field area. Unfortunately, however, this is not always the case since many farmers over-drain water either for self lucrative reasons (by over-irrigating without knowing which is the appropriate amount of water for their type of cultivation) or in order to compensate for water losses during water transfer from the lake to their field (ignoring the fact that all the excessive water is wasted, at least for the specific cultivation period). To that extent a 'greediness' variable g is introduced, to account for all these issues and in order to differentiate between the amounts of water each farmer drains. Study of research reports of the lake Koronia ecosystem (Tsiouris et al., 2002; Tzionas et al., 2004), revealed that five distinct categories should suffice to depict the different behaviours in the community of the farmers, with respect to their degree of greediness. As it will be shown in Section 3.3 (Initialization/Calibration), this categorization is adequate to encapsulate the various farmer behaviours of the nearby area of lake Koronia. The full set of values for the greediness variable is g = 1, 2, 4, 6 and 10, which practically means that farmers are clustered in 5 greediness categories; for the purposes of this paper a farmer of greediness category 4 is considered to drain 4 times more water than a farmer of greediness 1, and so on.

Additionally, a learning method is introduced within the model that based on machine learning principles, provides a means of modifying the behaviour of each farmer towards efficient resource allocation. In that sense, water otherwise wasted by greedy farmers over-irrigating their cultivation, is re-distributed to non-greedy farmers that need it in order to meet their goals sufficiently.

3.1.3. Process Overview and Scheduling

The model advances in daily time steps, denoted by the index t throughout this paper, and each simulation is considered to be concluded in 160 days that nearly corresponds to a cropping season of the year. During each time step the following procedure takes place:

- Each farmer takes an action i.e. requests a certain quantity of water from the lake. The action is selected out of a set of predefined available actions (i.e. water quantity requests) that are the same for all the farmers of the same greediness category.
- Each farmer receives the appropriate amount of water from the lake. This amount of water drained is either equal to the quantity requested, if the water level of the resource is above the predefined threshold T or equal to a predefined quantity WT, if the resource is nearly dried out.
- Each farmer perceives the condition of the resource. Due to the constraints imposed by the model (i.e. lack of communication and observation), this is done with a novel feedback signal that utilizes only local information.
- Based on the irrigation policy π_i a farmer employs (i.e. his behaviour), he takes a decision about what action to take the next day, i.e. what quantity of water to drain the following day in order to reach his goal. Practically, the irrigation policy adopted by the farmer defines the appropriate action to be taken, based on some criteria (e.g. profit, environment sustainability, environmental awareness and others) which vary for different policies π.
- The procedure is repeated iteratively until the end of the experimental period, i.e. for 160 days. It should be noted that this is the time window of interest that is examined by the model. In this sense, if the lake is not completely drained during this period, it is not considered to be depleted.

3.2. Design Concepts

Following the ODD design (Grimm et al., 2006), the general concepts that describe the underlying design of the model are given in this subsection.

3.2.1. Emergence

Several behaviours related with the environmental awareness of the population and the sustainability of the ecosystem have emerged through the model simulations. These behaviours can be depicted by using specific figures of merit. Such a figure is the community value V_C, introduced in this paper as an index that characterizes the overall performance of a policy with respect both to the sustainability of the resource and to the economic survival of the farmers population.

3.2.2. Adaptation

The behaviour of the farmers is adapted to two hierarchical levels: On the first hierarchical level, farmers adapt their behaviour based on their reasoning process, as defined by the adopted irrigation policy. Such an irrigation policy defines the appropriate action, i.e. the quantity of water, a farmer should drain. Three such policies are implemented in the current version of the model: a) a non-rational policy where farmers decisions are solely based on personal beliefs about their actual water needs. In this case, any available feedback from the environment is ignored. b) a profit driven policy that maximizes a multivariate utility function whose parameters are configured accordingly in order to maximize profits, similarly to utility functions found in the literature (Monticino et al., 2007; Brown et al., 2008; Lin et al., 2008) and c) an environmental friendly policy, which can be derived by an appropriate parameter reconfiguration of the multivariate utility function utilized in b, so as the sustainability of the resource is maximized.

On the second hierarchical level, the quantity of water to be drained by the adopted irrigation policy can be further adjusted towards efficient resource allocation, by utilizing the learning algorithm introduced in this paper.

3.2.3. Sensing

The main novelty of the model introduced in this paper lies in the lack of any communication between the farmers. Although it is widely accepted that humans communicate, develop relations of trust and discover rules of efficient resource allocation (Ostrom, 1990; Ostrom et al., 1999; Ostrom, 2002), it is more than certain, in our opinion, that the economic crisis will force farmers to take actions under extreme pressure, leading to self-lucrative behaviours (Sen et al., 1996). Thus, we consider that any communication or bonds of trust developed between the farmers cannot be considered reliable and that information is ignored. However, farmers are capable of making some implicit observations of the water level that are used by the self adaptive learning algorithm introduced in this paper. To that direction, an appropriate feedback signal has been defined in this paper that uses only local information, which can be used by farmers to perceive changes in the condition of the lake.

3.2.4. Interaction

Due to the imposed constraints, farmers can only interact with the water resource and not with each other.

3.3. Details of the Model (Sub-models)

3.3.1. Hydrological Model of the Lake

The hydrological model of the lake was developed based on the water balance over a period of time, as it is analytically described in (Katirtzoglou, 2001; Mylopoulos et al., 2007; Kolokytha, 2010). Literature is rather vague on this issue, since parameters such as the hydraulic communication between the shallow and the deep aquifer as well as the lake-aquifer interaction are not fully clarified (Mylopoulos et al., 2007). For the purposes of this paper, data and measurements presented in (Katirtzoglou, 2001; Mylopoulos et al., 2007; Kolokytha, 2010) were utilized for the estimation of the water balance equation, presented in Table 1. To our knowledge, these are the most accurate and official measurements of lake Koronia inflows/outflows. It should be noted that, outflows due to irrigation are purposefully excluded from the estimation of the water balance equation, since this is the exact procedure simulated in the proposed model by the farmers' actions. As a result, using the formula $100 \cdot (\text{total outflows/total inflows})$, a water balance is approximated were the total amount of outflows equals 78.2% of the total inflows.

Table 1. Inflows and Outflows of Lake Koronia

Inflows ($\times 10^6$ m^3 water)		Outflows ($\times 10^6$ m^3 water)	
Surface Water	25.3	Evapotranspiration	107.1
Rainfall	147.6	Lake evaporation	30.0
Groundwater	22.0	Groundwater outflow to Volvi	5.3
		Urban Needs	2.1
		Industrial needs	4.0
		Outflow to Scholarion aquifer	4.0
Total Inflows	194.9	Total Outflows	152.5

3.3.2. Actions Performed by the Farmers

Let A be the set of available actions, defined as

$$A = \begin{cases} \text{action}_1 = \dfrac{a^*}{6} \\ \text{action}_2 = \dfrac{a^*}{2} \\ \text{action}_3 = a^* = g \cdot \dfrac{G_i}{160} \\ \text{action}_4 = 4 \cdot a^* \\ \text{action}_5 = 6 \cdot a^* \end{cases} \tag{1}$$

where a^* describes the nominal action i.e. the amount of water that agent ideally requires to cover his irrigation needs, G_i is the goal of farmer i, i.e. the amount of water each farmer wishes to drain in 160 days, and g is the greediness variable used to create heterogeneous categories of agents that have different needs, require different water quantities, have a different action space and different goals (as described in Section 3.1). The action selected by farmer i on day t is denoted by $a_{i,t}$.

3.3.3. Farmers' Perception of the Water Level

Since farmers do not communicate with each other, information regarding the condition of the resource should stem from their interactions with it. The only available information to the farmer i is the water actually drained on day t, denoted by $r_{i,t}$, and given by the formula:

$$r_{i,t} = \begin{cases} a_{i,t} \text{ , if water level >T} \\ W_T \text{ , if water level <T} \end{cases} \tag{2}$$

And the overall amount of water each farmer has drawn so far. To compensate for the lack of communication and assist farmers perceive the state of the resource using only local knowledge, a novel feedback signal $f_{i,t}$, is introduced in a multi-agent model for the first time in this paper, defined as:

$$f_{i,t} = \frac{a_{i,t}}{p} \tag{3}$$

where $a_{i,t}$ the action selected by farmer i on day t and $p = l / X$ is a decaying parameter, involving the lake initial level X and current level l. Parameter p actually modifies the farmers reward according to the current lake level, thus implicitly providing knowledge about the current state of the resource. The parameter p is an indirect measure of the water level and could be related to other available physical quantities such as any visual feedback a farmer can have to get a hint about the water level (which is commonly the case) or even, the electrical current drawn by a water draining pump, that should typically increase as the water resource gets empty, etc.

A farmer i perceives the rate of change of the resource $c_{i,t}$ at day t, according to Equation 4, that provides an estimate about the rate it is drained rather than about the actual water level of the resource. Thus p does not need to be accurately estimated (as it could not be due to its vague nature):

$$c_{i,t} = \frac{d(f_{i,t})}{dt} \tag{4}$$

3.3.4. Irrigation Policies of the Farmers

For the purposes of this paper, farmers are considered to choose actions based on one of the following available behaviours (i.e. irrigation policies):

- *Non-rational policy (NR)*: This is the simplest form of reasoning a farmer can employ. No adaptation takes place and a farmer always drains the same amount of water considered to be adequate to cover his needs, despite the condition the resource is in. Under the non-rational policy, the farmers' behaviour is driven exclusively by profit in a self-lucrative manner. At any time, the farmer always chooses the actions that maximize their immediate rewards. This policy represents an extreme example of self-lucrative behaviour and, to our knowledge, it resembles the actual behaviour of farmers in the ecosystem of lake Koronia in Greece that lead to its depletion. (Tsiouris et al., 2002; Tzionas et al., 2004; Laycock, 2007).
- *Profit driven policy (PD)*: Under this policy, actions that maximize a specific utility function are selected, similarly to (Monticino et al., 2007; Brown et al., 2008; Lin et al., 2008) and the decision support systems presented in (Bazzani, 2003; Shajari et al., 2008). Utility functions are usually defined either from experience, surveys or focus group sessions, as is the case in (Monticino et al., 2007). In this paper, a multivariate utility function is introduced, of the form $U = K_W \cdot U_W + K_P \cdot U_P$, similarly to (Monti-

cino et al., 2007). It is consisted of two partial utility functions U_W and U_P modeling water preservation and profit, respectively. Regarding U_W, actions resulting to a low rate of draining of the resource should be mapped to higher values than actions that lead to a high rate of draining of the resource. Thus, the marginal utility function for water preservation was defined as a decreasing exponential function of the rate of change of the resource $c_{i,t}$, as defined in Equation 4. The marginal utility function for profit U_P was defined as an increasing exponential function of a farmers' demand $a_{i,t}$, similarly to (Lin et al., 2008). The total utility function takes the following form:

$$U(a,c) = k_W \cdot \left(1 - e^{-RC}\right) + k_P \cdot \frac{a^{1-R}}{1-R} \qquad (5)$$

where R serves as a parameter defining the slopes of the partial utility functions (set to the value of 6 for the purposes of this paper, similarly to Lin et al., (2008)). The weighting parameters k_W, k_P indicate the relative value that a farmer places on water preservation and profit, respectively. Setting $k_P > k_W$, as is the case in this policy where $k_P = 0.6$ and $k_W = 0.4$ similarly to Monticino et al. (2007), represents a policy where farmers are primarily interested in profit maximization.

- *Environmental friendly policy (EF):* This policy is a variation of the profit driven policy, as it is derived by a different configuration of the same multivariate utility function. In more detail, the weighting parameters are configured so that $k_W > k_P$, (i.e. $k_W = 0.6$ and $k_P = 0.4$ similarly to Monticino et al. (2007)) thus representing a policy where farmers are interested in water preservation rather than profit.

3.4.5. Initialization/Calibration

The model proposed in this paper was calibrated with respect to data and observations from the lake Koronia ecosystem. The information available about the ecosystem consists of a) the lake inflows/outflows over a short period of time (Katirtzoglou, 2001; Mylopoulos et al., 2007; Kolokytha, 2010), b) a dataset of water losses during the period 1988 to 1995 with respect to the increase of agricultural activities during that same period (Piesold et al., 1999; Hellenic Ministry of Agriculture, 2001; Ioannidou et al., 2003), c) observations that where derived by studying research reports of the lake Koronia ecosystem (Piesold et al., 1999; Hellenic Ministry of Agriculture, 2001; Ioannidou et al., 2003; Tzionas et al., 2004). It should be noted that the dataset of water losses was divided in a train set that was used for calibration (50% of the data), and a test set that was used for validation (remaining 50% of the data). Moreover, for calibration purposes, farmers considered to employ the non-rational policy, since this seems to correspond to their actual behaviour that lead to the explosion of agricultural activity and eventually dried out the lake (Tsiouris et al., 2002; Laycock, 2007). Additionally, it should be noted that water level values and thresholds were normalized in order to extract some quantitative results

The hydrological model was manually calibrated, as described in Section 3.3, and the greediness categories (i.e., g = 1, 2, 4, 6, 10) were derived after studying (Tsiouris et. al., 2002; Tzionas et. al., 2004). The remaining parameters were calibrated in a two-step procedure: During the first step, observations from (Piesold et al., 1999; Hellenic Ministry of Agriculture, 2001; Ioannidou et al., 2003) were employed to provide initial estimates of the model parameters. Subsequently, to improve the accuracy of the model, the train set of the water loses dataset was employed to fine-tune these initial estimates. The random calibration method was employed (Jaffe et al., 1988): the initial estimates were slightly altered in random, so that model – predicted water losses match the actual water losses observed in the lake. Farmers were distributed in various greediness categories, corresponding to a similar increase of agricultural activities and goodness of fit was measured using the mean absolute deviation between the model output and the data of the train set. The final model parameters after the calibration procedure are presented in Table 2.

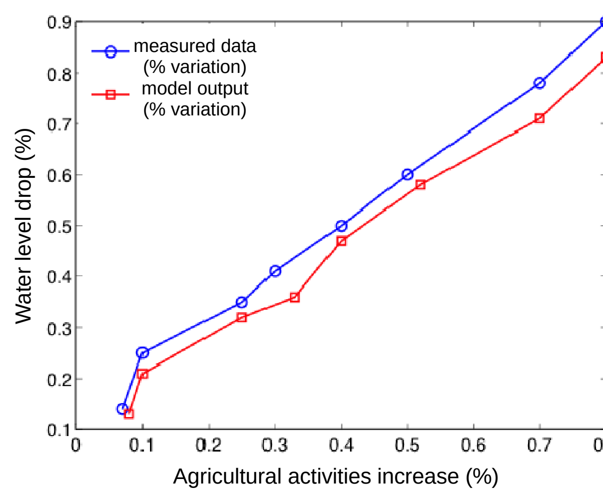

Figure 1. Percentage variations of water losses of lake Koronia with respect to percentage variations of agricultural activities in the area, for the time period of 1988 to 1995. The two data lines refer to the observed data and the output of the proposed model after calibration, based on observations from lake Koronia.

The model was validated after the calibration with the test set, following the same procedure, i.e. by redistributing farmers in various greediness categories, so that similar increase of agricultural activities was achieved. The results are illustrated in Figure 1, where the x-axis refers to the percenttage variations of agricultural activities and the y axis to the percentage variations of water losses. It is clearly demonstrated that the output of the model matches the measured data, as demonstrated by the two plots of Figure 1. In more detail, median absolute deviation between measured data (from the lake) and the model output is 0.015.

Thus, it is verified that the calibration of the model with respect to lake Koronia was satisfactory. Additionally, since the model predicted water losses match the actual water losses, it is verified that the values for the model parameters that were excluded from the calibration procedure (i.e. the hydrological model and the classification of farmers to five greediness categories) are sufficient.

3.4. Self-Adaptive Learning Algorithm

Learning in such dynamic settings, where all agents learn simultaneously without communicating with each other, is a difficult task. The lack of communication is an inhibitory factor for the coordination of the agents (farmers), as most multi-agent learning algorithms in the literature rely on observation or communication (Bușoniu et al., 2008). Moreover, measurement uncertainties and noise are almost certain to occur, due to the large number of agents (farmers) and the continuous nature of the problem under study. Furthermore, even though most multi-agent learning algorithms converge to optimal policies that focus on the maximization of the agent's utilities, a different approach should be followed. Due to the nature of the problem under study, it is critical that any viable solution should not only focus on the agent's utilities but on the preservation of the water resource as well. Such solutions however are difficult to be found, since they cannot be easily mapped to a unique equilibrium (Mannor and Shamma, 2007). In this sense, most existing multi-agent algorithms proposed in the literature are inappropriate to solve the problem under study. It is thus made clear that learning in settings such as the one entailed in this work is a hard task.

To achieve learning in such a constrained setting, we introduce a learning algorithm that modifies the actions prescribed by any policy, towards efficient resource allocation. This is achieved by adding an additional layer of complexity in the decision making process of the agents (i.e. their policy), which incorporates specific rules that are based on empirical findings of the decision support system for lake Koronia presented in (Tzionas et al., 2004). In more detail, the action $a_{i,t}$ of farmer i in day t is modified by the following factor:

$$\text{modFactor}_{i,t} = \left(\frac{1}{\left(\sum_{m=0}^{t} \gamma \cdot f_{i,m} \right) \cdot \left(\sum_{m=0}^{t} r_{i,m} \right) \cdot \left(c_{i,t} \right)} \right)^{Q} \tag{6}$$

where $\gamma = (1-0.95^{m})/(1-0.95^{t})$ is a discount factor commonly met in reinforcement learning algorithm used to weight the rewards a farmer has received in the past (latest rewards weight higher than old ones to allow the function to converge to an optimal solution), $\sum_{m=0}^{t} r_{i,m}$ is the actual sum of water drained by farmer i until day t (see Equation 2) , $c_{i,t}$ is the rate of change of the resource (see Equation 4) and Q is a regulating exponential factor, controlling the impact the overall modification factor has to the farmer's initial water request (similar to the learning rate of machine learning algorithms). To estimate the optimal value for Q a series of extensive tests were carried out, validating Q values over the range 0.01 to 0.5. This range refers to relative small Q values, since large learning rates may lead to steep state variations that would cause the instability of the learning algorithm. It should be noted however that the precise and analytical determination of Q is part of our future work. Out of the tested Q values, 0.2 gave the best results and that value was used throughout the experiments, when the self-adaptive learning algorithm is employed.

Table 2. Model Parameter Values Based on Data and Observations from Lake Koronia

Parameter	Value
Lake initial level (X)	10000 units
Lake level threshold (T)	2000 units
Water provided under threshold (W_T)	12.5 units
Number of agents (N)	50
Goal of agent i (G_i)	160 units
Number of greediness categories	5 (g = 1, 2, 4, 6, 10)
Water balance equation (irrigation outflows excluded)	Outflows = (0.78)· inflows

A closer examination of Equation 6 reveals that each farmer's policy is modified at a predefined learning rate, utilizing: a) local knowledge (as depicted in the amount of water gathered), b) the perceived rate of change of the resource, and c) experience, represented by the feedback accumulated over time by the farmer, according to machine learning principles (Bishop, 2007). As a result, the overall quantity of the water drained from the lake is modified.

In more detail, the terms of Equation 6 can be justified as follows:

- Since the rate of change of the resource is crucial for sustainability, farmers should modify their requests in a manner inverse proportional to it, as encapsulated in the term $1/c_{i,t}$.
- As farmers accumulate more and more water, they should limit their requests, so that water would suffice for others as well. This rule is depicted in term $1/\sum_{m=0}^{t} r_{i,m}$.
- Considering that the feedback signal $f_{i,t}$ increases as the water level of the resource decreases in time, term $1/\sum_{m=0}^{t} r \cdot f_{i,m}$ adapts the farmers behaviour so that the resource is exploited in a more conservative manner.

Although a farmer's behaviour is initially defined by the irrigation policy employed, Equation 6 allows adaptation of his behaviour in a self-regulated manner. The amount of water $a'_{i,t}$ a farmer will request under the proposed self adaptive learning algorithm will be:

$$a'_{i,t} = \text{modFactor}_{i,t} \cdot a_{i,t} \tag{7}$$

It should be noted that the proposed algorithm does not modify the greediness degree of each farmer and, subsequently, the overall greediness degree of the population. According to his greediness degree, each farmer self-adjusts the quantity of water he drains through Equation 7. As it will be shown in the results section, the terms employed in Equation 6 and discussed above, are adequate for simultaneously improving both the viability of the resource as well as the economic survival of the farmers population.

As far as the practical implementation of the proposed learning algorithm is concerned, Equation 6 and Equation 7 can be easily implemented by a computerized control system that could be installed in every field. Any computer enhanced with an input/ output interface could serve as such a computerized system, that would:

- gather measurements from the water pumps in order to keep track of all the required variables of Equation 6 (i.e. the change of rate of the resource, water drained so far, etc). Measuring such variables is quite simple given today's sensoring devices. Also, the cost of such sensoring devices is insignificant.
- control the amount of water drained by the pumps, by implementing Equation 7. Given the advancement of today's computer systems, any computer/laptop could be modified with appropriate input/output devices, to control the pumps through Equation 7.

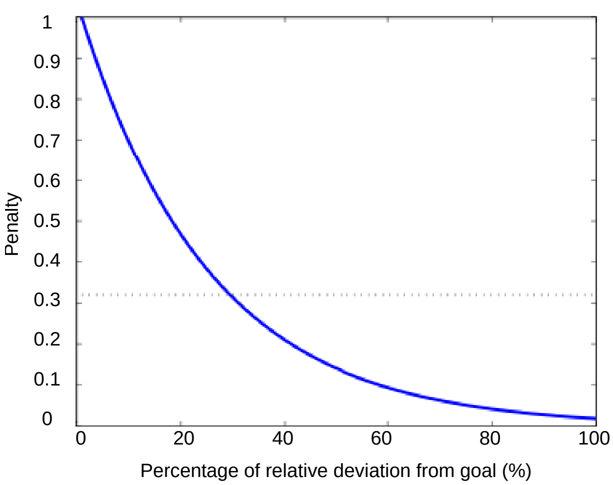

Figure 2. Penalty function used to evaluate the percentage of relative deviation from the optimal goal. The threshold line denotes the value under which the decay diminishes rapidly, resulting to a severe penalty.

Due to the design of the learning algorithm, since learning is based only on local information, no centralized computer is required to organize the farmers, nor do the computers have to communicate with each other. Thus, the need of a global regulator is avoided. As a result, costs are further reduced, as in any other case a computer network or an internet conection would be required, which is a highly inhibitory factor for most agricultural landscapes. In this sense, the implementtation of such a computerized control system is not only simple (given the simplicity of Equation 6 and Equation 7 and the estimation of its variables), but also affordable, as any computer can be used for the implementation and no network is required.

Although incentives provided by government or legislation would be more than welcome, the use of the proposed system does not require government intervention or farmer training, as it would automatically control the pumps to autoregulate the amount of water drained. Also, as previously mentioned, no significant costs are involved. Considering the effectiveness of the solution, that will be presented in the experimental results, we argue that there is no need of special motivation for the farmers to adopt this solution as a) it does not entail significant costs, b) does not require any special education, as it can be fully automatic, c) ensures that more farmer will have increased profits, d) preserves the resource (in most cases). Thus, it is self evident that enabling such a decentralized and distributed management procedure, benefits both the farmers' income and the sustainability of the environment, at all times.

4. Policy Assessment

4.1. Environmental Impact

In the context in this paper, the water remaining at the end of the season is the only parameter used for evaluating the impact the irrigation policies have to the environment.

4.2. Impact on the Population of the Farmers

The survival of a farmer's cultivation depends on the overall amount of the water gathered, since the production of a cultivation field is directly related to the quantity of irrigation water. The goal of each farmer during the cultivation period is to gather enough water for his cultivation to 'survive' and make a profit. Depending on the type of the cultivation, there is a threshold S_T defining the minimum quantity of water that the crop should be irrigated with during the cultivation period or else it is destroyed. Farmers that gather less water than this threshold are not considered to survive, either because their cultivation is destroyed or because their profit was not enough to sustain them financially.

The ratio of the number of farmers exceeding that minimum quantity of water S_T (N_I), to the total number of farmers (N), i.e.

$$S_R = \frac{N_1}{N} \tag{8}$$

is defined as the 'Survival Rate' S_R. It should be noted that S_R is estimated at the end of each simulation, thus farmers that do not eventually survive continue to drain water throughout the 160 days of a simulation.

4.3. Overall Impact

In order to assess the performance of the irrigation policies under study, with respect both to the farmers and to the environment, a new variable is introduced in this paper, i.e. the community value V_C. Actually, V_C is the weighted sum of all local values v_i that depict how well farmer i has performed during the cultivation period.

In order to evaluate v_i, and thus V_C, it is essential to calculate the deviation each farmer has from his initial goal (ga), formulated as:

$$ga = \frac{goal - actual\ water\ gathered}{goal} \tag{9}$$

In this sense, a ga value of 0.35 denotes that a farmer has reached 65% of his optimal goal. The local value v_i of each farmer is estimated as the product of two functions: a penalty function and a reward function. Considering that a cultivation may be completely ruined if not irrigated at least by a specific amount of water that depends on the type of the agriculture, it is evident that the profit and survival of a farmer highly depends on the water gathered throughout the season. This high correlation between the survival of the farmers and the percentage of the goal achieved, indicates that farmers that have not met their season goal, or at least a significant portion of it, should be punished since they do not contribute to the overall community value. Thus, the exponential decaying function shown in Figure 2 was selected for this purpose, i.e. $f_1(x) = 1/e^{2x}$. The form of this penalty function was further justified from empirical findings that emerged by studying the situations that arose in lake Koronia (Ioannidou et al., 2003). A closer examination reveals that the function output diminishes rapidly for ga values higher than 0.35, as highlighted in Figure 2. This implies that farmers that do not reach at least 65% of their goal should receive a severe penalty. It should be noted that by tuning the exponential decay and the 0.35 threshold, the model can be adjusted to include different type of cultivations.

Besides the penalty function, a reward function is required in order to favour farmers that managed to gather more water than others, attaining higher profits (within the expected profit limits related to their cultivation as set by S_T, under which the cultivation is considered to be destroyed, and 100%) and thus contributing to higher overall system values. For this purpose, function $f_2(x) = 2.8 \cdot 10^{-2} x + 0.8$, was selected, using linear regression principles. Empirical data were used for the linear regression, that were derived after studying the decision support system of Tzionas et al. (2004), that deals with the lake Koronia ecosystem. Evidently, the ‘strictness’ of the penalty and reward functions are encapsulated in their respective gradients.

As a result, the local value of a farmer i is defined, in the context of this paper:

$$v_i = f_2\left(\sum_{m=0}^{160} r_{i,m}\right) \cdot f_1(ga) \tag{10}$$

and the total community value V_C is thus given by:

$$V_C = \left(e^{\frac{l}{X}}\right) \cdot \left(\sum_{i=0}^{N} v_i\right) \tag{11}$$

where l is the final water resource level and X the initial water resource level. As expected, all local v_i values are weighted the same, for the estimation of the global V_C value. The exponential weighting coefficient preceding the summation of Equation 11 is introduced in order to punish farmer communities that adopt irrigation policies leading to low lake levels. In this sense, V_C also encapsulates the environmental impact of an irrigation policy, with policies of higher final water levels mapped to higher V_C values, and vice versa. Thus V_C is a non-linear combination of the two main measures used in this paper: a) the final water level related to environmental sustainability and b) the survival rates reflecting the farmers wealth.

5. Results and Discussion

In order to explore a) the performance of irrigation policies under the lack of communication amongst farmers and b) the consistency of the proposed self-adaptive learning algorithm in producing efficient resource allocations under any scenario, a series of extended simulations were conducted. These simulations include numerous computer-generated case studies, corresponding to the depletion of the resource, in analogy to actual situations that arose in the lake Koronia ecosystem (Piesold et al., 1999; Hellenic Ministry of Agriculture, 2001; Ioannidou et al., 2003; Tzionas et al., 2004; Mylopoulos et al., 2007; Kolokytha, 2010). Simulation scenarios that correspond to different combinations of farmer behaviours with respect to their greediness degree were created using a Monte Carlo approach similar to (Berger and Schreinemachers, 2006). A farmer was assigned to the equivalent greediness category by means of a random number generator producing numbers between 1 and 5. This procedure was repeated until all the possible combinations of farmers assignment to greediness categories was achieved. It should be noted that the proposed procedure was repeated for every irrigation policy under study, since all farmers are assumed to employ the same policy during each simulation. Thus, a farmer community was produced, that can be characterized by a ‘greediness degree’, which is the sum of all the greediness variables of the farmers consisting the population. Each one of these produced farmer communities is considered a simulation scenario. Although each simulation advances in daily time steps (until the limit of 160 days is reached) simulation results in this section are shown with respect to the greediness degree of the community. This is because that variable differentiates the farmer communities produced with the Monte Carlo procedure.

Considering that a population of 50 farmers was used for the conducted experiments, the expected range of values for the greediness degree lies between 50 and 500. It is evident that there are many combinations of possible farmer greediness values that could lead to a given greediness degree. Figure 3 depicts the number of such distinct combinations for each greediness degree.

Table 3. Min, Max and Median Standard Deviations of the Distinct Combinations for Each Greediness Degree, for all the Parameters under Study

	min	max	median
Water level	0	3.2	2.1
Survival Rate	0	1.5	0.9
Community Value	0	4.2	2.0

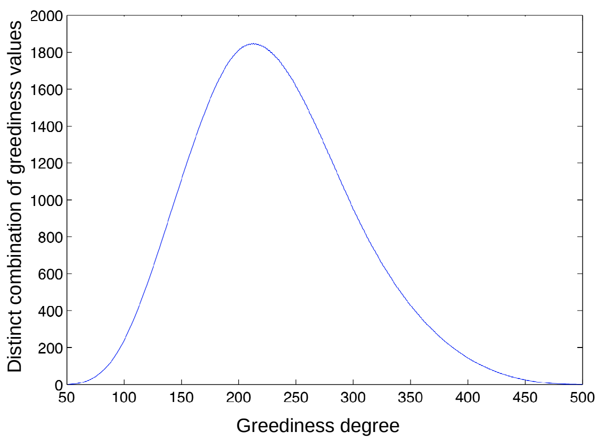

Figure 3. Number of distinct combinations of greediness values for every greediness degree.

It should be noted however, that the simulations for all the distinct combinations of a given greediness degree do not vary significantly. This is because the simulation outcome depends on the overall greediness degree. This is clearly depicted in Table 3, that presents the minimum, maximum and median standard deviations of the distinct combinations for each greediness degree, for all the respective parameters under study (presented in Section 4).

To that extent, in order to ensure the simplicity and clarity of the representations, each point plotted on the following figures corresponds to the average value of the respective parameters at the end of the 160 days, over the distinct combinations of each greediness degree.

The irrigation policies under study and the self adaptive learning algorithm introduced in this paper were evaluated utilizing the figures of merit presented in Section 4, in a two step procedure: in the first step, simulations employ each irrigation policy under study and farmers behave according to them. The goal of this step is to demonstrate the impact of each irrigation policy to the environment and to the farmers' community, given the imposed constraints (i.e. lack of communication/observation). In the second step, the policies employed in the first step are augmented with the self adaptive learning algorithm, thus adapting each farmer's behaviour towards efficient resource allocation. The goal of the second step is to evaluate the robustness of the proposed learning algorithm, i.e. its ability to produce efficient resource allocations under any simulation scenario and irrigation policy. It should be noted that this type of evaluation (i.e. in terms of robustness) is considered appropriate for multi-agent learning algorithms that deal with complex, real world problems which entail solutions that cannot be described by a unique typical equilibrium (i.e. Nash) (Mannor and Shamma, 2007).

5.1. Performance Evaluation with Respect to Water Level

The impact of each policy to the environment was assessed by evaluating the water level of the resource at the end of the season. The distribution of the water level at the end of the season with respect to the greediness degree of the community is presented in Figure 4. The horizontal axis denotes the greediness degree of a farmers community, i.e. a simulation scenario, and the vertical axis denotes the average water level over the distinct combinations of each greediness degree (i.e. simulation scenario).

It can be noticed that as the greediness degree increases, the water level decreases rapidly to 0 for all irrigation policies under study, denoting that in most experiments the resource is completely dried out (i.e. tragedy of the commons (Ostrom, 2002)). These results are also supported by findings in the ecosystem of lake Koronia, where irrigation policies resembling the NR policy (coupled with the lack of environmental awareness and the primitive infrastructures available to regulate water draining) were adopted resulting to its depletion (Piesold et al., 1999; Tzionas et al., 2004; Kolokytha, 2010).

A closer examination of Figure 4a reveals that the EF policy outperformed the other two policies. Before depletion, and for scenarios of the same greediness degree, more water is preserved under the EF policy. Alternatively, depletion of the resource is reached for scenarios of higher greediness degree, corresponding to higher water demands. This is in accordance with the EF policy design criteria (i.e. $k_W > k_P$) which ensures that environmental sustainability is the key factor of greater importance when selecting an action under that policy. For scenarios of high greediness degree however, corresponding to communities with increased water needs, the lack of communication between the farmers forced the EF policy to fail, verifying that the tragedy of commons is inevitable when communication is lacking (Ostrom, 2002). Both the PD and NR policies performed poorly, as it was expected, since actions under these policies are selected based on profit maximization criteria. In more detail, they both resulted to the depletion of the resource for small greediness degree values, with the PD policy providing worse results. This can be justified by considering that actions selected under the PD policy are of a more self-lucrative nature than the ones selected by the NR policy.

When the irrigation policies under study were augmented with the self-adaptive learning algorithm introduced in this paper, a significant degree of improvement for all the three policies under study was achieved (Figure 4b). It should be noted that there seems to be a sudden change in the behaviour, occurring around greediness degree 100. This is due to the

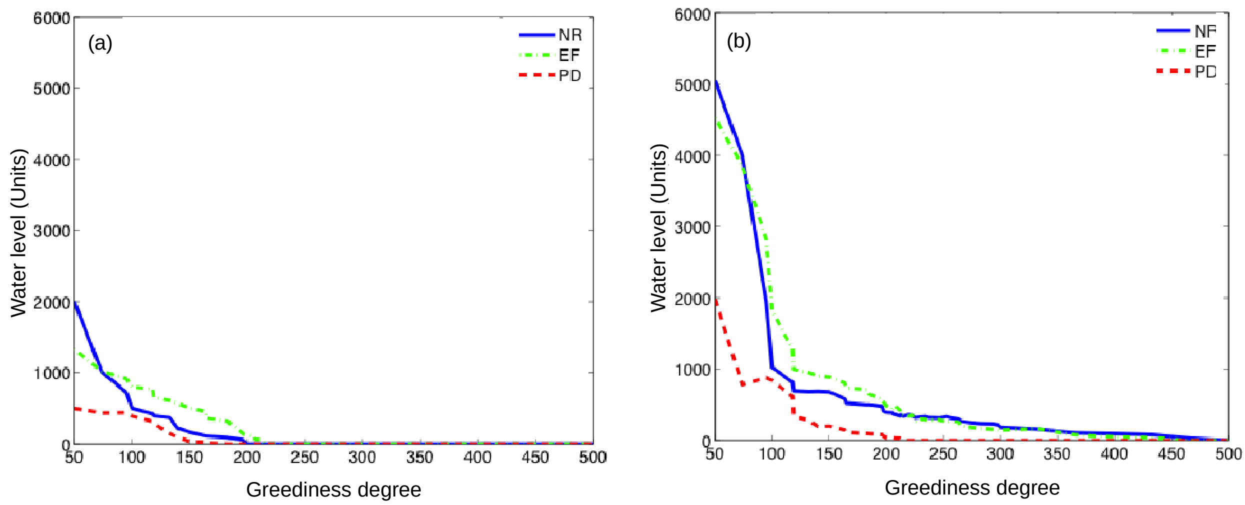

Figure 4. Water level after D days with respect to the greediness degree in the population (a) without the self adaptive learning algorithm and (b) with the self adaptive learning algorithm (diagrams are shown on the same scale for comparison purposes).

increased water needs of the population, as the overall requests of the farmers exceed the capacity of the lake, drain it below the threshold value T and their needs are further suppressed by the hydrological model of the lake (Equation 2). By cross examining the water level values of Figure 4a and Figure 4b it is clear that more water is preserved for the same greediness degree scenarios, under all policies. Additionally, for all the policies under study, the resource is depleted only at much higher values of greediness degree. It is evident then, that the employment of the self adaptive learning algorithm led to a significant increase in the preservation of the water of the resource, thus avoiding the tragedy of the commons despite the lack of communication amongst the farmers.

Table 4. Median and Average Water Levels at the End of the Season, for all Policies under Study (a) without the Self-adaptive Learning Algorithm and (b) with the Self Adaptive Learning Algorithm

(a)	NR	PD	EF
Median	0	0	0
Mean	118	93	156
(b)	NR	PD	EF
Median	327	18	388
Mean	564	171	612

In order to perform a quantitative comparison between the policies under study, the average and median values of the distributions of Figure 4 were extracted and presented in Table 4. In general, the median value is considered to be more representative in describing the central tendency of a distribution, since it is more robust to outliers. The median values of the water level at the end of the season are 0 for all policies under study (1st row of Table 4a), denoting that for most greediness degree scenarios, farmers dried out the resource. These results are also supported by findings in the ecosystem of lake Koronia, where irrigation policies resembling the NR policy, coupled with the lack of environmental awareness and the primitive infrastructures to regulate water draining, were adopted resulting to its depletion. To investigate the remaining scenarios, where water was preserved (for low greediness degree scenarios), the average of the water level at the end of the season was estimated (2nd row of Table 4a). The EF policy achieves the highest water level (156 units), verifying that action selection is targeted towards resource sustainability. The self-lucrative nature of farmers adopting the NR and PD policies results to lower water levels, i.e. 118 and 93 respectively. It is thus demonstrated that actions aiming at profit maximization when combined with the lack of communication amongst the farmers, lead to catastrophic results.

The 1st and 2nd rows of Table 4b correspond to the median and mean water level at the end of the season, respectively, under the self adaptive learning algorithm introduced in this paper. Median values are increased for all irrigation policies, demonstrating that more water is preserved in the resource (and thus less water is drained by the farmers). It should be noted that even policies that select their actions based on profit maximization criteria, i.e. PD policy, exhibit a significant relative increase in water level when augmented with the self-adaptive learning algorithm. These findings are also depicted in the average value of the water level, where 3.19 times more water was preserved in average, across all policies (comparing the 2nd rows of Tables 4a and 4b, respectively). Thus it is clear, that the self-adaptive learning algorithm coordinates the actions of the farmers towards the benefit of the environment, despite the specific policy employed.

5.2. Performance Evaluation with Respect to the Farmers Economic Survival Rate

The economic impact the policies under study have to the farmers community and their production is evaluated by utilizing the survival rate S_R of the population (Equation 8).

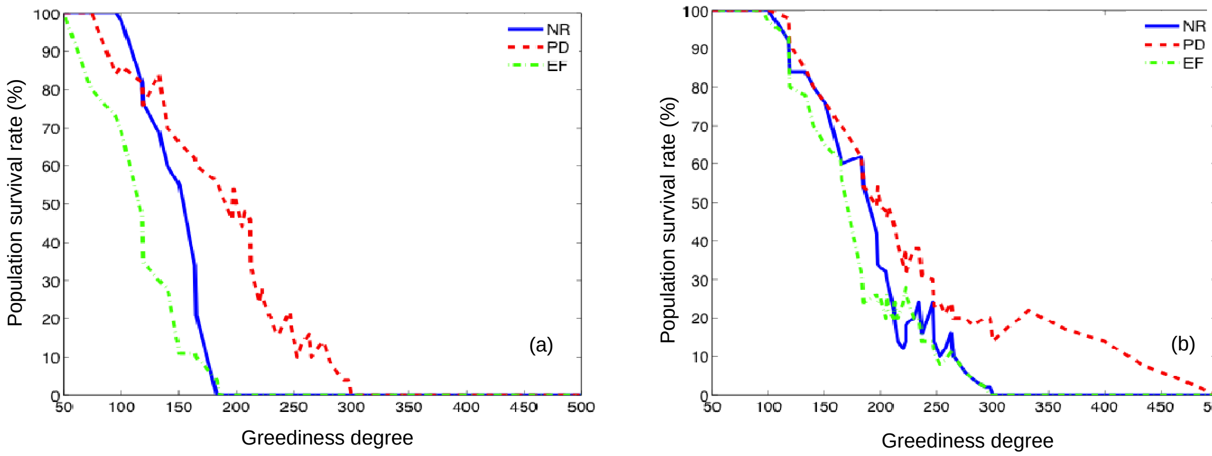

Figure 5. Survival rate with respect to the greediness degree in the population (a) without the self adaptive learning algorithm and (b) with the self adaptive learning algorithm (diagrams are shown on the same scale for comparison purposes).

Again, utilizing a Monte Carlo approach in order to produce farmer communities with varying greediness degree, the survival rate was estimated for a survival threshold $S_T = 60\%$ of the predefined goal G_i of each farmer. Figure 5 illustrates the average survival rate over the distinct combinations of each greediness degree (i.e. scenario), for all the irrigation policies under study.

For small greediness degree values (horizontal axis), which correspond to farmer communities with limited water needs, all farmers manage to survive, regardless the policy they adopt (demonstrated by high values on the vertical axis). As the greediness degree increases however, there is a decay in the survival rate distributions. Since the NR policy selects actions in a self-lucrative manner, it manages to sustain the economic survival of only a portion of the farmers and for low greediness degree scenarios (up to 200). It is outperformed by the sophisticated profit maximizing nature of the PD policy, that supports the economic survival of more farmers and for higher greediness degree scenarios (up to 300). However, their survival is achieved at the expense of the survival of other farmers in the community and also at the expense of the environment, since the NR and PD policies usually result to the depletion of the resource (see Figure 4a for the same greediness degree scenarios, e.g. 250). The EF policy provides the lowest survival rate for every greediness degree scenario, revealing that water preservation under the EF policy (see Figure 4a) is achieved at the expense of the economic survival of the farmers who are driven to an almost certain economic extinction. Given the specifications of each policy and the lack of communication amongst the farmers, it is evident that the water drained from the lake is not distributed evenly amongst the farmers, leading to the economic extinction of a significant portion of their population.

The respective survival rates are significantly improved when the self-adaptive learning algorithm is employed, as illustrated in Figure 5b. There is an absolute increase of the survival rate value under all three policies, as survival rates that correspond to the same greediness degree as in Figure 5a are now higher (as depicted in the vertical axes of Figures 5a and 5b). Additionally, for all the policies under study, the economic survival of a portion of the farmers is supported for higher greediness degree scenarios (as depicted in the horizontal axes of Figures 5a and 5b). Interestingly, in the case of the PD policy, a small portion of the farmers manages to survive regardless the greediness degree of the population. This clearly demonstrates that the self-adaptive learning algorithm proposed in this paper, provides a means of efficiently distributing the drained water, to the farmers community.

Table 5. Average Survival Rate of the Farmer Population for all Policies under Study (a) without the Self Adaptive Learning Algorithm and (b) with the Self Adaptive Learning Algorithm

(a)	NR (%)	PD (%)	EF (%)
Median	0	20	0
Mean	15	33	10
(b)	NR (%)	PD (%)	EF (%)
Median	20	38	20
Mean	32	47	29

Again, a quantitative comparison between the policies under study is performed, with the aid of the average and median values of the distributions of Figure 5 presented in Table 5. In more detail, when the NR and EF policies were employed none of the farmers managed to economically survive for most of the greediness degree scenarios (1st row of Table 5a). This is due to a) the non adaptive nature of the NR policy, that didn't modify the farmer's behaviour while the resource dried out, and b) the high environmental awareness implied in the EF policy, expressed by the inequality $k_W > k_P$ shown in Equation 5. On the contrary, a percentage of at least 20% of the farmers' population managed to economically survive under the PD policy, as a result of its profit maximization nature. The 2nd row of Table 5a accounts for greediness degree scenarios where at least a portion of the farmers managed to survive. Under the NR and EF policy, this portion is 15 and 10% in average, respectively. The PD policy ensures the

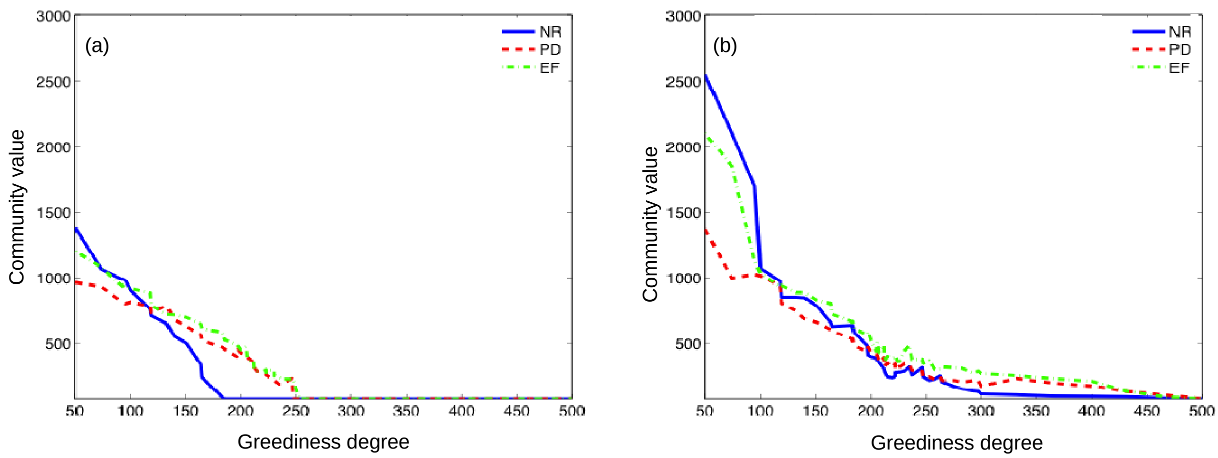

Figure 6. Community values with respect to the greediness degree in the population for all type of agents (a) without the self adaptive learning algorithm and (b) with the self adaptive learning algorithm (diagrams are shown on the same scale for comparison purposes).

economic survival of 33% of the population in average, but it must be stressed once again that this comes at the expense of the environment, since these scenarios lead to the depletion of the resource.

The survival rates were significantly improved by the use of the self-adaptive learning algorithm introduced in this paper, as presented in Table 5b. The median values are significantly higher, as at least 20% of the farmer population manages to gather enough water to sustain survival under the NR or EF policies. In the case of the PD policy, that already outperformed the other two, the introduction of the self-adaptive learning algorithm increased the survival rate to 38%. The average values of all policies are improved by a factor of 2.15, denoting that the number of farmers that additionally survive are more than twice. Returning to the previous discussion about the tragedy of the commons, where it was shown that less water is drained under the self adaptive learning algorithm, it must be pointed out that this significant increase in SR denotes that the water drained is also more efficiently distributed. This is because, water otherwise wasted by greedy farmers is now more efficiently utilized, as it is now redistributed to non-greedy farmers, assisting them to meet their goals and survive. Thus, one could consider that farmers implicitly coordinate their action in order for the whole community to survive.

5.3. Overall Performance Evaluation

The total community value V_C was introduced in Section 4 as a measure of the joint performance of the policies from an environmental and economic survival point of view, simultaneously. Figure 6 depicts the average community value over the distinct combinations of each greediness degree (i.e. scenario), for all the irrigation policies under study. High values of V_C correspond to simulations where most of the farmers have gathered enough water to sustain their production and at the same time, the resource was not depleted. On the other hand small values of V_C correspond to simulations where farmers did not manage to gather enough water to sustain their cultivation, and thus they are financially ruined, and at the same time there was a severe negative environmental impact since the resource was dried out.

The plots of Figure 6a reveal the fact that the penalty and reward functions (chosen in Section 4) were adequate to distinguish the different behaviours exhibited by the farmer populations in all simulations. The distributions for the EF and the PD policies follow a similar path, as discussed previously: under the EF policy the priority is given to the resource preservation and farmers do not gather enough water to economically survive whereas, on the contrary, under the PD policy the farmers do manage to survive but this is at the expense of the resource depletion. In accordance with the previous findings, the introduction of the self-adaptive learning algorithm results to a significant increase in the community value plots, as denoted in Figure 6b. By cross-examining the community value V_C of Figures 6a and 6b it is clear that the introduction of the self-adaptive learning algorithm, introduced in this paper, significantly increased its value for the same greediness degree scenarios, under all policies. Alternatively, the community value V_C diminishes for higher greediness degree scenarios, where there are higher water needs in the farmers population.

Table 6 presents statistical measures derived from Figure 6. Cross-examining Tables 4 and 5 with respect to Table 6, reveals that policies corresponding to low final water levels (Table 4) and low survival rates (Table 5) are assigned to low community values (Table 6). On the contrary, policies leading to high water levels and high survival rates are assigned to high community values. This was as expected, according to the definition of V_C in Equation 11. In more detail, the low survival rates and water levels related to the NR policy (Tables 5a and 4a respectively) are reflected in a low community value ($V_C = 124$). The PD policy may have a more negative environmental impact (Table 4a) but it achieves significantly higher survival rates (Table 5a), thus assigned to a greater V_C value, i.e. 204. The EF policy outperforms the other two policies, resulting to a community value of 241. This is as expected, since the EF policy may have performed worse in terms of the economic survival of the farmers (Table 5a), but it achieved the highest final water levels (Table 4a).

Table 6. Average Community Values, for All Policies under Study (a) without the Self Adaptive Learning Algorithm and (b) with the Self Adaptive Learning Algorithm

(a)	NR	PD	EF
Mean	124	204	241
(b)	NR	PD	EF
Mean	387	343	412

Similarly to the other performance evaluations, Table 6b presents the improvement achieved by the introduction of the self-adaptive learning algorithm. In more detail, the community values V_C were increased by a factor of 1.8 in average, demonstrating the fact that the performance of all policies is improved both with respect to the resource preservation as well as to the economic survival of the farmer population. Although a different policy outperforms the other two in each of the Tables 4 and 5, this is due to the fact that the respective policy evaluation criteria are different, i.e. final water level in Table 4 and the survival rate in Table 5. However, the determination of the best policy when taking under consideration both criteria is made clear when examining the V_C values presented in Table 6. There the EF policy augmented by the self-adaptive learning algorithm provides the best results, as reflected in the community value V_C. Thus, an overall balance between resource sustainability and survival of the farmers population is achieved, as discussed in the previous sections.

6. Conclusions

In this paper, a novel MAS simulation model was presented that simulates the exploitation of a water resource by a community of farmers. The novelty of the model lies in the fact that it employs agents/farmers that do not communicate with each other and they do not develop relations of trust with each other. This assumption was made since when farmers operate under economic pressure, they exhibit self-lucrative behaviours (Sen et al., 1996). Moreover, since if no other measures are taken the depletion of the resource is considered inevitable (Ostrom, 2002), a novel self-adaptive learning algorithm was introduced that provided a means of achieving efficient resource allocation using only local knowledge, despite the lack of communication.

Three typical irrigation policies were examined in a two-step procedure: The first step aims at demonstrating the impact of each irrigation policy to the environment and to the farmers community, given the imposed constraints. During the second step, each irrigation policy was augmented with the self-adaptive learning algorithm, in order to examine the robustness of the proposed algorithm. Extensive experiments were conducted following a Monte Carlo procedure, and objective figures of merit were used to estimate the impact of each irrigation policy to the environment, to the farmers community and to the overall ecosystem. It was verified that the lack of such communication lead to the depletion of the resource in most cases, despite the policy imposed (i.e. tragedy of the commons (Ostrom, 2002)) as none of the policies under study was capable to produce efficient resource allocation schemes.

When the proposed self-adaptive algorithm was employed (i.e. second step), results were significantly improved. Compared to the results of the first step (without the self-adaptive learning algorithm), the resource was depleted only at much higher values of greediness degree across all policies. At the same time, a higher portion of their community was supported for higher greediness degree scenarios. This means that, in most cases the resource was preserved and, at the same time, more farmers managed to sustain survival. Water was more efficiently distributed, since water otherwise wasted by greedy farmers was redistributed to non-greedy farmers, allowing them to meet their goals and survive. These findings were also depicted on the V_C value that provided a global measure of assessing the policies. It was thus demonstrated that in any case, even when farmers acted under profit maximization criteria, the introduction of the proposed learning algorithm ensured both the sustainability of the resource and the maximization of the farmers' profits. Considering the effectiveness of the proposed learning algorithm, the simplicity of its implementation and the minimal requirements that reduce costs, we strongly believe that farmers would adopt such a system without any special motivation. In this sense, the proposed model and learning algorithm are valuable tools for assessing alternative irrigation policies and providing the best policy for any given scenario. Further research on the subject is carried on, investigating the optimal values of several parameters of the proposed learning algorithm, as well as further extending the model to be applied to different ecosystems and cases of natural resource sharing.

References

Alexandridis, K. (2006). Exploring complex dynamics in multi agent-based intelligent systems: Theoretical and experimental approaches using the multi agent-based behavioral economic landscape (mabel) model. PhD thesis, Purdue University.

Athanasiadis, I., Mentes, A., Mitkas, P., and Mylopoulos, Y. (2005). A hybrid agent-based model for estimating residential water demand. *Simulation,* 81(3), 175-187. http://dx.doi.org/10.1177/0037549705053172

Barnaud, C., Bousquet, F., and Trebuil, G. (2008). Multi-agent simulations to explore rules for rural credit in a highland farming community of Northern Thailand. *Ecol. Econ.,* 66(4), 615-627. http://dx.doi.org/10.1016/j.ecolecon.2007.10.022

Barreteau, O., and Bousquet, F. (2000). SHADOC: a multi-agent model to tackle viability of irrigated systems. *J. Ann. Oper. Res.,* 94(1-4), 139-162. http://dx.doi.org/10.1023/A:1018908931155

Bars, M.L., and Attonaty, M. (2001). A multi-agent system to the common management of a renewable resource: Application to water sharing. IEEE Int Conf Tools with Artificial Intelligence, 42-49.

Bazzani, G. (2003). A DSS for an integrated analysis of irrigation at farm and catchment level: DSIRR. In: Int. Cong. on Modelling and Simulation (MODSIM '03), 1655-1660.

Becu, N., Perez, P., Walker, A., and Barreteau, O. (2001). Catchscape: An integrated multi-agent model for simulating water management at the catchment scale, a Northern Thailand case study. In: Int. Cong. on Modelling and Simulation, 1141-1146.

Berger, T., and Schreinemachers, P. (2006). Creating agents and landscapes for multiagent systems from random samples. *Ecol. Soc.,* 11(2), 19.

Berger, T., Birner, R., Diaz, J., McCarthy, N., and Wittmer, H. (2007). Capturing the complexity of water uses and water users within a multi-agent framework. *Water Res. Manage.,* 21(1), 129-148. http://dx.doi.org/10.1007/s11269-006-9045-z

Bishop, C. (2007). Pattern Recognition and Machine Learning (Information Science and Statistics). Springer.

Bousquet, F., Page, C.L. (2004). Multi-agent simulations and ecosystem management: A review. *Ecol. Model.,* 176(3-4), 313-332. http://dx.doi.org/10.1016/j.ecolmodel.2004.01.011

Bousquet, F., Bakam, I., Proton, H., and Page, CL. (1998). Cormas: Common-pool resources and multi-agent systems. In: Proc. 11th Int. Conf. Industrial and Engineering Applications of Artificial Intelligence and Expert Systems (IEA/AIE '98), Springer-Verlag, 826-837.

Brown, D., Robinson, D., An, L., Nassauer, J., Zellner, M., Rand, W., Riolo, R., Page, S., Low, B., and Wang, Z. (2008). Exurbia from the bottom-up: Confronting empirical challenges to characterizing a complex system. *Geoforum,* 39(2), 805-818. http://dx.doi.org/10.1016/j.geoforum.2007.02.010

Buşoniu, L., Babuska, R., and De Schutter, B. (2008). A comprehensive survey of multi- agent reinforcement learning. *IEEE Trans. Syst. Man Cy. (C Appl. Rev.),* 38(2), 156-172. http://dx.doi.org/10.1109/TSMCC.2007.913919

Chemonics International Inc (2004). Analysis of the decentralization of environmental management in Egypt. Tech. rep., Egyptian Environmental Policy Program Executive Committee.

Collier, N.T., and North, M.J. (2011). Repast SC++: A Platform for Large-scale Agent-based Modeling, Large-Scale Computing Techniques for Complex System Simulations, Wiley (In Press).

Deadman, P. (1999). Modelling individual behaviour and group performance in an intelligent agent-based simulation of the tragedy of the commons. *J. Environ. Manage.,* 56(3), 159-172. http://dx.doi.org/10.1006/jema.1999.0272

Feek, W., and Morry, C. (2003). Communication and Natural Resource Management. Food And Agriculture Organization Of The United Nations. PMid:14656624

Feuillette, S., Bousquet, F., and Goulven, P.L. (2003). Sinuse: a multi-agent model to negotiate water demand management on a free access water table. *Environ. Model. Software,* 18(5), 413-427. http://dx.doi.org/10.1016/S1364-8152 (03)00006-9

Ginot, V., and Page, C.L. (1998). Mobidyc, a generic multi-agents simulator for modeling populations dynamics. In: Proc. 11th Int. Conf. Industrial and Engineering Applications of Artificial Intelligence and Expert Systems (IEA/AIE '98), Springer-Verlag, 805-814.

Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., Goss-Custard, J., Grand, T., Heinz, S.K., Huse, G., Huth, A., Jepsen, J.U., Jrgensen, C., Mooij, W.M., Müller, B., Peer, G., Piou, C., Railsback, S., Robbins, A., Robbins, M., Rossmanith, E., Rüger, N.R., Strand, E., Souissi, S., Stillman, R., Vab, R., Visser, U., and DeAngelis, D. (2006). A standard protocol for describing individual-based and agent-based models. *Ecol. Model.,* 198(1-2), 115-126. http://dx.doi.org/10.1016/j.ecolmodel.2006.04.023

Gurung, T. (2004). Use of multi-agent system to improve irrigation water sharing in lingmuteychu watershed, Bhutan. Master's thesis, Chiang Mai University.

Hardin, G. (1968). The tragedy of the commons. *Science,* 162(3859), 1243-1248. http://dx.doi.org/10.1126/science.162.3859.1243

Hellenic Ministry of Agriculture (2001). Water quality characteristics for the lakes and rivers of Greece. Tech. rep., www.minagric.gr.

Ioannidou, I., Paraskevopoulos, S., and Tzionas, P. (2003). Fuzzy modeling of interactions among environmental stressors in the ecosystem of lake koronia, Greece. *Environ. Manage.,* 32(5), 624-638. http://dx.doi.org/10.1007/s00267-003-0001-x

Izquierdo, J., Montalvo, I., Perez-Garcia, R., and Gutierrez-Perez, J.A. (2010). A multi-agent framework for an IEDSS in urban water management. In: Int.l Cong. on Environmental Modelling and Software, International Environmental Modelling and Software Society.

Jaffe, R., Paniconi, C., and Wood, E., (1988). Model Calibration Based on Random Environmental Fluctuations. *J. Environ. Eng.,* 114 (5), 1136-114. http://dx.doi.org/10.1061/(ASCE)0733-9372(1988)114:5(1136)

Janssen, M.A., and Ahn, T.K. (2006). Learning, signaling, and social preferences in public- good games. *Ecol. Soc.,* 11(2), 21.

Klügl, F., and Puppe, F. (1998). The Multi-Agent Simulation Environment SeSAm. In: H. Kleine Büning (Hrsg.): Proceedings des Workshop "Simulation in Knowledge-based Systems", Paderborn.

Kolokytha, E. (2010). European policies for confronting the challenges of climate change in water resources. *Curr. Sci., Special Issue: Clim. change water resour.,* 98(8), 1069-1076.

Laycock, A. (2007). Irrigation Systems Design, Planning and Construction. Cromwell Press.

Li, A. (2011). Modeling human decisions in coupled human and natural systems: Review of agent-based models, *Ecol. Model.,* 229, 25-36.

Lin, S., Mullen, J., and Hoogenboom, G. (2008). Farm-level risk management using irrigation and weather derivatives. *J. Agric. Appl. Econ.,* 40(2), 485-492.

Luke, S., Revilla, C., Panait, L., and Sullivan, K. (2004). MASON: A New Multi-Agent Simulation Toolkit. Proc. of the 2004 Swarm Fest Workshop.

Mannor, S., and Shamma, J. (2007). Multi-agent learning for engineers, *J. Artif. Intell.,* 21(7), 417-422. http://dx.doi.org/10.1016/j.artint.2007.01.003

Matthews, R., Gilbert, N., Roach, A., Polhill, J., and Gotts, N. (2007). Agent-based land-use models: a review of applications. *Landscape Ecol.,* 22(10), 1447-1459. http://dx.doi.org/10.1007/s10980-007-91 35-1

Ministry of Rural Development and Floods (2000). Greek national action plan fro combating desertification (in greek).

Monticino, M., Acevedo, M., Callicott, B., Cogdill, T., and Lindquist, C. (2007). Coupled human and natural systems: A multi-agent-based approach. *Environ. Model. Software,* 22(5), 656-663. http://dx. doi.org/10.1016/j.envsoft.2005.12.017

Mylopoulos, N., Mylopoulos, Y., Veranis, N., and Tolikas, D. (2007). Groundwater modeling and management in a complex lake-aquifer system. *Water Res. Manage.,* http://dx.doi.org/10.1007/s11269-006-9025-3

Oel, P.V., Krol, M., Hoekstra, A., and Taddei, R. (2010). Feedback mechanisms between water availability and water use in a semi-arid river basin: A spatially explicit multi-agent simulation approach. *Environ. Model. Software,* 25, 433-443. http://dx.doi.org/10.1016/j.envsoft.2009.10.018

Ostrom, E. (1990). Governing the Commons: The Evolution of Institutions for Collective Action. New York: Cambridge University Pres. http://dx.doi.org/10.1017/CBO9780511807763

Ostrom, E. (2002). Managing resources in the global commons. *J. Bus. Admin. Pol. Anal.,* 30-31, 401-413.

Ostrom, E., Gardner, R., and Walker, J. (1994). Rules, Games, and Common-Pool Re-sources. University of Michigan Press.

Ostrom, E., Burger, J., Field, C., Norgaard, R., and Policansky, D. (1999). Sustainability - revisiting the commons: Local lessons, global challenges. *Science,* 284(5412), 278-282. http://dx.doi.org/10.1126/science.284.5412.278

Pauly, D., Christensen, V., and Walters, C. (2000). Ecopath, Ecosim, and Ecospace as tools for evaluating ecosystem impact of fisheries. *ICES J. Mar. Sci.,* 57, 697-706. http://dx.doi.org/10.1006/jmsc.2000.0726

Piesold, K., Karavokiris, G., and Anelixi, A.A.S. (1999). Environmental rehabilitation of lake koronia. Tech. rep., European Commission Directorate General XVI, Regional Policy and Cohesion, Greece. PMid:10234958

Queensland Government (2009). Callide valley groundwater management area water sharing rules. Tech. rep., Department of Environment and Resource Management Queensland State (2009) Callide Valley groundwater management area water sharing rules. The State of Queensland (Department of Environment and Resource Management).

Rasmussen, L.N., and Meinzen-Dick, R. (1995). Local organizations for natural resource management: Lessons from theoretical and empirical literature. Environment and Production Technology Division, International.

Schlüter, M., and Pahl-Wostl, C. (2007). Mechanisms of resilience in common-pool resource management systems: an agent-based model of water use in a river basin. *Ecol. Soc.,* 12(2), 4.

Sen, S., Sen, I., Roychowdhury, S., and Arora, N. (1996). Effects of local information on group behavior. In: Proceedings of the International Conference on Multiagent Systems, AAAI Press, 315-321.

Shajari, S., Bakhshoodeh, M., and Soltani, G. (2008). Suitability of multiple-criteria decision making simulations to study irrigation water demand: A case study in the Doroudzan riber basin, Iran. *Am. Eurasian J. Agric. Environ. Sci.,* 1, 25-35.

Shivakoti, G., and Ostrom, E. (1992). Institutions and Collective Action: Self-Governance in Irrigation. CA: ICS Press.

Shivakoti, G., and Ostrom, E. (2002). Improving Irrigation Governance and Management in Nepal Oakland. CA: ICS Press.

Smith, J. (1998). Does decentralization matter in environmental management? *J. Environ. Manage.,* 22(2), 263-276. http://dx.doi.org/10.1007/s002679900102

Suleiman, A.A. (2005). Syrian agriculture at the crossroads. FAO agricultural policy and economic development series. *Agric. Syst.,* 83(3), 336-337. http://dx.doi.org/10.1016/j.agsy.2004.08.007

Thebaud, O., and Locatelli, B. (2001). Modeling the emergence of resource sharing conventions: an agent-based approach. *J. Artif. Soc. Soc. Simul.,* 4, 2.

Tsiouris, S., Mamolos, A., Kalburtji, K., and Barbayiannis, N. (2002). Fertilizer management in watersheds of two ramsar wetlands and effects on quality of inflowing water. *J. Environ. Manage.,* 29(5), 610-619. http://dx.doi.org/10.1007/s0026701-0062-7

Tzionas, P., Ioannidou, I., and Paraskevopoulos, S. (2004). A hierarchical fuzzy decision support system for the environmental rehabilitation of lake Koronia, greece. *Environ. Manage.,* 34(2), 245-60. http://dx.doi.org/10.1007/s00267-004-0307-3

Walker, W. (1989). Guidelines for designing and evaluating surface irrigation system. FAO Irrigation and Drainage Paper, 45, 137.

World Bank (2007). Reengaging in Agricultural Water Management, 1st edn. World Bank Group.

Yi-Chen, E.Y., Cai, X., and Stipanovic, D. (2009). A decentralized optimization algorithm for multiagent systembased watershed management. *Water resour. Res.,* 45(8), 1-8. http://dx.doi.org/10.1029/2008WR007634