The evolution of institutions for common-pool resource management: An agent-based model

Giangiacomo Bravo
giangiacomo.bravo@unito.it

Draft, April 23, 2009
Do not quote without author's permission

Abstract

The joint exploitation of open-access natural resources is often modeled as a social dilemma with no escape for rational actors. Nevertheless, real individuals are not helplessly trapped in this dilemma and are often able to sustainably manage their resources by building endogenous institutions. The agent-based model presented here simulates the management of a common-pool resource by focusing on the relation between agents' beliefs and institutions. The conditions where agents can build management institutions lead to much better outcomes than the base model, where agents can only rely on individual beliefs in order to limit the resource consumption. This happens despite the fact that agents act in a competitive environment. Overall, higher sustainability can be obtained thanks to the establishment of institutions: just as observed in empirical settings.

Keywords: common-pool resources; agent-based simulation; institutional building; sustainable development; forest management

1 Introduction

Starting from Hardin's seminal article *The Tragedy of the Commons* (1968), the joint exploitation of an open-access natural resource is often modeled as a social dilemma, with no escape for rational actors. Nevertheless, subsequent work has shown that real individuals are not helplessly trapped in this dilemma and can instead sustainably manage their resources by building endogenous institutions (e.g. Berkes et al. 2003; Lam 1998; Ostrom 1990; Ostrom et al. 2002; Tang 1992).

This paper focuses on the institutional supply process by explicitly modeling the relation between the evolution of beliefs among the actors involved in exploiting the common resource and the development of management institutions. The idea of a macro-level institutional change driven by changing knowledge, beliefs, informal constraints and other micro-level attributes of actors is rooted mainly in North (2005) and Ostrom (2005) work. More specifically, actors are supposed to hold beliefs and mental models of the world, which influence both their perceptions (input from the external environment) and their choices (output to the external environment). Beliefs and mental models are formed and updated using two sources of information: the feedback received from the external environment and the shared belief systems. The process of forming beliefs and mental models is therefore internal, but is also influenced larger cultural/institutional structures. While the institutional structure influences individuals' internal models, it is not an independent variable. North (2005) argues for a strict relation between the institutional framework and individuals' belief systems The institutional framework is, from many points of view, an external manifestation of the latter and its evolution strongly depends on changes occurring in belief systems. At the same time, the institutional structure constrains choices and actions, directly or indirectly influencing individual's beliefs. This picture is, of course, an over simplification both of empirical reality and of North's ideas. Nevertheless, for the purposes of our work, it is enough to outline the existence of a mutual feedback loop between micro-level shared beliefs and macro-level institutions, which constrains individuals' actions, but are simultaneously influenced by the aggregation of beliefs.

This framework, however, remains a largely informal one. The nature of the relation between belief change and institutional evolution needs to be defined with precision. What are the forces driving micro-level changes? How do individuals' internal changes translate into macro-level rule transformations? How can individuals acting in a competitive environment cooperate in order to build shared institutions? In order to try to answer to these questions, we created an agent-based model which simulates the management of a common-pool resource (CPR). The choice of the CPR situation is due to its importance in both theoretical and practical terms. Theoretically, it is a well studied case of social dilemma, whose implications extend far beyond the strict description of the situation (e.g. Ostrom 2005; Ostrom et al. 1994). Empirically, the management of shared natural resources is one of the key challenges facing humanity and its development (e.g. Diamond 2005; Stern 2007; Smil 2002; Volk 2008).

The model proposed here uses a forestry scenario. It can however be easily adapted to the peculiarities of other resources as well. The simulation focus lies in the relation between micro-level beliefs held by agents and macro-level institutional changes. What is especially important is the explicit modelization of internal beliefs. More specifically, each agent, besides having access to public information regarding the actual state of the resource, also holds individual beliefs about how this state should be and about the best actions to reach this desirable state. The agent's beliefs subsequently aggregates into institutional rules that modify the agent's actions and may also lead to the development of new beliefs, thanks to their influence on agents' behaviors and earnings.

Note that, in order to keep the model simple, many rules that are fundamental in real life are only assumed in our simulations. For instance, boundary rules defining who has the right to use the common resource are implicitly incorporated in the model by assuming a stable population of agents throughout the whole run. Similarly, collective-choice rules (Ostrom 2005, 58–60) are limited to a single qualified-majority rule. Despites those limitations, the proposed model produced interesting results. First, due to the competition among agents, the spontaneous beliefs are insufficient to allow a sustainable management of the common resource. Second, the introduction of the possibility of building institutions significantly improves resource use both economically and ecologically under a wide range of parameters. Third, the dynamics of the model substantially reproduce those found in real settings: open-access resources are doomed to destruction, but strong institutions are created following shocks due to resource overuse and subsequently persist over time (e.g. Berkes 1998; Carlsson 2003; Johannes 2002; Pinkerton 1998).

The paper is organized as follows. Section 2 presents the theoretical and empirical background. The base model — where beliefs evolve, but no institutions are created — is defined in Section 3. Management rules are introduced in Section 4, which presents the Institutional evolution (hereafter *Inst*) model. Section 5 concludes the paper by discussing its results.

2 Research background

Common-pool resources (or simply "commons") are natural or man-made resources shared among different users. This produces competition that often (although not necessarily) leads to their degradation or even to destruction. Many natural resources fall in this category and are today "chronically" overused. Examples are forests, fisheries, water basins, biodiversity and even the atmosphere. Following Hardin's original statement, the management of CPRs is often depicted as a social dilemma and formalized using different variations of n person Prisoner's dilemmas or Public good games. These all share the idea that the rational equilibrium of the game is well below the collective optimum theoretically achievable by restricting resource use to a sustainable level (e.g. Casari and Plott 2003; Dawes 1973; Milinski et al. 2002; Ostrom 1990; Ostrom et al. 1994). In contrast with theoretical predictions, much empirical research, in particular Ostrom (1990) seminal work, shows that successful management of the resources can be achieved by building endogenous institutional. More specifically, the "tragedy" is avoided thanks to institutions that define clear exploitation rights and create incentives to prevent resource overuse. In other words, the tragedy of the commons is the tragedy of open-access resources, not necessarily of well managed CPRs.

Most studies on CRR management is field-based work. Nevertheless, experimental research also plays an important role. An important finding of CPR experiments is that, when the protocol allows communication and/or mutual sanctioning, the participants actively try to devise a shared set of rules to avoid resource overuse (e.g. Cardenas 2000, 2003; Janssen et al. 2008; Janssen and Ostrom 2008; Ostrom et al. 1994). One advantage of using experimental procedures is that the very process of institutional building can be observed in real time. In the course of standard field-work, researchers can rarely follow the discussion and bargaining processes that lead to resource managing rules and subsequently observe the actual effects of the new institution. On the other hand, many CPR experiments allow communication among the participants between each exploitation round. This means to observe the process whereby participants agree on a given management rule can be observed. It is worth noting that, by agreeing on a shared set of rules to avoid resource overuse, the participants solve a second order social dilemma. Formally, this is not easier to solve than the original resource dilemma. However, both CPR and Public good experiments show that subjects are willing both to invest to create rules that foster group cooperation and to punish noncooperators (e.g. Barclay 2006; Cardenas et al. 2000; Fehr and Gäcther 2000, 2002; Henrich et al. 2006). The fact that second order social dilemmas are relatively easy to solve is probably due to the psychological satisfaction that humans derive both from cooperating with other cooperative individuals and from punishing free-riders (de Quervain et al. 2004). At a deeper level, it is probably part of our evolutionary adaptation for cooperation in small groups of hunter-gatherers that represent the human ancestral environment (Boyd et al. 2003; Fehr and Fischbacher 2003; Gintis et al. 2003; Richerson and Boyd 2005).

Of particular interest is the recent series of experiments by Marco Janssen and colleagues using an innovative "dynamic interactive spatial commons" experimental platform (Janssen et al. 2008; Janssen and Ostrom 2008). In contrast with standard CPR experiments, based on abstract descriptions of the game structure and discrete time events, Janssen developed a visual platform allowing real-time interaction among participants. They play by moving their "avatars" on a two-dimensional space where the resource units harvested by the participants "grow" (for a detailed description of the experimental platform, see Janssen et al. 2008). The main aim in building the new platform was to "create a natural resource harvesting situation with continuous opportunities for repeated decisions regarding the speed and amount of harvesting [...], and with an intuitive, interactive way to harvest" (Janssen et al. 2008, 294). Besides improving the participant's comprehension of the situation, the platform allows more realistic management rules, for instance to define property rights based on spatial boundaries or specific harvesting periods, followed by resource restoration phases.

In the initial round of both experiments, participants in groups of five were allowed to move their avatars freely in order to collect tokens representing the resource, each worth $0.01. The resource was a renewable one and new tokens appeared in real-time on the virtual space following a density-dependent probabilistic function. Each experimental round lasted a few minutes. Under this conditions, the participants completely destroyed the resource in usually less than two minutes, hence realizing sub-optimal earnings (on average, 1/4 of the potentially feasible ones) just as in a standard CPR experiment.

In the subsequent round of the first experiment (Janssen et al. 2008), participants were allowed to vote for the implementation of an externally enforced private property rule (i.e. to split the commons into equal parts among them).[1] Voting for the rule was costly and the resulting institution was actually implemented only when the majority of subjects were in favor of the new rule. The implementation of the institution had therefore the structure of a second-order social dilemma: more precisely, a public good game with a provision threshold. Nine out of the twenty groups that enjoyed the voting possibility actually decided to implement the institution. Those groups significantly improved the management of the commons and their earnings comparing with both the initial round and the groups that failed to implement it.

The second experiment (Janssen and Ostrom 2008) substituted the voting procedure with a ten minutes communication phase that occurred between the different exploitation rounds. The communication allowed participants to exchange information and to devise their informal management rules. Nevertheless, no formal regulation or enforcement of any rule was possible and participants played the subsequent rounds using a protocol identical to the first round one. In most groups, participants actively used the communication time in order to coordinate their action. One common result was the splitting of the commons into private "properties", using different visual attributes of the resource system.[2] Other groups tried to envisage strategies enabling a regeneration of the resource, e.g. waiting before the start of harvesting, establishing resting periods, or envisaging "zigzag" movements of avatars to avoid the formation of large bare areas. Despite the fact that the envisaged institutions could not be formally enforced, most groups performed better in the second round, improving both resource conditions and participants' earnings. Even better was the performance during the third round, which followed a second communication period, used mostly to "tune-up" the management rules devised earlier. It is worth noting that, while the average performance during the second round of the second experiment was similar to the second round one of the first experiment (where the private property rule was formally enforced), the performance during the third round of the second experiment was even higher. This shows how various informal arrangements in a formally open-access situations can work even more effectively than a single externally-enforced rule.

This second experiment is especially interesting in showing the effectiveness of informal institutions devised by participants. Note that subjects taking part in the experiment were undergraduate students without any experience of real-world natural resource management. However, they rapidly "discovered" two of the strategies mostly used by real CPR managing communities, namely spacial allocation of the resource and control of the harvesting time. More generally, the experiment shows how easily participants agreed on building a management institution as a "natural" way of coordinating their action and fostering cooperation, despite the second-order social dilemma structure of the situation. This replicates the finding both of previous experiments and of a huge number of field studies, However, the possibility to observe the actual rule-shaping process, including its rapidity and its effectiveness, clearly brings added value to this study.

While Janssen and colleagues' experiments largely inspired our work, before presenting it it is worth summarizing some previous studies that used simulation techniques to model CPR exploitation. In one of the first studies using simulation to analyze socio-ecological dynamics with an explicit "tragedy of the commons" structure, Grant and Thompson (1997) developed a system dynamics model to check the different performance of optimizing (rational) vs. reciprocity-based strategies. Their results showed how reciprocity outperforms optimizing both economically and

[1] Actually, not all the participants' groups were allowed to vote for the rule. Control treatments included groups playing all the rounds like the first one, and others where the private property rule was imposed by experimenters. Since they are not of direct interest for our research, those treatments will not be discussed here.

[2] This rule was implemented especially when the participants' group included one or more subject that already participated to the first experiment.

ecologically. However, the model was extremely simple with only two "agents"[3] sharing the common resource, both using the same strategy.

In another study, Castillo and Saysel (2005) followed the long ecology tradition of system dynamic modeling to computationally replicate the results of a field experiment conducted with fishermen and crab hunters of the Colombian Providence Island, which replicated both the structure and findings of the Ostrom et al. (1994) and Cardenas (2000, 2003) experiments. Their model simulated the experimental data with only a relatively poor degree of accuracy. However, an interesting hypothesis deriving from the model was that crab hunters were more cooperative than fishermen. While plausible in the light of the different gender/age structure of the two populations (fishermen were adult males, while crab hunters mostly females and children), this hypothesis was unfortunately not empirically checked by the authors.

Using system dynamics to model CPR management suffers from significants shortcomings. While it may be appropriate to model the underlying dynamics of ecological systems, system dynamics can only roughly portray most of the individual attributes of resource users (preferences, strategies, norms of behavior, etc.) and completely fails to capture the dynamics of interaction. Agent-based modeling (ABM) is an approach better suited to model most aspects of social interaction. Using ABM, both individual actions and decision-making routines can be explicitly included in the model (Gilbert 2008). Probably the first work using ABM to model CPR situations was Deadman et al. (2000). Deadman modeled agents that replicated most of the findings of Ostrom et al. (1994) experiments, including the strong effect of communication on cooperation and sustainable use of the resource. It is worth noting that nothing was included in the model that directly specified system behavior, which resulted instead from the aggregation of individual agent choices. Nevertheless, the model succeeded in replicating the experimental findings, especially when the "communication" routine used a "central authority" to inform agents of the strategy that best performed in past rounds. What is especially interesting here is that the "central authority", although unable to enforce the proposed strategies, represented a rough sketch of an institution. This is similar to what happens in experiments, where individuals use the communication periods to devise arrangements that avoid resource overuse. Just as in the Deadman model, subjects participating in CPR experiments cannot formally enforce rules. Nevertheless, most of the time, they succeed in improving their performance thanks to the establishment of informal institutions.

In a recent paper, Janssen and Ostrom (2006) explicitly modeled the emergence of institutions in a population of heterogeneous agents. In their model, the CPR has a "physical structure" similar to that subsequently used as experimental platform.[4] Besides playing the CPR game, agents had to decide whether or not to implement an institution able to regulate the exploitation level to a fixed quantity. The decision routine of agents was based on two factors: the state of the resource and the amount of trust existing in the system. The latter depended on the outcome of trust games that agents played as a side activity that accompanied the main CPR game and on the heterogeneity of agents, modeled using "tags" identifying arbitrary belonging group. The main result of Janssen and Ostrom's work was that agents had to experience one or more resource crisis before being willing to create an institution. Nevertheless, in most conditions once the institution was in place it appeared to be able to coordinate the agents' actions and to significantly improve the economical and ecological performance of the system.

---

[3] Note that agents were not explicitly modeled. Rather their behavior was implicitly included in the system dynamic description.

[4] The development of the experimental platform was inspired by the simulation results (Marco Janssen, pers. com.).

3 The base model

Our model retains a number of elements from Janssen and Ostrom's platform, including the spatially defined resource and the possibility of institutional building. Nevertheless, its focus lies in the relationship between the internal (micro) states of the agents and system (macro) outcomes, including the establishment of institutions. The core of the model is the formalization of the "beliefs" (or mental models) that agents hold regarding the desirable state of the resource and the best way to achieve it. All outcomes, including the harvesting levels and (in the *Inst* model) the establishment of a management institution, depend on those beliefs.

According to North (2005), a strong relation exists between the belief systems and the institutional framework that humans use to coordinate their behaviors. Beliefs and other informal constraints influence human behavior both directly, through shaping the perception of the world and of what actions are appropriate in a given situation, and indirectly, by affecting the development of the institutional structure. Both influences have been considered in our model, even if the latter will be included only in the *Inst* one. At the same time, the macro states of the model, i.e. the competition among agents and the conditions of the resource, influence the agent beliefs. The resulting feedback loop — which is not explicitly modeled, but results from the interaction of the agent actions with the characters of the ecological system — is the main driver of the model, largely influencing its final outcome.

3.1 Base model definition

A total of $n = 100$ agents act on a regular lattice of degree $k = 8$, having the form of $m \times m$ toroidal surface with $m = 50$.[5] Each patch represents the forest area that can be logged in one round. Patches have the attribute $b_{xy} \in [0, b^{max}]$, with $x, y \in \{1, \cdots, m\}$, which stands for the total tree biomass present at a given moment, with b^{max} representing the maximum possible level of biomass per patch. At the beginning of the simulation, the whole area is green with b_{xy} randomly distributed in the $[\frac{1}{2}b^{max}, b^{max}]$ interval, simulating the exploitation of a mature forest area.

If not logged, trees in each patch with $b_{xy} > 0$ grow at a fixed rate (0.5 units per round) up to the point when they reach b^{max}. If the patch is empty (i.e. $b_{xy} = 0$), trees start to regrow with a probability depending on the state of the neighboring patches. More formally, the function defining the regrowing probability of an empty patch is

$$p = p^* \frac{N+1}{k+1} \tag{1}$$

where $p^* = 0.05$ is the basic probability of re-growth, N is the number of neighboring patches with $b_{xy} > 0$ and k is the degree of the graph, i.e. the neighborhood dimension. This implies that the regrowth probability of an empty patch with all its neighbor "green" (i.e. with $b_{xy} > 0$) is 0.05 per round, while the probability for a patch surrounded by other empty patches is only $0.00\bar{5}$ per round.

The function presented in (1) is similar to the one used by Janssen et al. (2008) for their "spatial commons" experiments, with the only difference that the re-growth probability for an isolated patch is small, but strictly above zero. Janssen's choice of having a zero probability of regrowth for isolated patches was motivated by the need of creating a vivid "tragedy of the commons" situation for participants. However, the function in (1) reproduces the natural recovery capacity of forests (at least of temperate ones) better, even after large scale clearing, due to seed conservation in the soil, seed dispersion by birds and other animals or other mechanisms.

[5] Notice that test exploration of the dynamics of the same model with up to 1000 agents, moving on a correspondingly increased surface, led to results similar to the ones presented below.

Agents are defined by two characters representing their beliefs. First, they have a general belief about the "right" overall level of biomass, which represents a sort of cognitive model about how the "world" should be. Formally, each agent $i \in \{1, \cdots, n\}$ has a character β_i that represents the fraction of the initial tree biomass that should be ideally conserved in the commons. At the beginning of each run, this character is drawn randomly from a normal distribution with mean 0.5 and standard deviation 0.25 and it remains subsequently constant. A second belief regards the level of cutting κ_i that is able to maintain the actual tree biomass at the desired level. Practically, this character is the minimal level of tree biomass that a given patch should have in order to be logged. All agents have $\kappa_i = 0$ when they enter the game, but this value can subsequently change (see below).

The simulation is divided into periods, each lasting 10 rounds. Each simulation run covers 2000 periods. Periods represents a relatively extended period of time. Each agent i has a payoff, which is set to zero at the beginning of each period and subsequently depends on its actions. In every round, i pays a fixed cost c independent of its actions. Its earnings depend instead on the amount of biomass on its current patch. More specifically, the agent starts logging only if $b_{xy} > \kappa_i$, where x and y are the coordinates of the current agent location. If the condition is true, b_{xy} is added to its current payoff. If the biomass in its current patch is less or equal to κ_i, the agent checks if any of the patches in the interval $x \pm 2, y \pm 2$ has biomass above that value. If this is true, it moves to one of those patches, pays the fixed cost and has no earnings in the current round. If none of the patches has enough biomass, the agent moves randomly to one of the patches included in the $x \pm 2, y \pm 2$ interval and has no earnings. Earnings and costs are summed within each period to form the agent payoff.

At the end of each period a few updates are made. First, agents update their κ_i belief. More specifically, each agent checks its current period payoff (P_t) and compares it with the one achieved in the previous period (P_{t-1}). If the current payoff is greater or equal to the previous one, it maintains its previous beliefs. If it is lower, the agent changes κ_i with a probability proportional to the absolute value of the difference between the two results. In formal terms, the probability q of changing κ_i is given by

$$q = \frac{|P_t - P_{t-1}|}{|P_t| + |P_{t-1}|} \tag{2}$$

with q rounded to one if greater than this figure. A random extraction determines whether the agent will actually change its belief. If this happens κ_i is modified according to β_i. More specifically, if $B_t > \beta_i B_0$, where $B_0, B_t = \sum_{x,y=1}^{m} b_{xy}$ represent the total resource biomass in the initial (0) and in the current (t) period respectively, the agent decreases κ_i by a random value in the interval $[0, 9]$. If $B_t < \beta_i B_0$, it increases κ_i by the same amount. The idea behind this function is that agents suffering from a payoff reduction are "unsatisfied" and are motivated to modify their beliefs and, consequently, their behavior. Since they posses a cognitive model of the "right" state of the world (i.e. β_i), if the current proportion of biomass is lower than this benchmark, they will ascribe the payoff reduction to an excessive cutting level and will therefore increase their own κ_i. Vice-versa, if the proportion of biomass is higher than the benchmark, they will reduce their κ_i and start to log patches with lower biomass.

At the end of the belief update routine, agents undergo a process of selection, representing market competition. This leads to the bankruptcy of unsuccessful agents and their substitution by new entrants. These are assumed to copy the best practices available in the system. The process occurs in two steps. First, one of the agents with the highest period payoff and one with the lowest period payoff are selected. Then a copy of the former (i.e. its β_i, while κ_i is always equal to zero when a new agent enters the game) replaces the latter. In this step, there is a 1% probability of "mutation", representing copy errors or simply new entrants with innovative beliefs. At the end of the selection routine all payoff are reset to zero and a new period starts.

The base model will be tested on low/high cost and low/high maximum biomass conditions, namely $c \in \{1,5\}$ and $b^{max} \in \{10,20\}$. Note that the actual value of the simulation parameters has little empirical meaning. The main goal of the parameter setting is to build an environment of difficult forest management, possibly leading to overharvesting of the commons and to low average payoffs. The base model will therefore serve as a test platform to check the effects of the subsequent introduction of the management institutions. All the models presented here have been implemented using the C++ language.[6]

3.2 In search for optimality

In order to make the results comparable and independent from the specific choice of the parameters, they will be presented as proportions of benchmark levels. As benchmark levels for the number of green patches and the total biomass, we will use the values of those variables at the start of the simulation. In order to evaluate the payoff benchmarks, for each parameter combination we ran simulations where all the agents were forced to harvest patches only when their biomass was above of a fixed threshold equal for all. In practice, we set the κ_i at a fixed level, equal for all agents. We explored all possible levels of $\kappa_i \in \{0, 0.5, \ldots, (b^{max} - 0.5)\}$, corresponding to all possible levels of biomass in each patch except b^{max} (remember that setting $\kappa_i \geq b^{max}$ means that the agents simply are not allowed to harvest), by running 50 simulations for each of them. Under this condition, the system reaches an equilibrium in about 20 periods. We therefore limited the number of periods of the benchmark model to 120. Figure 1 presents the average payoffs for each condition, calculated from period 20 onwards.[7]

[Figure 1 about here.]

The maximal payoff from each combination of c and b^{max} coincides with the optimal one since no κ_i level can guarantee a higher one at equilibrium. The main result of this benchmark model is probably that the optimal payoffs always occurred for κ_i close to b^{max}. More specifically, this value resulted from $\kappa_i = 0.95b^{max}$ for all parameter combinations. The minimum payoffs always corresponded to agent that did not harvest anything during the whole period and were obviously dependent on c (Tab. 1). All results presented from now on will be re-scaled to the interval corresponding to the minimum and the maximum payoff for each (c, b^{max}) combination.

[Table 1 about here.]

3.3 Base model results

In the base model, the resource overharvesting reached a point where only a few patches per round developed new vegetation, which was almost immediately cut down (Fig. 2). This was especially evident with $b^{max} = 10$, while the process proceeds somewhat slower when $b^{max} = 20$. However, even here, a clear trend was present, leading to the complete destruction of the resource.[8]

[Figure 2 about here.]

[6]A first exploratory implementation of the model was based on NetLogo 4.0.3 platform (Wilensky 1999). Subsequently, a C++ implementation was chosen to increase computational efficiency. Besides some minor differences due to the impossibility of reproducing the same routines exactly using two largely diverse codes, the models implemented on the two platform led to analogous results. All codes are available from the author upon request.

[7]All statistical analysis and plots have been produced using the R 2.8.0 platform (R Development Core Team 2008)

[8]Some tentative runs, encompassing a higher number of periods, showed indeed that the process of resource depletion of the resource ends up as in the $b^{max} = 10$ case.

Table 2 shows the average results for the final 100 periods of the simulations. Both the number of green patches and the total biomass were reduced to a small proportion of the initial quantities. Payoffs were far from the optimum (from 1/8 to 1/3, depending on the parameter conditions), showing that indiscriminate cutting does not pay off. Especially interesting was the effect of the selection mechanism on the β_i of agents. This variable was steadily reduced over time, although at a faster rate for lower maximal biomass, up to the point that it eventually reaches zero or even slightly negative values. This implies that the selection process leads to agents believing that the "right" state of the resource is one with no biomass on it. The κ_i value declines accordingly, up to the point when it almost reaches zero (for $b^{max} = 10$) or, at least, a small proportion of the maximal biomass (for $b^{max} = 20$). Note that in the latter case an increase of the number of periods also leads κ_i close to zero.

[Table 2 about here.]

All explored parameter configurations led to a qualitatively similar picture, even if increasing b^{max} at least reduced the speed of depletion. A closer look at the dynamic of the simulation showed a failed attempt of sustainable resource management. More specifically, the strong decline of the biomass in the initial periods of the simulation led to a temporary increase of the agent κ_i (Fig. 2e) and to a temporary improvement. However, this trend lasted for only 20–30 periods and the subsequent biomass decline was no longer reverted. This temporary inversion of the depletion trend was due to the different speed of change of the agents' beliefs. While agents can rapidly adapt their κ_i to any new situation, the β_i changes are driven by the selection process, which involves only one agent per period. Of course, a faster selection process would be easy to implement. Nevertheless, the interplay between hard-wired, slow-changing deep beliefs and easy-to-change operational procedures is probably important for real human beings. Moreover, the resulting dynamic is instructive in itself: while agents initially try to cooperate by increasing their κ_i, the very fact that they share an open-access resource where the dominant strategy (i.e. the strategy leading to the maximal payoff in the *current* period) is defection leads in the medium term to the selection of "rational" agents overharvesting their resource. This is the "tragedy of the commons" in all its harshness.

4 The institutional evolution model

4.1 *Inst* model definition

The institutional evolution model retains all the elements of the base model and introduces the effects of an endogenous management institution. The institution regulates the behavior of agents, but its evolution depends on their beliefs, which evolve as in the previous model.Nevertheless, the agents' beliefs no longer influence the agents' behaviors directly, which depends instead on the current institutional rule. This rule is modeled using a system-level variable K that determines the minimal level of b_{xy} needed to log a given patch. At the beginning of the simulation $K = 0$, i.e. the existing institution allows agents to cut any patch, but this value can subsequently change as a function of the agents' beliefs (Fig. 3).

[Figure 3 about here.]

Institutional change depends on whether the agents are satisfied with the current institution. An agent is unsatisfied when its current payoff is declining compared with the previous period (just as in the base model) *or* when the current institution is too far from its beliefs. More specifically, a "tolerance level" τ is now introduced into the model and an agent i become unsatisfied if $|K - \kappa_i| >$

τ. Note that even unsatisfied agents do not try to cheat by cutting patches with $b_{xy} \leq K$ (this point will be further discussed below). Nevertheless, they become willing to change the institution.

When the number of unsatisfied agents exceed 2/3 of the population, a new institution replaces the former one with the new K equal to the average κ_i of agents. The relatively high share of the population needed to change the institution reflects the fact that, in real CPR situations, institutional change is usually costly and a large consensus is needed to achieve this goal, at least when there is no subset of actors able to impose their rule on the majority (e.g. Ostrom 1990, 2005; Singleton and Taylor 1992).[9] Once the new institution is in place, the agents change their behavior accordingly. Besides the values of c and b^{max} derived from the base model, the effect of a low/high tolerance level, namely $\tau \in \{\frac{1}{3}b^{max}, \frac{2}{3}b^{max}\}$, will be explored as simulation parameter.

4.2 *Inst* model results

A summary of the results for the *Inst* model is shown in Table 3. For all c, b^{max} and τ levels, the introduction of the possibility of building a management institution increased the sustainability of resource exploitation. Both the number of green patches at equilibrium and the total system biomass were much higher than in the base model, ranging on average from 25% to 48% of the initial number of green patches and from 18% to 30% of the initial biomass. A c change did not influence those figures in a straightforward way. A higher value of b^{max} significantly increased both the number of green patches and the total biomas at equilibrium. It is worth noting that higher tolerance produced a lower number of green patches and lower biomass levels. This point will be discussed below.

[Table 3 about here.]

The average payoffs were also much higher than in the base model, ranging now around 90% of the optimum. This is a noticeable result. Thanks to institutional building, agents were able to cooperate and to considerably improve their payoffs; an outcome perfectly in line with the empirical litterature on CPRs. A closer look at the dynamics of the model (Fig. 4 and 5) helps us to understand better how this happened. Unlike the base case, in the *Inst* one the average β_i of agents tended to remain constant throughout the simulation. The average figure hinders the fact that in some runs the β_i tended to increase, while in others to decrease (although not reaching values as low as in the base model).

Despite the fact that the selection routine works as in the base model, it no longer determines the population dynamic that is now driven mainly by random drift. This happens because the management institution work tends to reduce the effects of the selection mechanism. The possibility of logging is no longer a matter of individual κ_i, but depends instead on the system-level value given by K. Agents with high β_i (and therefore high κ_i) no longer achieve payoffs significantly below the ones of agents with low β_i and consequently they are not excluded from the game. In other words, the management institution makes the selection mechanism less effective in positively selecting selfish agents: a process that spread noncooperation in the base model. As a result, cooperation can spread in the *Inst* model leading to a long term situation close to optimality.

[Figure 4 about here.]

[Figure 5 about here.]

[9] Note that decreasing the share of unsatisfied agents needed to change the current institution, e.g. by using a simply majority rule, leads to a greater number of institutional changes, and consequently to K closer to the average κ_i values of each period, without altering much the general outcome of the simulation.

The dynamics of the model showed increased κ_i in the first part of the simulation similar to that in the base model. However, in the *Inst* model this produced a corresponding rise of K: an outcome that stopped the β_i decline that occurred without the institution-building possibility. Once in place, the institutions tended to remain stable, even if some runs showed subsequent K changes and, in some rare cases, the system failed to reach a clear state of equilibrium (Fig. 4e and 5e).

As in the base model, the raise of fixed cost that agents have to pay in each round does not seem to change the situation much. Increasing the value of b^{max} tended to increase (even not to a large extent) the agent payoffs, expressed as proportion of the optimal one. What is especially interesting is the significantly *negative* effect on an increase of agents' "tolerance" . This is true both for payoffs (Wilcoxon rank sum test, $W = 30376$, $p < 0.001$ one sided) and for the two "ecological" variables, green patches and biomass ($W = 25986$, $p < 0.001$ one sided, and $W = 26599$, $p < 0.001$ one sided, respectively). The main reason behind this result is that, being more tolerant, agents were satisfied even with institutions relatively far from their own beliefs. As a consequence, once in place institutions changed less frequently and failed to adapt to further agent κ_i increases, even when a higher K could lead to better overall results.

5 Discussion

The main result of this simulations was probably the formal acknowledgment of the importance of institutional building. Indeed all conditions where agents were allowed to build a management institution produced much better outcomes than those in the base scenario, where agents could only rely on their individual beliefs in order to limit their harvesting. This is not a trivial result, since institutional building is driven by the evolution of agents' beliefs, which occurs in a competitive environment where agents achieving low payoffs are systematically removed from the system and replaced by more successful ones.

In the base model, the selection process tends to reduce the agents' β_i over time and consequently to increase their willingness to log any green patch. In the long run, this can only lead to a "tragedy of the commons" situation. This is no longer true when agents have the possibility of building a management institution. The rule that agents can implement simply states the minimum level of biomass that a given patch should possess in order to be logged. Initially this value is zero, corresponding to a rule allowing any action (Crawford and Ostrom 1995). This situation rapidly leads to a sharp decline of both the biomass in the system and the agents' payoffs. However, in the *Inst* model, the biomass decline changes the agents' beliefs, which produces stricter logging rules that increase both biomass and payoffs.

The fact that agents should "experience" a strong decline in biomass before changing their beliefs and becoming willing to implement management institutions is consistent with the findings of Janssen and Ostrom (2006) model. Unlike artificial agents, real individuals can forecast consequences of their actions and, to some extent, they may anticipate resource depletion by limiting their harvesting levels before reaching dangerous conditions. Nevertheless, just like the artificial agents, they rarely appear to seize this opportunity. Usually, individual have to "learn" about the fragility of their natural resources by depleting them before starting to reduce and regulate their harvesting levels (e.g. Berkes and Folke 1998; Berkes et al. 2003; Diamond 2005; Johannes 2002). This dynamic was perfectly reproduced by our model.

Resource depletion and subsequent institutional implementation also occurred in Janssen and Ostrom's (2008) "spatial commons" experiment. In the first round, participants rapidly overharvested the common resource. Subsequently, the introduction of a communication phase allowed them to devise and implement informal institutions that largely improved their results. The similarities with our simulations are interesting, starting with the incapacity of both real individuals and artificial agents to overcome the commons dilemma without shared rules of behavior.

Nevertheless, once they reached an agreement on a harvesting rule, both succeeded in sustainably managing the resource. Our model was not intended to simulate the experimental results and included no data from its participants. Still, the similarities in the payoffs achieved in the two conditions are noticeable. Janssen and Ostrom report that, before the first communication period, the average earnings of the participants were, on average, about 1/4 of the potential amount. Similarly, in the base model our agents "earned" between 1/8 to 1/3 of the optimum, depending on the value of c and b^{max}. In the last round of the experiment, after two communication periods, participants approximately doubled the number of collected tokens and, consequently, their earnings (Janssen and Ostrom 2008, Tab. 2). Our agents, enjoying a formal institution that did not allow for any rule breaking, performed even better and achieved at least 70% of the optimum.

Unfortunately, we do not know much about the beliefs and the mental models of the subjects that participated in the experiment. Nevertheless, by explicitly modeling them the simulation suggests that a situation where agents possess a general idea of the desirable status of the commons and a more specific belief regarding the strategy that can to lead to this is sufficient to explain the observed dynamics, but only when subjects are free to devise and implement management institutions. It is worth noting that this conclusion is also consistent with the picture of institutional building and change offered by North (2005). Future experiments may be explicitly aimed at measuring the beliefs of real individuals in a situation akin to our model. The participants' κ_i could easily be measured by letting them vote for a cutting-threshold rule, while their β_i can be inferred from questionnaire answers. This would be a crucial test for the validity of our model and, at the same time, an extremely valuable source of data for modeling agents that behave closer to real individuals.

An interesting side result of the model was the lower payoffs achieved in the *Inst* condition by increasing the tolerance level. This is somewhat counter-intuitive, since higher tolerance is often thought to be positively linked with cooperation. However, in this case higher tolerance simply meant that agents were already satisfied with the institutions that produced outcomes far from optimum. In other word, constant research for better situations driven by low-tolerant, unsatisfied agents is important to allow constant tuning-up of institutions, making them better adapted to their working environment.

The current model is thought to be a starting point for future research. It would be especially useful to proceed in two parallel directions by relaxing two of its most demanding assumptions, i.e. the absence of social influence and the impossibility for agents to break institutional rules. Unlike real settings, our agents changed their beliefs individually without considering what other agents do or "believe". Introducing some form of social influence, e.g. on belief formation, would allow selective learning and shared improvement of agents' "knowledge". This is an important part of human cultural systems (Richerson and Boyd 2005) and would greatly improve the empirical plausibility of our model. It is probable that social influence will lead to a faster rate of belief update and therefore of institutional change, especially in the initial part of the simulation. This will probably make the model behavior more subject to rapid shifts, better reproducing the dynamics observed in real settings.

In our model, agents always conform to institutional rules. This is obviously not true in reality, where cheating and rule breaking is widespread when not controlled. Relaxing the assumption of a perfectly controlling institution is therefore important in making the model closer to empirical situation. For instance, it is possible to imagine that agents holding beliefs in contrast with the institutional rule may become willing to cheat by logging forbidden patched (besides trying to change the institution itself). This will put in question the management capacity of the institution and lead to the need for a monitoring and sanctioning system, whose effectiveness may also depend, at least partially, on the agents' beliefs. Janssen and Ostrom (2006) included the possibility of mutual monitoring. More specifically, monitoring actions depended on a parameter specific for each agent, which was hence subject to evolutionary dynamics. Instead of assuming a separate monitoring propensity, it may be more realistic to relate it to the agents' beliefs. More specifically, just as agents are more prone to break the rules further their beliefs are from the current institution, they may be more willing to monitor compliance the more their beliefs are closer to the institution. In other words, agents may monitor with a probability that is inversely proportional to the distance between κ_i and K. This will lead to an interesting situation where agents with beliefs far from the current institution will try to cheat, while the ones holding beliefs consistent with it will try to control logging behavior. At the same time, the beliefs of both populations will participate in the institutional changing process. Due to the large number of interactions between these factors, it is difficult to forecast the overall behavior of the system, but it is likely that its implementation will lead to significant insights regarding the functioning of real institutional schemes.

These improvements will surely make the model more realistic. Nevertheless, the current version already highlights some basic dynamics of the relationship between individuals and institutions. What is especially relevant is the formal display of the importance of the feedback loop existing between beliefs and institutional rules. Institutions for CPR management depend on the beliefs of actors exploiting the resources. At the same time, by affecting individuals' actions and outcomes, institutions actively participate in changing their beliefs in a cycle that can be adaptive, but also highly destructive. Being aware of this cycle is crucial to better understand the dynamics of real institutions and therefore to improve our capacity of recognize what is going on both in the management of shared natural resources and in other situations characterized by one or more underlying social dilemmas.

Aknowledgements I wish to thank Vincent Buskens, Marco Janssen, Lucia Tamburino and Werner Raub for their highly valuable comments on a previous version of this paper. The simulation work was completed during a month when I was a visitor at the Sociology Department of Utrecht University. I am grateful to Vincent Buskens and Werner Raub for inviting me. I also gratefully acknowledge all the suggestions offered by the participants to a seminar held at the same department during my visit. I would also like to thank Robert Coates for his linguistic revision.

References

Barclay, Pat 2006. 'Reputational Benefits for Altruistic Punishment.' *Evolution and Human Behavior* 27: 325–344.

Berkes, Fikret 1998. 'Indigenous Knowledge and Resource Management Systems in the Canadian Subartic.' In *Linking Social and Ecological Systems: Management Practicies and Social Mechanisms for Building Resilience*, eds. Fikret Berkes and Carl Folke, pp. 98–128. Cambridge: Cambridge University Press.

Berkes, Fikret, Johan Colding and Carl Folke, eds. 2003. *Navigating Social-Ecological Systems: Building Resilinece for Complexity and Change*. Cambridge: Cambridge University Press.

Berkes, Fikret and Carl Folke, eds. 1998. *Linking Social and Ecological Systems: Management Practicies and Social Mechanisms for Building Resilience*. Cambridge: Cambridge University Press.

Boyd, Robert, Herbert Gintis, Samuel Bowles and Peter J. Richerson 2003. 'The Evolution of Altruistic Punishment.' *Proceedings of the National Academy of Science USA* 100(6): 3531–3535.

Cardenas, Juan-Camilo 2000. 'How Do Groups Solve Local Commons Dilemmas? Lessons from Experimental Economics in the Field.' *Environment, Development ans Sustainability* 2: 305–322.

Cardenas, Juan-Camilo 2003. 'Real wealth and experimental cooperation: experiments in the field lab.' *Journal of Development Economics* 70: 263–289.

Cardenas, Juan-Camilo, John K. Stranlund and Cleve E. Willis 2000. 'Local Environmental Control and Insitutional Crowding-Out.' *World Development* 28(10): 1719–1733.

Carlsson, Lars 2003. 'The strategy of the commons: history and property rights in central Sweden.' In *Navigating Social-Ecologica Systems: Building Resilience for Complexity and Change*, eds. Fikret Berkes, Johan Colding and Carl Folke, pp. 116–131. Cambridge: Cambridge University Press.

Casari, Marco and Charles R. Plott 2003. 'Decentralized Management of Common Property Resources: Experiments with a Centuries-Old Institution.' *Journal of Economic Behavior and Organization* 51: 217–247.

Castillo, Daniel and Ali Kerem Saysel 2005. 'Simulation of common pool resource field experiments: a behavioral model of collective action.' *Ecological Economics* 55: 420–436.

Crawford, Sue E.S. and Elinor Ostrom 1995. 'A Grammar of Institutions.' *American Political Science Review* 89: 582–600.

Dawes, Robyn M. 1973. 'The Commons Dilemma Game: An *N*-Person Mixed-Motive Game with a Dominating Strategy for Defection.' *ORI Research Buletin* 13: 1–12.

de Quervain, Dominique J.-F., Urs Fischbacher, Valerie Treyer, Melanie Schellhammer, Ulrich Schnyder, Alfred Buck and Ernst Fehr 2004. 'The Neural Basis of Altruistic Punishment.' *Science* 305: 1254–1258.

Deadman, Peter, Edella Schlager and Randy Gimblett 2000. 'Simulating Common Pool Resource Management Experiments with Adaptive Agents Employing Alternate Communication Routines.' *Journal of Artificial Societies and Social Simulation* 3(2): 2.

Diamond, Jared 2005. *Collapse. How Societies Choose to Fail or Succeed.* New York: Viking Press.

Fehr, Ernst and Urs Fischbacher 2003. 'The Nature of Human Altruism.' *Nature* 525: 785–791.

Fehr, Ernst and Simon Gächter 2000. 'Cooperation and Punishment in Public Goods Experiments.' *American Economic Review* 90: 980–994.

Fehr, Ernst and Simon Gächter 2002. 'Altruistic Punishment in Humans.' *Nature* 415: 137–140.

Gilbert, Nigel 2008. *Agent-Based Models*. London: Sage.

Gintis, Herbert, Samuel Bowles, Robert. Boyd and Ernst. Fehr 2003. 'Explaining Altruistic Behavior in Humans.' *Evolution and Human Behavior* 24: 153–172.

Grant, W.E. and P.B. Thompson 1997. 'Integrated ecological models: simulation of socio-cultural constraints on ecological dynamics.' *Ecological Modelling* 100: 43–59.

Hardin, Garrett 1968. 'The Tragedy of the Commons.' *Science* 162(3859): 1243–1248.

Henrich, Joseph, Richard McElreath, Abigail Barr, Jean Ensminger, Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwins Gwako, Natalie Henrich, Carolyn Lesorogol, Frank Marlowe, David Tracer and John Ziker 2006. 'Costly Punishment Across Human Societies.' *Science* 312: 1767–1770.

Janssen, Marco A., Robert L. Goldstone, Filippo Menczer and Elinor Ostrom 2008. 'Effect of Rule Choice in Dynamic Interactive Spacial Commons.' *International Journal of the Commons* 2(2): 288–312.

Janssen, Marco A. and Elinor Ostrom 2006. 'Adoption of a New Regulation for the Governance of Common-Pool Resources by a Heterogeneous Population.' In *Inequality, Cooperation and Environmental Sustainability*, eds. Jean-Marie Baland, Pranab Bardhan and Samuel Bowles, pp. 60–96. Princeton: Princeton University Press.

Janssen, Marco A. and Elinor Ostrom 2008. 'TURFs in the Lab: Institutional Innovation in Real-Time Dynamic Spatial Commons.' *Rationality and Society* 20(4): 371–397.

Johannes, R.E. 2002. 'Did Indigenous Conservation Ethics Exist?' *Traditional Marine Resource Management and Knowledge Information Bulletin* 14: 3–7.

Lam, Wai Fung 1998. *Governing Irrigation Systems in Nepal: Institutions, Infrastructure, and Collective Action*. Oakland: ICS Press.

Milinski, Manfred, Dirk Semmann and Hans-Jürgen Krambeck 2002. 'Reputation Helps Solve The Tragedy of the Commons.' *Nature* 415: 424–426.

North, Douglass C. 2005. *Understanding the Process of Economic Change*. Princeton, NJ: Princeton University Press.

Ostrom, Elinor 1990. *Governing the Commons. The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

Ostrom, Elinor 2005. *Understanding Institutional Diversity*. Princeton: Princeton University Press.

Ostrom, Elinor, Thomas Dietz, Nives Dolšak, Paul C. Stern, Susan Stonich and Elke U. Weber, eds. 2002. *The Drama of the Commons*. Washington DC: National Academy Press.

Ostrom, Elinor, Roy Gardner and James Walker 1994. *Rules, Games, and Common-Pool Resources*. Ann Arbor: The University of Michigan Press.

Pinkerton, Evelyn 1998. 'Integrated management of a temperate montane forest ecosystem through wholistic forestry: a British Columbia example.' In *Linking Social and Ecological Systems: Management Practicies and Social Mechanisms for Building Resilience*, eds. Fikret Berkes and Carl Folke, pp. 363–389. Cambridge: Cambridge University Press.

R Development Core Team 2008. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. http://www.R-project.org.

Richerson, Peter J. and Robert Boyd 2005. *Not by Genes Alone: How Culture Transformed Human Evolution*. Chicago: The University of Chicago Press.

Singleton, Sara and Michael Taylor 1992. 'Common Property, Collective Action and Community.' *Journal of Theoretical Politics* 4(3): 309–324.

Smil, Vaclav 2002. *The Earth's Biosphere: Evolution, Dynamics, and Change*. Cambridge: The MIT Press.

Stern, Nicholas 2007. *The Economics of Climate Change*. Cambridge: Cambridge University Press.

Tang, Shui Yam 1992. *Institutions and Collective Action: Self-Governance in Irrigations*. San Francisco: ICS Press.

Volk, Tyler 2008. *CO_2 Rising: The World Greatest Environmental Challenge*. Cambridge, MA: The MIT Press.

Wilensky, Uri 1999. 'Netlogo.' Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. http://ccl.northwestern.edu/netlogo/.

Figures

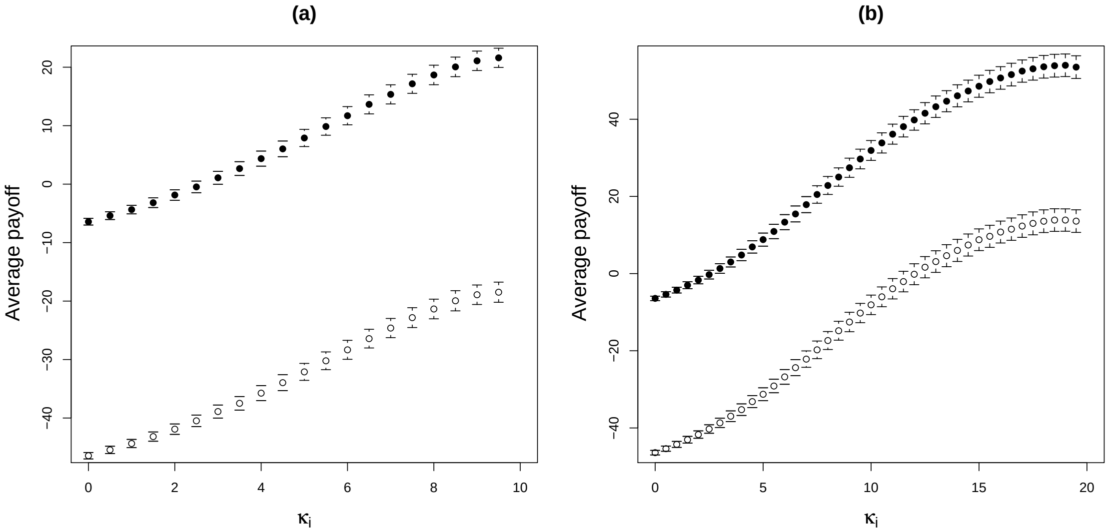

Figure 1: Average payoffs and standard deviations in 50 runs for each condition of the benchmark model. The (a) panel presents the results for $b^{max} = 10$, the (b) panel for $b^{max} = 20$. Conditions with $c = 1$ are plotted using a filled circle, conditions with $c = 5$ using an empty one.

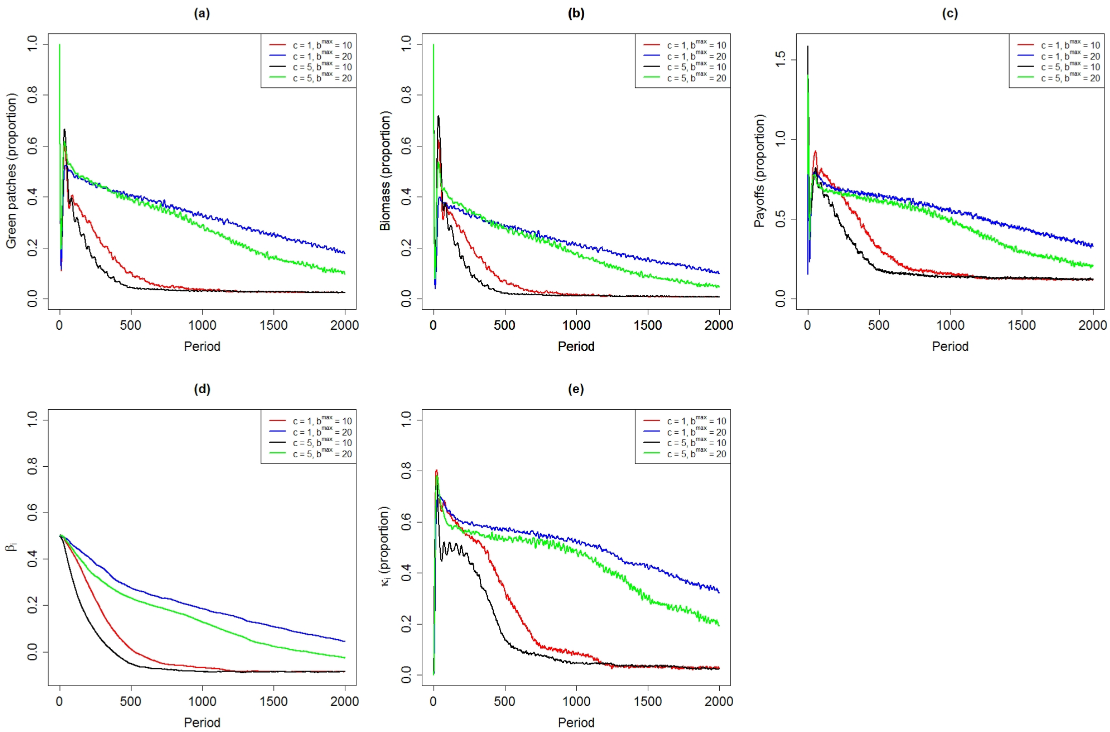

Figure 2: Average dynamics in 50 repetitions per parameter configuration of the base model. Green patches and total biomass are expressed as proportion of the initial conditions; payoffs are are expressed as proportion of the optimum; κ_i is expressed as proportion of b^{max}.

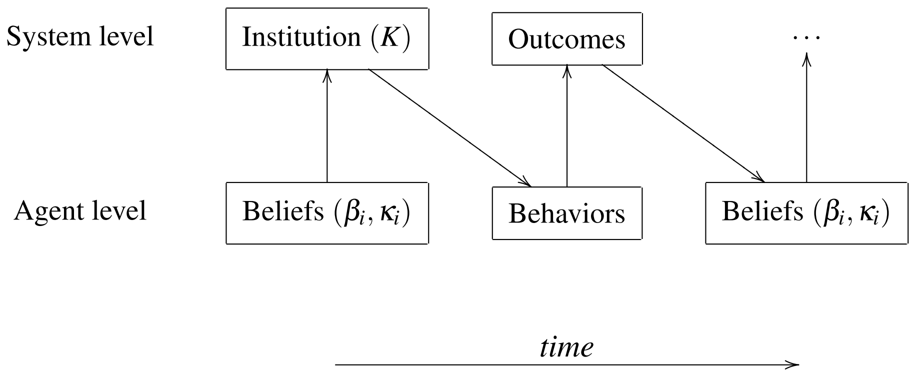

Figure 3: Relation among beliefs, institutions and outcomes in the *Inst* model.

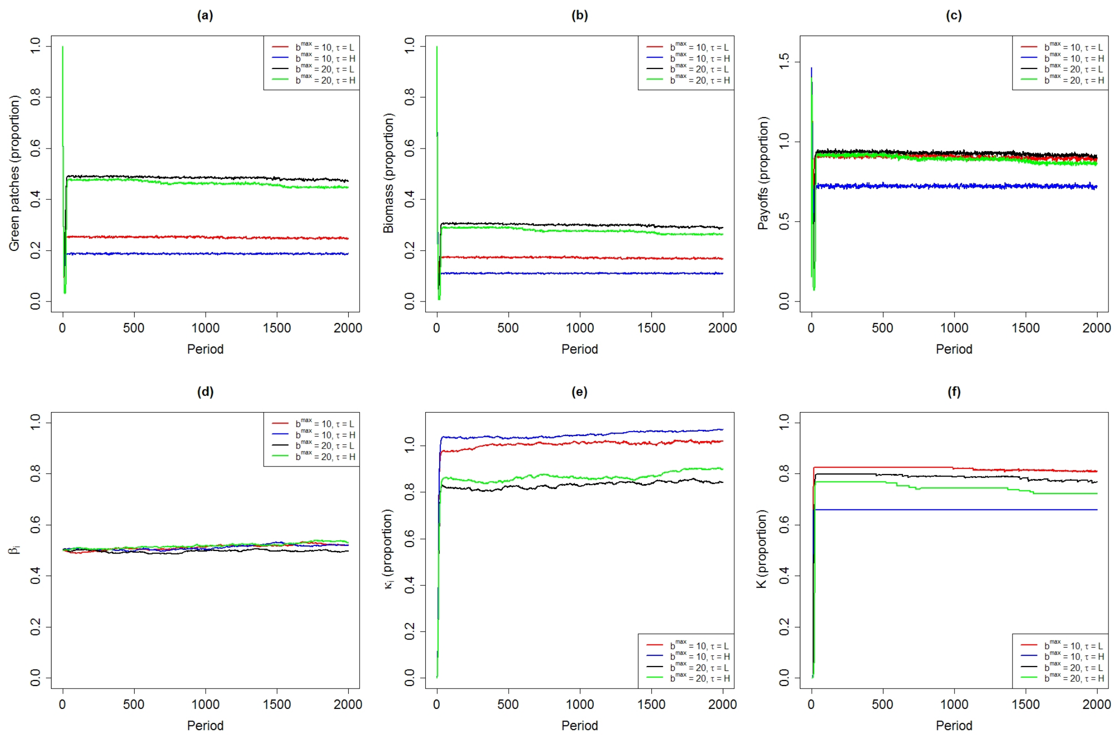

Figure 4: Average dynamic in 50 repetitions per parameter configuration of the *Inst* model with $c = 1$. Green patches and total biomass are expressed as proportion of the initial conditions; payoffs are are expressed as proportion of the optimum; κ_i and K are expressed as proportion of b^{max}.

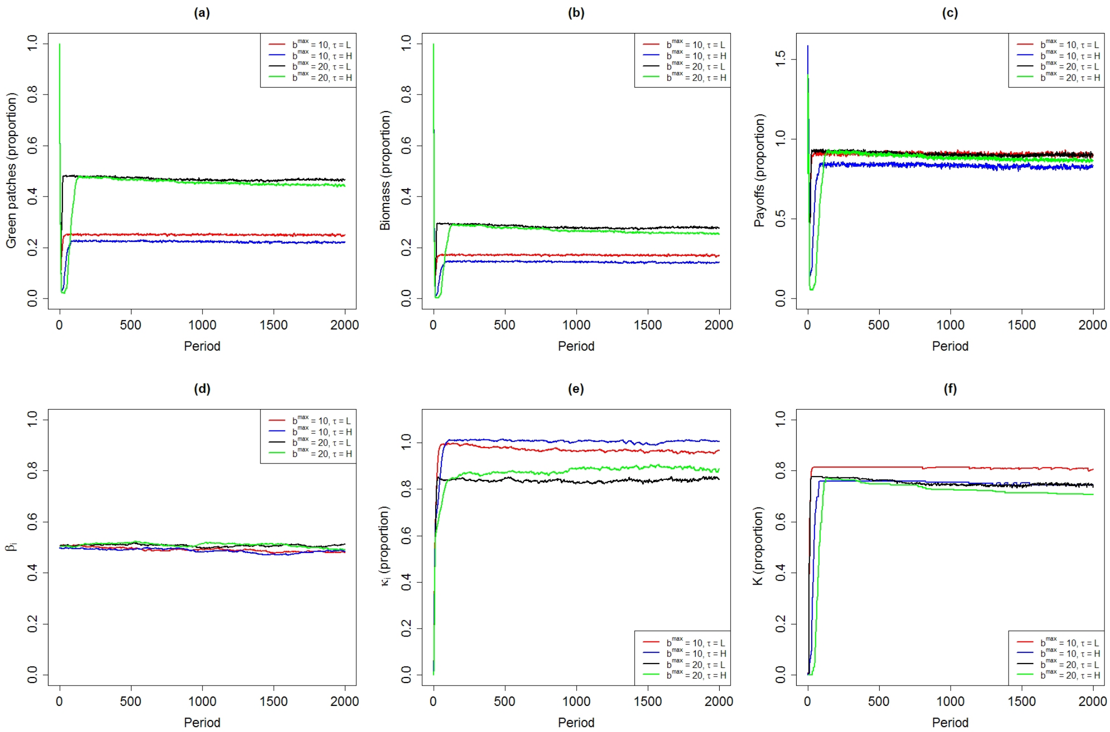

Figure 5: Average dynamic in 50 repetitions per parameter configuration of the *Inst* model with $c = 5$. Green patches and total biomass are expressed as proportion of the initial conditions; payoffs are are expressed as proportion of the optimum; κ_i and K are expressed as proportion of b^{max}.

Tables

c	b^{max}	min.	max.
1	10	-10	21.6
1	20	-10	54.0
5	10	-50	-18.5
5	20	-50	13.9

Table 1: Minimum and maximum average payoffs for each condition.

c	b^{max}	Green patches	Total biomass	Average payoffs	Average β_i	Average κ_i
1	10	0.025	0.009	0.123	-0.086	0.029
		(0.012)	(0.008)	(0.042)	(0.081)	(0.094)
1	20	0.189	0.112	0.345	0.051	0.339
		(0.132)	(0.092)	(0.221)	(0.131)	(0.233)
5	10	0.026	0.009	0.124	-0.086	0.025
		(0.016)	(0.012)	(0.052)	(0.088)	(0.096)
5	20	0.109	0.054	0.213	-0.021	0.212
		(0.106)	(0.067)	(0.190)	(0.123)	(0.234)

Table 2: Average results for the last 100 periods in 50 repetitions per parameter configuration of the base model. Green patches and total biomass are expressed as proportion of the initial conditions; payoffs are expressed as proportion of the optimum; κ_i is expressed as proportion of b^{max}. Standard deviations are in parentheses.

c	b^{max}	τ	Green patches	Tree biomass	Average payoffs	Average β_i	Average κ_i	Average K
1	10	L	0.248	0.182	0.895	0.520	1.018	0.810
			(0.025)	(0.024)	(0.079)	(0.162)	(0.133)	(0.070)
1	10	H	0.187	0.119	0.721	0.520	1.068	0.657
			(0.021)	(0.021)	(0.067)	(0.100)	(0.065)	(0.060)
1	20	L	0.475	0.300	0.912	0.495	0.841	0.770
			(0.048)	(0.044)	(0.084)	(0.142)	(0.181)	(0.080)
1	20	H	0.448	0.272	0.867	0.535	0.901	0.721
			(0.073)	(0.067)	(0.130)	(0.123)	(0.134)	(0.117)
5	10	L	0.249	0.182	0.903	0.480	0.960	0.806
			(0.013)	(0.013)	(0.037)	(0.124)	(0.166)	(0.033)
5	10	H	0.220	0.152	0.825	0.486	1.008	0.744
			(0.024)	(0.023)	(0.076)	(0.123)	(0.155)	(0.063)
5	20	L	0.465	0.287	0.901	0.509	0.848	0.745
			(0.034)	(0.035)	(0.057)	(0.116)	(0.153)	(0.061)
5	20	H	0.444	0.263	0.866	0.492	0.882	0.707
			(0.041)	(0.044)	(0.069)	(0.132)	(0.186)	(0.075)

Table 3: Average results for the last 100 periods in 50 repetitions per parameter configuration of the *Inst* model. Green patches and total biomass are expressed as proportion of the initial conditions; payoffs are expressed as proportion of the optimum; κ_i and K are expressed as proportion of b^{max}. Standard deviations are in parentheses.