Rules, Knowledge and Complexity: How Agents shape their Institutional Environment♠

Alex Smajgl*, **Luis R. Izquierdo****, **Marco Huigen****

* *CSIRO Sustainable Ecosystems, University Drive,Townsville QLD 4810, Australia*
tel/fax: +61 7 4753 8650
e-mail: alex.smajgl@csiro.au
** *Universidad de Burgos, Edificio la Milanera, 09001, Burgos, Spain*
*** *Universität Hohenheim, Stuttgart, Germany*

Submitted: 02/02/2010
Accepted: 10/03/2010
Appeared: 20/03/2010

Abstract—This paper aims to contribute to the advancement of modelling *endogenous* changes of institutional rules (Ostrom 2006) within social groups. We understand that changes in institutional rules can be modelled exogenously or endogenously. Modelling rule changes *exogenously* means that it is the modeller who explicitly defines when and what institutional rules come into place. Modelling changes of institutional rules *endogenously* involves modelling groups of individuals who are able to *autonomously* invent, assess, and propose social rules. Agent-based modelling is used to simulate this emergent property of a social system in the context of water use. This paper explains the relevance of epistemological characteristics, such as the ability to reason or to learn, for modelling institutional change. The paper argues that rule changes are relevant for many applied models and shows that endogenous rule mechanisms can lead to different model results.

Keywords: Institutions; Norms; Agent Based Modelling; Emergence; Social Simulation; Adaptation and Learning.

1. INTRODUCTION

Agent-based modelling is becoming increasingly popular (Gilbert 2008), particularly in the context of applied decision making (see, for instance Bousquet and Le Page (2004)). In these situations, institutional rules (such as legislation) have often been modelled as separate entities that can shape agents' behaviour, but which cannot themselves be altered by the agents. There are many present-day socio-ecological problems that are governed by (both formal and informal) social rules which evolve so quickly that one may observe dramatically different institutional environments within the temporal scope of the modelling exercise. The example of climate change is a case in point: Societies implemented a diverse list of regulatory or market-based mitigation strategies, including cap-and-trade systems. In these fast-moving social environments, the modeller cannot ignore the mechanisms that trigger institutional changes; ideally one should model these mechanisms explicitly. This is the so-called endogenous approach to modelling rule changes, where the mechanisms by which individuals shape and change their institutional environment is explicitly modelled.

The aim of this paper is to investigate the ability of agents to adapt not just at the level of their action choice but also at the collective choice level. This aim involves two steps, implementing endogenous rule generation mechanisms and testing the relevance of such mechanisms for agent-based simulation. Previous work (Smajgl, Izquierdo, and Huigen 2008) developed an initial approach using a generic structure for rules and norms (Crawford and Ostrom 1995). Implementing such dynamics has been a challenge as it requires artificial agents to perceive changes and develop a new set of rules that leads to improvements from the individual or the collective perspective. Such a mechanism has to be *endogenous* if the aim is not to pre-code institutional solutions but to see rules or norms emerging. Generally, variables are classified as endogenous if they are explained by the model while exogenous variables are defined from outside (OECD 2008). A rule would be exogenous if the modeller defines its implications and potentially the conditions under which it applies. Rules are endogenous if the modeller provides only the building blocks that agents may use during the simulation run to generate a rule. The rule emerges as an endogenous variable of the collective. While the modeller still defines the building blocks, emerging rules cannot be easily foreseen, but they emerge as a mechanism of 'higher endogeneity'.

♠ The authors are grateful for the funding provided by the CSIRO OCE Science Team (Julius Career Award). L.R.I. gratefully acknowledges financial support from the Spanish Ministry of Education (grant JC2009-00263).

This paper is focused on rule generation by an agent that identifies externalities and hence the need for system-level change. We discuss this using an example of a shared aquifer from which agents extract water. Rules are generated by unsatisfied agents who evaluate them before communicating them to other agents. Generally required steps such as diffusion dynamics, negotiation processes (see for instance Wanyama and Farr (2007)) are simplified here as they do not represent the main focus of this paper.

A core part of rule generation is the agent's cognitive ability to learn causal relations from past experience, often labelled as reasoning, inference or pattern recognition (Pearl 2000). Epistemological aspects of developing this cognitive ability are described in this paper and implemented as an algorithm mimicking the development of simple heuristics. This paper argues that individual knowledge and institutional dynamics are tightly interwoven. The underlying hypothesis is that allowing institutional rules to emerge endogenously leads to simulations that are substantially different from model runs without the ability of artificial societies to adapt by creating new rules and norms. For an applied modelling context in which real world adaptations of norms and rules are likely to occur, this is a critical aspect if models are developed to improve system understanding of model users. This paper shows for applied agent-based simulations that institutions do not matter under all assumptions. Results show that it depends on the agent ability to learn and on the heterogeneity of the agent population. The model and its code can be downloaded at

http://www.luis.izquierdo.name/models/EndoMAL.

The next section lays out the institutional background of this work, followed by a discussion of aspects of learning and causality. Following these two sections, the model design is presented. The conclusion revisits the hypothesis and discusses epistemological consequences of modelling rules endogenously.

2. SYSTEM-LEVEL RULES AND AGENT ACTIONS

Human activities often affect other humans' welfare or utility; when such influences are not appropriately compensated in economic terms, they lead to so-called externalities (Hanley, Shogren, and White 1997). In order to constrain certain actions and enable others (Bromley 2006) social groups develop rules and norms, often labelled as *institutions*. Ostrom (1992, p.19) defines institutions as "the set of rules actually used (the *working rules* or *rules-in-use*) by a set of individuals to organise repetitive activities that produce outcomes affecting those individuals and potentially affecting others." Rules differ from norms as they contain an additional element that defines a penalty mechanism (Ostrom 1992), as explained below.

The driving mechanism for institutional change is how individuals and groups of individuals perceive and assess the prospective outcome of undertaking alternative actions (Ostrom, Gardner, and Walker 1994). Changes in such rule sets can trigger changes in other rules or norms as the conditions of decision making changes (Kiser and Ostrom 1982; Smajgl and Larson 2007). Such flow-on effects in real socio-ecological systems create a dilemma for applied simulation models as they often include the assessment of changes in incentive structures or actual rules. Keeping institutional adaptation exogenous by either neglecting such processes or by defining specific conditions for activating pre-selected rules is unsatisfactory for simulations that require periods of many years (Smajgl 2004;Smajgl, Izquierdo, and Huigen 2008). In real situations, institutions are likely to change, as societies adapt to, for instance, technological, environmental and cultural changes. Therefore, applied models are in danger of neglecting a critical social mechanism. This is so especially if the applied model assumes learning agents but limits the agents ability to create feedback to the system level, since in reality individual system understanding can translate into changes in rules and norms.

Agent-based models (ABMs) are computational models which contain an explicit and individual representation of the entities of the target system being modelled and of their interactions (Gilbert 2008). Agents in the model can represent individual entities such as humans with various levels of cognitive capacity, but also groups of individuals and non-cognitive environmental entities. As the system representation is developed from the perspective of individual entities (bottom-up approach) agent-based modelling allows for the analysis of "evolving systems of autonomous interacting agents" (Tesfatsion 2002). As Deadman (1999) points out, instead of defining the overall behaviour, in ABMs "this overall behaviour emerges as a result of the actions and interactions of the individual agents." This makes agent-based modelling effective in analysing complex adaptive systems (Miller and Page 2008). In-depth descriptions of agent-based modelling can be found in, for instance, Holland and Miller (1991), Holland (1992) and Gilbert (2008).

In order to equip our agents with sensible actions and the ability to formulate rules regarding these actions we use Crawford and Ostrom's concepts of agents' strategies and rule evolution. Crawford and Ostrom (1995) define a *Grammar for institutions* and provide a general structure of rules (ADICO sequence). These are built from the elements: *Attributes*, *Deontic*, *Aim*, *Conditions* and *Or else*. Following this grammar we define an agents' action then as a rule minus the 'Deontic' and the 'Or else'. In the figure below an illustrative example of a rule according to the ADICO-sequence is given.

Attribute A	Deontic D	Aim I	Conditions C	Or else O
"All irrigators	must not	demand ground-water from the aquifer	at all times the water level is below X	or else the farmer will be levied a fine of Y by the local police."

Assuming such a generic grammar, a generic mechanism can be developed that enables agents to combine new rules

(ADICO sequences). This paper explains this mechanism and adds the cognitive capacity agents require to put such a mechanism into the context of their goals and actions.

The structure proposed by Crawford and Ostrom (1995) has to be disaggregated for the computational processes required. Specifically the *Aim* and the *Condition* have to be split into four sub-elements (fields) each to allow the implementation of a generic mechanism. This paper proposes the I_field to be split into

I_verb	I_bound	I_quantity	I_object
"demand	less than	50 Ml	water"

and the C_field into

C_object	C_bound	C_quantity	C_time
"waterstock	lower than	500,000 Ml	at any time"

Such a structure allows agents to develop new rules by combining such field elements in a discovery process. A new rule would always have the following fields:

ADICO =
[A_Field].[D_Field].[I_verbField].[I_boundField].
[I_quantityField].[I_objectField].[C_objectField].
[C_boundField].[C_quantityField].[C_timeField].
[O_quantityField]

Conditions define when rules are active. The *Aim* fields impose restrictions for the individual action choice. As this paper aims for *endogenous* rule generation a feedback mechanism is required that translates rules back into actions. This can be achieved by a generic structure for defining agent actions that is constructed by a subset of ADICO fields:

action =
[I_verbField].[I_objectField].[I_quantityField].
[C_objectField].[C_timeField]

This means that before executing actions, agents identify all ADICOs in place with the same fields and adjust if necessary the I_quantity. In this paper actions always have all quantities aligned, such as for *demand water*, but the approach also allows for actions defined as strings.

However, such processes require a high level of cognitive capacity in which agents realise that externalities exist (imposed by other agents' actions), that a system-level intervention (institutional emergence) could help the agent to achieve better outcomes according to her own goals, and their ability to evaluate the outcome of a new rule, which requires a certain system understanding and existing causalities. Developing such causalities requires designing an algorithm which allows agents to link actions and system variables with agent goals (i.e. payoff).

3. INDIVIDUALS AND THEIR SYSTEMS UNDERSTANDING

Developing system understanding requires the ability of agents to reason and, in our case, to develop an understanding of what is causing their profit to increase or to decrease. Because the agents' understanding of causality must not be predefined, two underlying conditions on the reasoning mechanism are distinguished. These conditions then split the logical reasoning sequence into two phases. The first phase assumes the agent believes to have control over its profit and the second phase comes when externalities are perceived.

For the first phase we assume that the agent compares, in each period, the realised profit with her own aspiration (Simon 1955). If aspirations are met the agent will continue to execute the same actions as before (e.g. she demands the same amount of water as last time). Conceptually the strategy is reinforced and could be labelled as reinforcement learning (Erev and Roth 1998; Roth and Erev 1995) without actual calculation and updating of probabilities, such as in the traditional Bush-Mosteller model (Bush and Mosteller 1955). (For a broad overview of learning algorithms see Fudenberg and Levine (1998).) Within this group of routine-based learning (Brenner 2006) the aspiration level is introduced to assume satisficing (Simon 1955), which avoids a constant strive for 'optimal' profit. Based on the concept of habituation (Clark 1999; Hodgson 2004) the aspiration level is updated based on recent profits (Witt 1986). In this case it is assumed that agents define their aspiration value as the average profit of the previous five periods.

The second phase commences when the agent realises that her own actions are not the only determinants for their payoff. It is assumed that an agent reaches this point after having attempted to achieve her aspiration for a certain length of time unsuccessfully (further explained below in the sub-section ADICO generation). In this second phase the agent learning changes: The agent reads previous payoffs and the state of other variables, for example rainfall and water level from its memory. It is assumed that in the real world, people's cognitive ability is constraint regarding memory and information-processing ability (Todd 2001) and that people build decision rules based on simplifying heuristics (Gigerenzer, Todd, and ABC Research Group 2000). The heuristic that is defined for this paper assumes that the agent observes pairs of variable values for each memorised time step (unidirectoral bivariate test). The correlation between two variables is computed as the number of time-steps where the two variables moved in the same direction (i.e. either both increased or both decreased) minus the number of time steps where the variables moved in opposite directions (i.e. one increased and the other decreased). Time steps where at least one of the variables did not change are not taken into account. These correlations shape the agent's belief of what is causing the achievement of agent goal(s) (similar to a Granger causality (Granger 1969)). Hence, this variable is referred to as the agent's causality variable.

For example, farmers start with a certain knowledge about crops and irrigation, which is mimicked for the purpose of this paper by running 100 time steps of history to allow agents to develop their individual perceptions of causal relationships. Results from the actual simulations are kept in the same database and agents are able to look back for the length of their memory. This implies that agents form their own beliefs, which can vary according to experiences agents perceived over time. Such an approach assumes that causality is based on the system representation agents develop, which Williamson (2006) classifies as "epistemological causality" based on Hume's view on how humans perceive causalities. This entails that agents, as humans, can develop 'incorrect' causalities. (For interesting experimental results on how humans develop understanding of causalities see Schlottmann et al. (2006).) The simple heuristic this paper applies deals with the diagnosis aspect implicitly as it allows agents to draw conclusions regarding what influences an agent goal such as payoff. More complex algorithms that allow agents to observe causality in a more explicit way have been developed (Ciampoliniet al. 2001; Pearl 2000; Peng and Reggia 1990) but are not necessary for the purpose of this paper. It is not the aim of this paper to develop *good decisions* based on *correct causalities* (Mazlack 2008; Nakashima, Matsubara, and Osawa 1997; Williamson 2006). In this hypothetical case agents may develop rules that are sub-optimal based on unrealistic causalities as long as the algorithm allows agents to understand system level dynamics.

4. MODEL DESIGN

4.1 Base Functions

The EndoMAL model has been implemented in netlogo and can be downloaded at

http://www.luis.izquierdo.name/models/EndoMAL.

The model assumes a hypothetical case in which farmers apply seed and water to their land in order to achieve a satisfying profit. No other goal is considered. Satisfaction is achieved when profit equals at least the agent-specific aspiration threshold. Based on the concept of habituation (Clark 1999; Hodgson 2004) agents define their aspiration value as the average profit of the previous five periods. Each agent is endowed with the heuristic algorithm described above, which allows them to develop causalities that are kept in a knowledge matrix with a memory. If not stated differently, the simulations below assume a memory of 50 time steps.

Each agent decides first how much seed to demand from the seed market by multiplying the size of land owned with how much seed is needed per ha, which is assumed to be 1,000.

Secondly, the agent calculates how much water to demand from the water stock all agents share. If the water stock is large enough, all agents receive what they demand, otherwise every agent's demand is reduced proportionally. For each time step the agent is satisfied, demand remains the same for the following period. If the agent is not satisfied, she looks back in history and tries to find a period when profit equals or exceeded the current aspiration and demands the amount of water she received at that time step. If such a period does not exist the agent changes her demand of the previous period by a random amount between +10% and -10%. The agent cannot observe other agents' water demand. Other information from the history, such as water stock conditions is not yet visible to the agent. This requires the agent to learn that the system level variable can impose externalities to her profit function. This triggers the rule generation process, as described below. For the purpose of focussing on the *institutional* process other aspects of how individual learning translate into collective action are ignored.

Thirdly, agents apply seeds and water to their land parcel.

Fourthly, the agents supply their produce, which is based on a crop growth function, represented by Figure 1. It is assumed that crop growth is solely dependent on land size and the amount of applied water. The underlying crop growth function increases exponentially until the optimal irrigation level and a decrease for excessive irrigation. The exponential function is

$d = \alpha^{1+\beta\cdot(x-1)}$ with α and β as shape parameters, and x as amount of applied water.

Decreasing returns can be achieved by inserting the exponential function into the exponent of the same mathematical form:

$$f = \alpha^{1+\beta\cdot(x-d)} \tag{1}$$

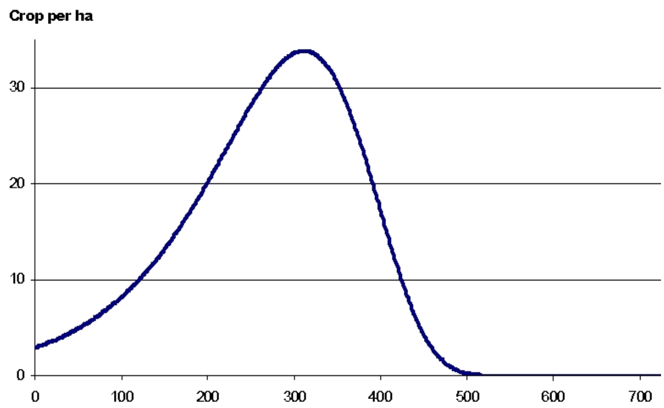

Fig. 1. Crop growth as a function of water applied for one ha irrigated land

Profit is calculated as $\pi = r \cdot p_{crop} \cdot f_{crop} - c_{seed}$, with r denoting land size and p_{crop} denoting the price for the crop the agent sells. For simplicity, a constant crop price of 1 is assumed. Seed costs are assumed to be 10.

At the end of each time step rainfall increases the water stock and defines the start conditions of the next time step. If the sum of received water over all agents exceeded precipitation, the water stock at the beginning of a period is lower than at the beginning of the previous period (or equals nil).

The user interface allows the user to determine the range of possible values for rainfall, the habituation period, the length of memory, length of history building phase, number of farmers, initial water stock, shape parameters of profit function, range of random water demand in history, and a range of values for the shape parameters in the profit function.

4.2 ADICO generation

Agents that do not achieve their aspiration enter a phase in which they adjust their action trying to achieve their aspiration threshold. How long agents explore their action space depends on the agent's belief in her own ability to achieve her goal, which is defined as the difference between her causality values (waterdemand $\Rightarrow$ profit minus waterstock $\Rightarrow$ profit). If the difference Δ_{cause} is high, it is assumed that the agent believes to have control over its own profit; the lower the delta Δ_{cause} between these two causalities the weaker this belief. It is assumed that unsatisfied agents try to achieve their aspiration for Δ_{cause} -time steps by adjusting their water demand before starting to believe that externalities prohibit satisfaction.

Agents that remain Δ_{cause} -time steps or longer unsatisfied generate a new rule in each time step they remained unsatisfied. An agent reads the most recent profit that achieved the current profit aspiration. If an agent cannot find such a record no ADICO is generated. If an agent finds a time step with such a profit she reads the amount of water she received (called my-past-sat-record) and places it in the I_quantity field. Additionally, the agent reads the conditions at that time (regarding C_object) and places in the value for water stock at that time step in the C_quantity field.

The agent then reads the causality value. If it is positive the following ADICO is created:

"all agents" "must" "demand" "less than or equal to"
([water received] of my-past-sat-record) "water stock" "lower than"
([C_quantity] of my-past-sat-record) "at all times" "Big Fine"

If causality is negative, the following ADICO will be created:

"all agents" "must not" "demand" "less than or equal to"
([water received] of my-past-sat-record) "water stock"
"greater than"
([C_quantity] of my-past-sat-record) "at all times" "Big Fine"

The generic design does not restrict ADICO sequences to emerge for water demand only. In case of the second scenario comparison noise is implemented at the level of seed application, which also leads to rules on seed application. Agents develop causality values for the land-profit link based on their action 'Apply seed to land' and infer a rule (or norm) such as 'All agents must not apply less or equal than 1000 seeds to their land or else big fine'.

This approach implies several simplifications:

Firstly, agents translate *system failure* into 'all agents'. This step can be much more sophisticated by allowing agents to read, for instance, attributes of other agents in order to distinguish who is causing the externality. Other simplifications are that the deontic of a rule is limited to 'must' and 'must not'. Lastly, the *or else* element of the grammar of institutions defines the penalty. For the purpose of this paper it is assumed that this field remains unspecified as all agents are assumed to comply with an implemented rule.

Modifying this assumption would require dissecting the [O_field] into several components and linking them to the relevant decision making algorithms and profit functions. One field component is the type of penalty, allowing for options such as a monetary fine, change of access rights (loss of license), or a labour related penalty (i.e. community work). A second field component is the actual quantity such as the fine value, similar to I_quantity. A third component could allow for a temporal penalty dimension defining how often this penalising action is to be executed, i.e. a one-off payment. A fourth component is the receiving entity, for instance to whom to pay the fine. Depending on the states of these four penalty components agents would have to employ different mechanisms to allow generically for a penalty to emerge.

Most relevant literature is focused on measuring the impact of punishment or sanctioning with pre-determined fines (in Public Good Games, i.e. (Fehr and Gächter 2000) or Common Pool Resource experiments, i.e. (Ostrom, Gardner, and Walker 1994) but not on developing evidence for what type of penalties people choose or how much they would fine (exceptions are, for instance (Cardenas 2000; Cardenas, Janssen, and Bousquet 2008;Janssenet al. 2008)). Relevant literature and especially experimental evidence suggest at least two important factors in framing the penalty:

Firstly, people seem to use so-called anchors when having to define uncertain quantities (Tversky and Kahneman 1974) which suggests that people will use, for instance their own income as a reference base for defining fines.

Secondly, the expected response by others will lead to strategic behaviour (in dynamic games), which makes the implementation highly context dependent.

The high relevance of context and the scarcity of experimental data for robust assumptions push the implementation of endogenous penalty formation beyond the aim of this paper. Therefore, the *or else* component is left unspecified and unconditional compliance is assumed for all agents. The downside of this simplistic assumption is that it makes it difficult to distinguish between norms and rules (Crawford and Ostrom 1995). Hence, implementing mechanisms that allow for a distinction will be focus of future work.

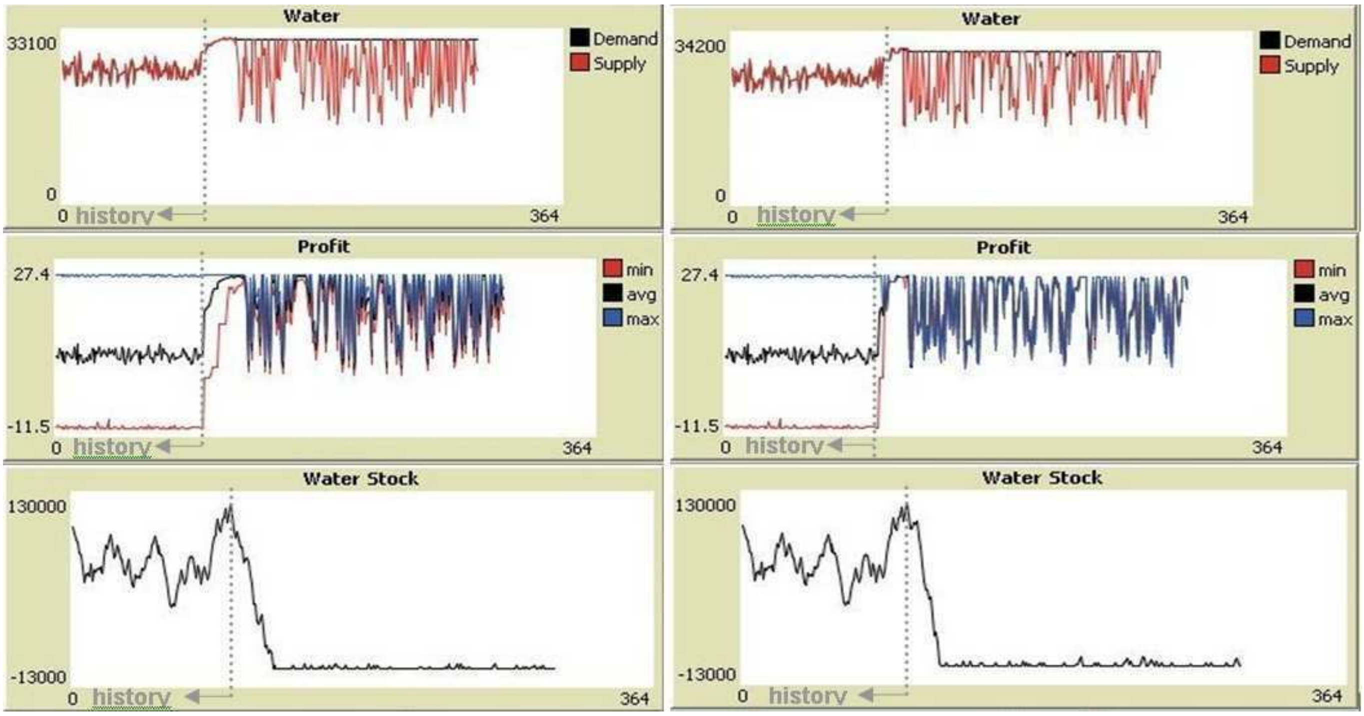

Fig. 2. Simulation results for a homogenous agent population without endogenous rule change on the left and with endogenous rule change to the right

4.3 ADICO Evaluation and implementation

The agent that discovered a new ADICO combination has to evaluate it before communicating the benefit of such system-level change to other agents. For the purpose of this paper it is assumed that the agent identifies the most recent time step with values that complied with the new rule and evaluates the rule positively if the profit at that time step exceeds the most recent profit. If the current profit is higher, they will not propose the rule. If they cannot find any such experience it is assumed that they propose the rule.

After agents evaluated the rule they composed positively, they communicate it. For simplicity reasons, this paper does not implement any diffusion dynamics, networks, or power relationships, as the focus is set on rule generation aspects. Rule implementation is based on the assumption that agents hand over their rule to a rule-implementation agent, which selects randomly one of the received ADICO combinations (per possible type) and implements the rule if at least 50% of the agent population is unhappy. Otherwise no rule is changed or implemented and all proposals are deleted. This assumption allows for a better testing of the hypothesis that endogenously emerging rules significantly change overall simulation results. Every non-random mechanism would have greater impact. However, new rules become active for the next period and each agent complies by checking if a rule applies (*Conditions*) and adjusts if necessary the behaviour (*Aim*).

5. RELEVANCE OF RULES FOR AGENT-BASED SIMULATION

This section uses representative runs to revisit the hypothesis that endogenous rule generation matters to agent-based simulation. Single representative runs have been chosen over multiple runs to test changes in patterns instead of absolute results. To test different conditions three scenario comparisons are provided: homogenous agents, homogenous agents facing complexity, and heterogeneous agents.

For the first comparison it is assumed that all agents share the same profit function. Knowledge, experiences and beliefs can vary due to random water demand in the history building phase. Under these assumptions endogenous rule creation is nearly irrelevant as depicted in Figure 2. After a history-building period of hundred time steps agents quickly identify a water demand that satisfies their aspiration, with or without an endogenous rule mechanism. In a period of decreasing water stock, agents remain demanding the satisfying quantity until the water stock is depleted. In this new purely rain-fed phase satisfaction is highly variable. Patterns are very similar for simulations with and without ADICO mechanism. This is due to the causal relations agents remember: As the phase of stable profits exceeds the memory most causal relationships have to be explored again. In other words, each agent starts learning from scratch and as all agents share the same profit function they identify the same causalities. In this situation of an almost homogenous population, rules do not seem to make a big difference.

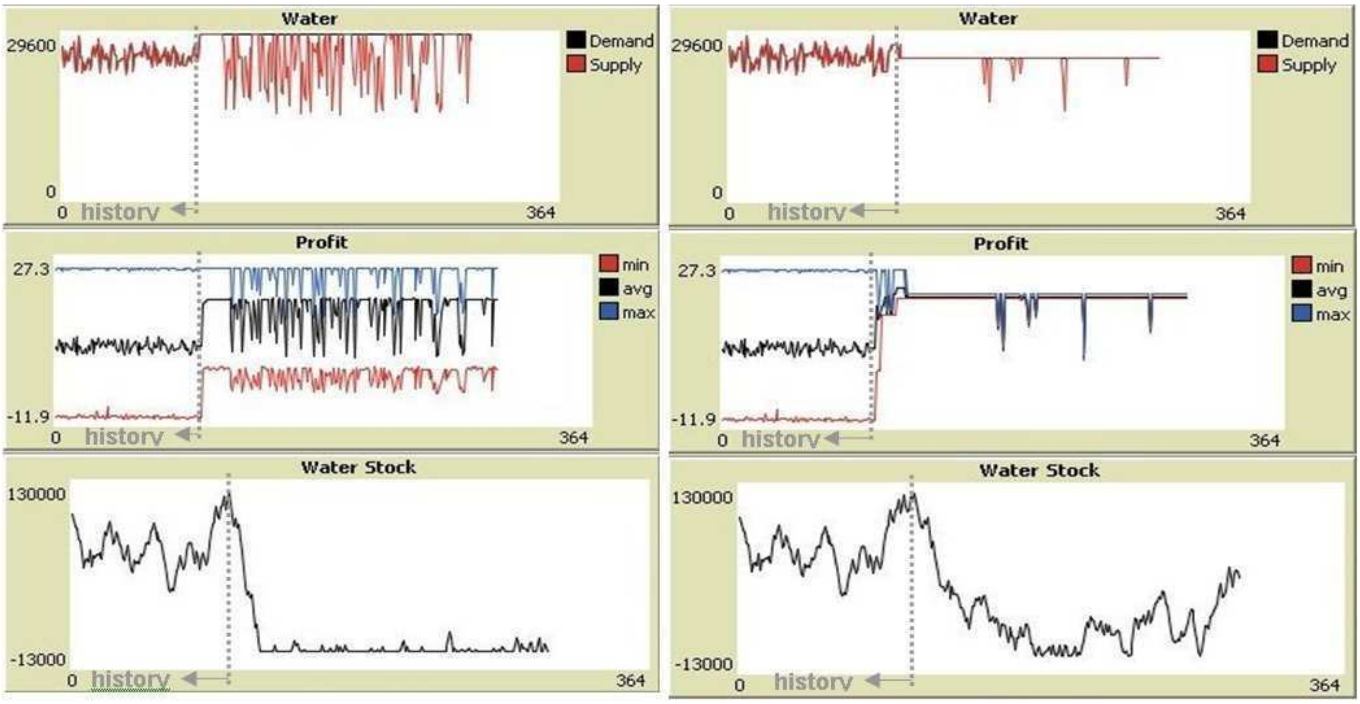

Fig. 3. Simulation results for a homogenous agent population without endogenous rule change to the left and with endogenous rule change to the right

If assumptions for this first comparison are modified regarding a longer memory (i.e. 100 combined with an earlier depletion of the water stock, i.e. rain 150-200 and initial water stock 100,000) the results for simulations with and without endogenous rules have also very similar patterns. The only difference is that profit distributions are narrower in the case with endogenous rule formation. This means that even with heterogeneity in agent knowledge, overall patterns are not different as long as the profit function is the same. This makes sense as rules signal strategies with higher payoff (for the rule proposing agent) and as all agents share the same optimum, rules do not make a big difference.

The second comparison makes the same assumptions as the first comparison apart from the costs for seeds c_{seed} that vary from time step to time step and between agents randomly between 9.5 and 10.5 (instead of constant 10). This level of noise makes it difficult for agents to develop unambiguous causalities. Figure 3 depicts on the left side for the case without endogenous rules that agents stick to the first value(s) they identify as beneficial. Just a few agents are close to the optimum, which leaves most agents demanding water far below their optimum, while the water stock keeps rising. Noise prohibits agents from learning satisfying strategies.

In the case where agents are able to generate rules endogenously the pattern is very different, as Figure 3 depicts on the right side. All parameter values are identical to the assumptions leading to results shown to the left in Figure 3 apart from the active ADICO mechanism. Figure 3 (right side) shows agents' demand is very close to the optimum right after the history building phase. However, already in the first time steps after the history some agents are unsatisfied and rules are generated and proposed. But as for the first nine time steps the majority of agents is satisfied no rules are implemented. In time step 110 the first two rules are implemented:

- ["all agents" "must" "demand" "less than or equal to" 314.45468191076947 "water" "water-stock" "lower than" 882845.3984020619 "at all times" "Big Fine"]
- ["all agents" "must not" "demand" "less than or equal to" 226.54562307620918 "water" "water-stock" "greater than" 882845.3984020619 "at all times" "Big Fine"]

These rules define an upper and a lower bound for all agents' water demand. From time step 124 on, the bounds are changed to the identical value of 314.45 (rounded), which defines the flat line that maintains until time step 878 when the water stock is depleted. The following new experience leads to the identification of a new threshold and due to the high level of dissatisfaction a new upper bound of 267.47 (rounded) is generated and implemented, followed in the next period with a lower bound of the same value. This rule set defines the constraints for water demand until the end of this simulation.

This result shows that one property of rules is that the individual knowledge on causality is brought into a social space and even if rules for implementation (as in this case) are chosen randomly, group-wide learning and hence profit is rapidly improving. Due to this property, patterns emerging in the simulations with endogenous rule mechanisms differ substantially from those without endogenously generated rules.

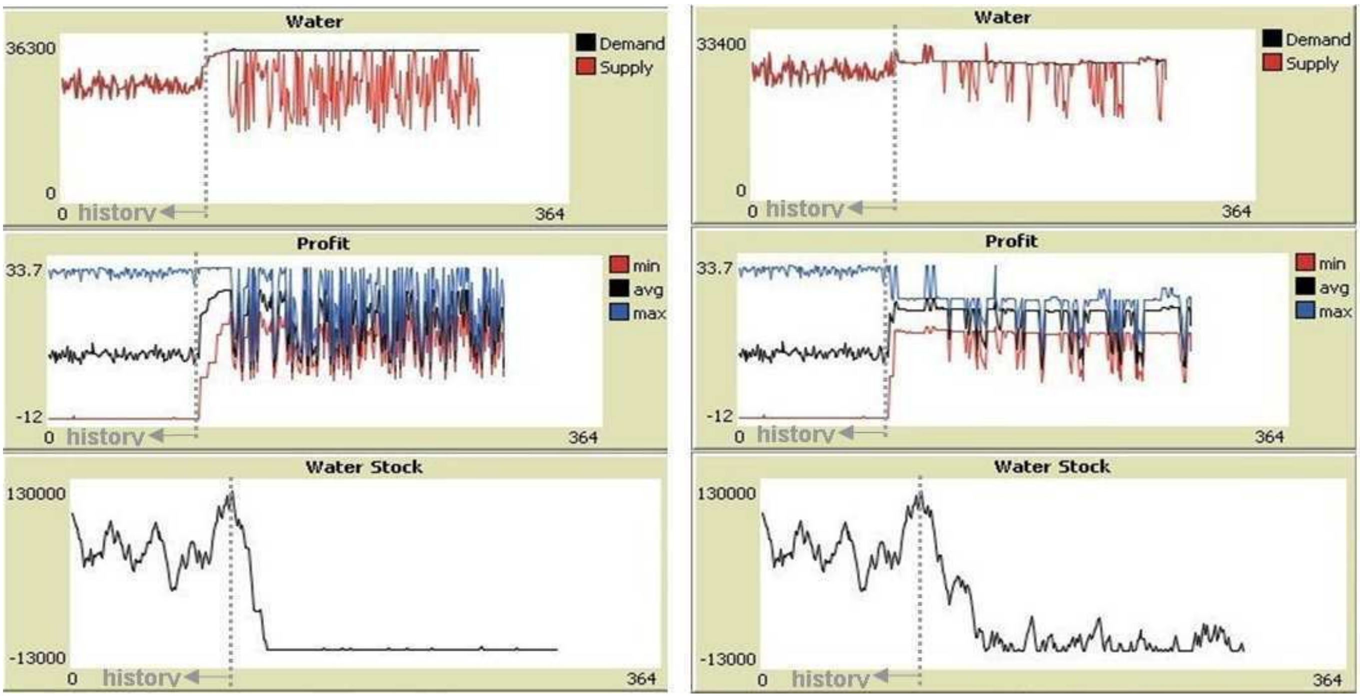

Fig. 4. Simulation results for a heterogeneous agent population without endogenous rule change to the left and with endogenous rule change to the right

This difference seems to matter particularly to agent-based simulations that aim to represent learning in a complex environment. If agents can easily identify causal relationships they can adapt based on their own knowledge base. In a more complex environment though, agents can pick up contradicting signals. As rules are based on at least one agent's system understanding individual learning is enhanced. This assumes that all agents share the same goal and that their profit function is identical as one agent's knowledge leads another agent to her optimum. This can be very different in case of heterogeneity, which is simulated in the next comparison.

The third comparison eliminates the noise and assumes instead heterogeneous profit functions, with α values between 2.5 and 3.5. Due to the different optima, agents generate rules that are beneficial to them (as in the simulations before) but are not likely to be beneficial to all other agents. This means that any new rule that is not generated by a specific agent is likely to reduce her profit.

In the case of no endogenous rule formation every agent learns quickly their optimum (no noise), as shown in Figure 4, and keeps it until water becomes scarce. Any changes in profit are mainly driven by the availability of water.

Endogenous rule formation leads to similar patterns, as depicted by Figure 4. One clear difference is the upper bound for profits: Without endogenous rules individual profits experience an upper bound by water scarcity. Endogenous rules add limitations proposed by farmers with lower optima. The underlying rule selection process is random, as described above. Every other selection criterion, such as an average of all rules, or related to power, or in protection of the water stock, would create an even bigger difference between the case with and without rule formation.

5. CONCLUSION

This paper developed a model with endogenously evolving rules that impose conditions on agents' action choice. The hypothesis that such mechanisms are relevant for agent-based modelling was tested in three comparisons. It has been demonstrated that in the case of homogenous agent populations such mechanisms do not affect overall simulation results. In the case of learning in noisy (complex) systems, agent-based models (even with homogenous agents) will have very different results if norms and rules can be formulated endogenously. This is due to the property of implicit knowledge about causal relationships that norms and rules contain. This can endure longer than the individual property and described therefore, a beneficial aspect of individual agents (in a homogenous population). In case of heterogeneity, rules imposed by others are likely to reduce profit of agents that learn their best actions.

The quintessence is that agent-based simulations that aim for improving the understanding of stakeholders in a real world context might be misleading if heterogeneity, complexity and dynamic norms and rules define relevant characteristics of the socio-ecological system to be modelled.

From a model user perspective endogenously emerging rules can have a drawback. In an applied context many agent-based models operate in a what-if modus that allows model users to define a scenario and observe the consequence of their change. The ADICO based mechanism allows agents to suggest changes to system level properties of the simulation. This reduces the ability of model users to control the simulations and requires additional monitoring mechanisms that allow the observation of endogenous adaptation processes that emerge.

Future research aims for improving the generic framework to allow many kinds of actions to be implemented. The showcased model also assumes the same learning mechanism for time steps before and after an agent reveals system-wide causalities. Depending on the context this requires more sophisticated changes in learning. Another limitation of the current model version is that new actions are likely to require their own decision-making algorithm, which seems to constrain the ability to develop a generic approach for all aspects of an agent-based model. Realistically, rules can also be generated by analogies drawn from a similar domain, for instance from one natural resource to another, which will be implemented in future model versions.

ACKOWLEDGMENT

The authors wish to thank Erin Bohensky, Art Langston and John Ward for their constructive comments and Christy Hortin-Matthes for her editorial work.

REFERENCES

Bousquet, F. and Le Page, C. (2004). Multi-agent simulations and ecosystem management: A review. *Ecological Modelling*, 176, 313-332.

Brenner, T. (2006). Agent Learning Representation: Advice on Modelling Economic Learning. In Tesfatsion, L. and Judd, K.L. *Handbook of Computational Economics*. 895-947. Amsterdam, North Holland.

Bromley, D. (2006). *Sufficient Reason: Volitional Pragmatism and the Meaning of Economic Institutions*. Princeton University Press.

Bush, R.R. and Mosteller, F. (1955). *Stochastic Models for Learning*. Jon Wiley & Sons.

Cardenas, J.C. (2000). How do groups solve local commons dilemmas? Lessons from experimental economics in the field. *Development and Sustainability*, 2, 305-322.

Cardenas, J. C., Janssen, M., and Bousquet, F. (2008). Dynamics of Rules and Resources: Three New Field Experiments on Water, Forests and Fisheries. In List, J. and Price, M. *Handbook on Experimental Economics and the Environment*. Cheltenham, UK, Edward Elgar Publishing.

Ciampolini, A., Lamma, E., Mello, P., and Torroni, P. (2001). LAILA: a language for coordinating adbuctive reasoning among logic agents. *Computer Languages*, 27, 137-161.

Clark, A.E. (1999). Are wages habit forming? evidence from micro-data. *Journal of Economic Behaviour & Organisation*, 39, 179-200.

Crawford, S.E. and Ostrom, E. (1995). A Grammar of Institutions. *American Political Science Review*, 89, 582-600.

Deadman, P. (1999). Modelling individual behaviour and group performance in an intelligent agent- based simulation of the tragedy of the commons. *Journal of Environmental Management*, 56, 159-172.

Erev, I. and Roth, A. (1998). Predicting how people play games: Reinforcement Learning in Experimental Games with Unique Mixed Strategy Equilibria. *The American Economic Review*, 88, 848-881.

Fehr, E. and Gächter, S. (2000). Cooperation and Punishment in Public Goods Experiments. *American Economic Review*, 90(4), 980-994.

Fudenberg, D. and Levine, D.K. (1998). *The Theory of Learning in Games*. MIT Press.

Gigerenzer, G., Todd, P.M., and ABC Research Group. (2000). *Simple heuristics that make us smart*. Oxford University Press.

Gilbert, N. (2008). Agent-based models. SAGE Publications.

Granger, C.W.J. (1969). Investigating causal relations by econometric models and cross-spectral; methods. *Econometrica*, 37, 424-438.

Hanley, N., Shogren, J.F., and White, B. (1997). *Environmental Economics*. Macmillan.

Hodgson, G.M. (2004). Reclaiming habit for institutional economics. *Journal of Economic Psychology*, 25, 651-660.

Holland, J.H. (1992). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. MIT Press.

Holland, J.H. and Miller, J.H. (1991). Artificial adaptive agent in economic theory. *American Economic Review*, 81, 365-370.

Janssen, M., Goldstone, R.L., Menczer, F., and Ostrom, E. (2008). The effect of rule choice in dynamic interactive spatial commons. *International Journal of the Commons*, 2(2), 288-312.

Kiser, L. L. and Ostrom, E. (1982). The Three Worlds of Action: A Metatheoretical Synthesis of Institutional Approaches. In Ostrom, E. *Strategies of Political Inquiry*. 179-222. Sage, Beverly Hills,CA.

Mazlack, L. J. (2008). Imprecise Causality in large Data Sets. Unpublished Work.

Miller, J.H. and Page, S.E. (2008). *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press.

Nakashima, H., Matsubara, H., and Osawa, I. (1997). Causality as a key to the frame problem. *Artificial Intelligence*, 91, 33-50.

OECD. (2008). Glossary of statistical terms. http://stats.oecd.org/glossary/detail.asp?ID=794 .

Ostrom, E. (2006). Multiple Institutions for Multiple Outcomes. In Smajgl, A. and Larson, S. *Sustainable Resource Use: Institutional Dynamics and Economics*. 23-50, 2006, Earthscan, London.

Ostrom, E. (1992). *Crafting institutions for self-governing irrigation systems*. Institute for Contemporary Studies Press.

Ostrom, E., Gardner, R.H., and Walker, J. (1994). *Rules, games, and common-pool resources*. The University of Michigan Press.

Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge.

Peng, Y. and Reggia, J.A. (1990). *Abductive Inference Models for Diagnostic Problem-Solving*. Springer Verlag, Berlin, New York.

Roth, A. and Erev, I. (1995). Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Run. *Games and Economic Behaviour*, 6, 164-212.

Schlottmann, A., Ray, E.D., Mitchell, A., and Demetriou, N. (2006). Perceived physical and social causality in animated motions: Spontaneous reports and ratings. *Acta Psychologica*, 123, 112-143.

Simon, H.A. (1955). A behavioural model of rational choice. *Quarterly Journal of Economics*, 69, 99-118.

Smajgl, A. (2004). Modelling evolving rules for the use of common-pool resources in an agent-based model. 10th Biennial Conference of the International Association for the Study of Common Property. 2004. Oaxaca, Mexico.

Smajgl, A., Izquierdo, L., and Huigen, M. (2008). Modelling endogenous rule changes in an institutional context: The ADICO sequence. *Advances in Complex Systems,* 2(11), 199-215.

Smajgl, A. and Larson, S. (2007). Institutional Dynamics and Natural Resource Management. In Smajgl, A. and Larson, S. *Sustainable Resource Use: Institutional Dynamics and Economics*. 3-19. 2007. EarthScan, London.

Tesfatsion, L. (2002). Agent-Based Computational Economics: Modelling Economies as Complex Adaptive Systems. Department of Economics, Iowa State University. 7-12-2002.

Todd, P. M. (2001). Fast and Frugal Heuristics for Environmentally Bounded Minds. In Gigerenzer, G. and Selten, R. *Bounded Rationality: The Adaptive Toolbox*. [4], 51-70. 2001. The MIT Press, Cambridge, Massachusetts.

Tversky, A. and Kahneman, D. (1974). Judgement under Uncertainty: Heuristic and Biases. *Science*, 185, 1124-1131.

Wanyama, T. and Far, B.H. (2007). A protocal for multi-agent negtiation in a group-choice decision making process. *Journal of Network and Computer Applications*, 30, 1173-1195.

Williamson, J. (2006). Dispositional versus epistemic causality. *Mind Mach*, 16, 259-276.

Witt, U. (1986). 'Firms' market behaviour under imperfect information and economic natural selection. *Journal of Economic Behaviour and Organisation*, 7, 265-290.

AUTHORS PROFILE

Alex Smajgl is an environmental economist and complex systems modeller. His applied work is focused on understanding socio-ecological interactions. His theoretical work is focused on institutional dynamics.

Luis R. Izquierdo is a complex systems modeller with a strong background in mathematics, computer modelling, and economics. His main area of expertise is game theory and the analysis of complex systems.

Marco Huigen is an agent-based modeller currently working in the GLOWA-DANUBE project and the PAK346 project at the University of Hohenheim. Both projects focus on the modelling of interactions between climate change and agricultural behaviour.