APPENDIX

Description of the agent based model.

The baseline model: open access

Variables utilized in every version of the model are described in detail in Table A.1. In the baseline model, the state of the world has the following features. 100 agents operate on a regular lattice of degree $l = 8$. The lattice has the structure of an m x m toroidal surface, with $m = 50$. The surface is divided in patches. Each patch is a forest area that can be logged in one round. Patches have the attribute *trees*, which belongs to the [0, *max-tree-growth*] interval. It represents the total tree biomass present in a given moment in the patch and, if its value is higher than zero, it takes a green colour. *max-tree-growth* is the maximum possible level of biomass per patch, and it is controlled by an external slider. This choice is made in order to have the possibility to represent different kinds of forest, containing more or less vegetation. At the beginning of the simulation the forest is mature, with the value of *trees* randomly distributed in the [½ *max-tree-growth, max-tree-growth*] interval. If not logged, biomass in each patch grows at the fixed rate of 0.5 units per round up to the point where they reach *max-tree-growth*. If the patch is empty, biomass regrows with a probability depending on the state of the neighbouring patches, according to the function

*growth-prob * ((living-neighbours + 1) / 9)*

where *growth-prob* is the basic regrowth probability and has the value of 0.05, *living-neighbours* is the number of non-empty neighbour patches and 9 means 8 + 1, with 8 being the number of neighbour patches. This concept of neighbourhood is analogous to that used in Janssen and Ostrom (2006). This means that if all the neighbour patches are green, the regrowth probability of an empty patch is 0.05, while if it is surrounded by empty patches the probability will be 0.005555. This function is used by Janssen et al. (2008) for the "spatial commons experiments". One difference is that here the regrowth probability is strictly above zero because of mechanisms, assumed to be present, such as the natural recovery capacity due to seed conservation in the soil and seed dispersion by animals.

Each agent has three features. The first is called *reference-trees* and represents a subjective idea about the fraction of the initial tree biomass that should be ideally conserved. This symbolizes a cognitive model that each agent has about the "right" state of the world. At the beginning of each round this is drawn randomly from a normal distribution having mean 0.5 and standard deviation 0.25 and it remains subsequently constant. The second is *minimal-cut* and represents a preference about the minimal level of tree biomass that a patch should have in order to be logged. For every agent it is equal to zero when agents enter the game. This conditions means that at the beginning of the game loggers believe that they can always cut. This variable will update during the simulation according to the state of the forest and to the economic profit of the agent. I will describe this mechanism later. The third feature is the *payoff* : it is assumed that when an agent logs a patch he earns a monetary profit. At the beginning of the simulation payoff is equal to zero for every agent.
The execution of the model operates as follows. Each simulation covers 2000 periods. Each period has 10 rounds. One round corresponds to one "tick" in NetLogo. In every round agents move within the simulated forest and each of them pays a fixed monetary charge. This variable is called

cost and it is controlled by an external slider, so that it is possible to assume high or low costs for displacement and logging in general. When an agent arrives on a patch he has to decide if logging or not. If the condition

[trees] of patch-here > minimal-cut

is true, than the agent cuts and the quantity of *trees* is added to his *payoff*. If the condition is not true, the agents controls if any of the neighbor patches has biomass above that threshold. If he finds any, he moves on one of these patches, pays the fixed charge and realizes no earning in the current round. If none of the patches has sufficient biomass, the agents move randomly and earns no profits. The *payoff* of each agent is given by the difference between his earnings and costs.

At the end of each period there is an update of the subjective preferences of each agent about the right threshold of biomass quantity that should be present on a patch in order to decide whether to log or not. If the current *payoff* is higher or equal to that of the previous round, the agent maintains his *minimal-cut*. This means that if the agent is satisfied about his profit from the logging activity, he has no reason to modify his opinion about the importance of preserving part of the forest intact. Otherwise, the agent changes his *minimal-cut* with a probability *q*:

q (payoff - old-payoff) / (abs payoff + abs old-payoff)

where *old-payoff* is the payoff of the previous round and *abs* means “absolute value”. A random extraction determines if the agent will actually change his *minimal-cut*. If this happens his *minimal-cut* is modified according to his *reference-trees.* More specifically, if the total number of green patches is higher than the fraction that should ideally be conserved according to the agent’s vision (*reference-tree)*, the agent decreases his *minimal-cut* by a random value in the interval [0,9]. If the contrary happens, that is to say, if the total number of green patches is lower than the agent’s *reference-tree*, he increases his *minimal-cut* by the same amount. The meaning behind is that agents facing a payoff reduction become unsatisfied and are motivated to modify their subjective values and, therefore, their behavior. If the share of the biomass left is lower than the agent’s *reference-trees* (which indicates the share of the forest that should be conserved according to the agent’s vision), he attributes the earning reduction to an excessive cutting and will increase his own *minimal-cut*, becoming more environmentalist, and viceversa. The interplay between slow-changing deep values (*reference-tree*) and easy-to-change operational procedures (*minimal-cut*) reflects reality.

At the end of the values update, a selection process among the agents takes place, through the bankruptcy of unsuccessful agents. First, one of the agents with the highest period payoff and one with the lowest payoff in the period are selected. Secondly, a copy of the former (i.e. its *reference-trees*, while *minimal-cut* is always equal to zero when a new agent enters the game) replaces the latter. There is a one per cent probability of “mutation”, that is to say “copy errors” or new entrants with innovative values. At the end of the selection process all payoffs are put equal to zero and a new period starts.

The results of the open-access version of the model show a complete depletion of the forest and very low payoffs for the agents. Both the number of green patches and the total biomass are reduced to a small proportion of the initial quantities. The dynamics of the socio-ecological system shows that a strong decline of the biomass in the very first period leads to a temporary increase of the agents' *minimal-cut*. However this lasts only for a few periods. Subsequently both payoffs and forest indicators go to zero. This temporary inversion of the depletion trend happens because of the different speed of change of the agents' values. While agents can quickly adapt their *minimal-cut* to the new situation in every period, changes in *reference-trees* are driven by the selection process, which involves only one agent per period. At the end of the simulation also the agents' *minimal-cut* and *reference-trees* go to zero. This implies that the selection process leads to the prevalence of the agents with higher earnings, which, in turn, are agents believing that the "correct" state of the forest is one with no trees on it. Since we are in an open access situation, with every agent deciding his behaviour only according to his personal values, agents with a low *minimal-cut* will log always more (Bravo 2011). At the end the typical tragedy of the commons occurs, with depletion of the forest.

Endogenous institution

In this version of the model one new variable is introduced. At a certain point of the process, agents agree on a shared cutting rule. As explained earlier, an agent is unsatisfied when his current payoff is lower than the one of the previous round. When the number of unsatisfied agents exceeds 2/3 of the population, the mean of the *minimal-cut* of each agent forms the new variable *current-institution*. This variable indicates the biomass threshold that a patch should contain in order to be logged and this cutting rule becomes compulsory for the whole community. An agent determines his behaviour on the basis of the shared *current-institution* and not anymore on the basis of his personal *minimal-cut*. At this point an additional criterion for agent dissatisfaction is in place: the distance between *current-institution* and *minimal-cut*. Therefore if an agent faces a payoff reduction or if his personal environmental values are too far from the institutional rule in place, he is unsatisfied. Again, when a high number of unsatisfied agents is reached the institutional rule is updated according to the mean of the agents' new *minimal-cut*. This new institutional rule will determine agents' behaviour.

The results of this model version show much higher levels of total biomass and of earning of the agents, if compared with the open access situation. These results are in line with the empirical literature (Bravo 2011) and show that an institution endogenous to the community may solve the tragedy of the commons. Observing the dynamics of the model it is possible to understand how these results emerged. Unlike the open access model, the average *reference-tree* of the agents remains constant until the end of the simulation. The establishment of the management institution diminishes the effect of the selection mechanism, even if this is the same than in the previous version of the model. Like in the previous model version, at the beginning of the simulation there is an increase of the average *minimal-cut*. However, here this leads to an increase of the shared institution and all the agents will cut less. The endogenously created institution makes the selection mechanism less effective in allowing the survival of more selfish characters among the agents and the defection of the others. The logging decision is no longer matter of personal *minimal-cut* of the agents, but depends on the system level *current-institution*. More environmentalist agents (with higher *minimal-cut*) no longer reach payoffs much lower than the "selfish" ones and therefore they are not excluded from the simulation. This happens because the cutting behaviour does not fluctuate anymore following the heterogeneous *minimal-cut*, therefore the payoffs are more stable as well and the selection mechanism less efficient. This means also that agents with a more "forest-friendly" vision (*reference-trees*) are not so easily excluded by the simulation (Bravo 2011).

Exogenous institution

In this version of the model I represent a situation in which an exogenous entity decides on the cutting threshold. Therefore the variable *current-institution* is not anymore made by the mean of agents' *minimal-cut*. It is now determined by an external slider controlled by the researcher. Its range goes from zero (which means that a patch should contain at least zero biomass in order to be logged) to [*max-tree-growth* – 0.5] (which means that a patch should contain at least its maximum biomass level minus 0.5 units in order to be logged). The meaning behind is that, in the first case, an agent is always allowed to cut, and in the latter case an agent is not allowed to cut at all. In this way it is possible to observe both situations of a "strict" cutting rule (with the *current-institution* slider set to 9) and of a "soft" cutting rule (with the *current-institution* slider set to 2).

When the cutting rule is set to level 9, it represents a classical situation of "fortress" style protected area, where resource extraction is almost completely forbidden (Hayes 2006, Campbell and Vainio-Mattila 2003). The important difference with reality is that at this stage we still assume that cheating does not exists and that every agent follows the imposed cutting rule. In line with common sense intuitions, the results of this simulation show a good state of the forest, but a very low level of agents payoff. Otherwise, if we shift the cutting level to a "soft" rule (level 2), according to which it is possible to cut a high number of patches, the forest is depleted and the payoffs become even negative, because after a certain number of periods agents do not find any more trees to cut. This outcome is similar to that in the open access scenario.

Cheating

At this point the possibility of violating the cutting rules is introduced in the model. In both scenarios, with endogenous and with exogenous institution, agents log a patch either if the cutting rule is fulfilled, or if they are unsatisfied. Again, an agent is unsatisfied either if his current payoff is lower than the one of the previous round, or if the cutting rule is too far away from his personal vision.

The only difference between the two settings is that for the endogenous institution version, the *current-institution* is the mean of the individual *minimal-cut*, while in the exogenous institution version it is determined by the external slider. Enforcement has not been introduced yet, therefore the impact of the possibility of violating the rule is very strong, regardless of what kind of institution is in place: in both cases the forest is completely logged and the payoffs of the agents are negative.

Enforcement

At this stage I introduce settings with rule violation and enforcement, regarding both kinds of institutions, exogenous and endogenous. The enforcement intensity is again determined in both cases by an external slider, since I assume that it depends on the availability of resources of the institution in charge, being endogenous or exogenous to the community, and it is not dependent on the performance of the participatory conservation experience. Additionally, agents now face a random probability to be effectively caught after the violation of the rule, as it is shown in the code.

When the agent enters a patch he logs it if the rule satisfaction condition holds. Otherwise he moves when satisfied, or he logs anyway if he is unsatisfied. If this latter case happens, if the probability to be caught is higher than the enforcement level effectively in place, the agents *dies*, which means he disappears from the next simulation rounds.

Literature cited

Bravo, G. 2011. Agents' beliefs and the evolution of institutions for common-pool resource management. *Rationality and Society* 23(1): 117-152.

Campbell, L. M. and A. Vainio-Mattila. 2003. Participatory development and community-based conservation: opportunities missed for lessons learned? *Human Ecology* 31: 417-436.

Hayes, T. M. 2006. Parks, people, and forest protection: an institutional assessment of the effectiveness of protected areas. *World Development* 34(12) : 2064–2075.

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

Janssen, M. A., and E. Ostrom E. 2006. Adoption of a new regulation for the governance of common-pool resources by heterogeneous population. Ch. 4. Pages 60-96 in J. M. Baland, P. Bardhan, and S. Bowels, editors. Inequality, cooperation and environmental sustainability. Princeton University Press, Princeton, USA.

Table A.1. Variables used in the simulation in Netlogo.

Variables names in NetLogo	Features	Explanation
max-pxcor (in "settings")	50	Maximum x coordinate for patches of the mxm toroidal surface
max-pycor (in "settings")	50	Maximum y coordinate for patches of the mxm toroidal surface
trees	Belongs to [0, bmax]. At the beginning of the simulation it is randomly distributed in [1/2 bmax, bmax]	Tree biomass present in a given moment on the patch. x y are the spatial coordinates
max-tree-growth	slider	Maximum possible level of biomass per patch
pcolor	60 + 5 * (trees / max-tree-growth) = the more trees the patch has, the lighter it is.	Colour of the patch
		Re-growing probability of an empty patch
living-neighbors		count neighbors with [trees > 0]
growth-prob	0.05	Basic probability of re-growth in p = p* (N+1)/(*k*+1)
reference-trees	At the beginning of each run, it is drawn randomly from a normal distribution with mean 0.5 and standard deviation 0.25. It remains constant.	Individual belief of each agent: fraction of the initial tree biomass that should ideally be conserved.
minimal-cut	= 0 when agents enter the game; it is updated frequently.	Individual belief of each agent: minimal level of tree biomass that a patch should have in order to be logged. If it is low, it means you can cut all. If it is high, it means you can not cut. Level of cutting that is able to maintain the actual tree biomass at the desired level.
payoff	= 0 at the beginning of each period; after it depends on agent's actions.	Agent's earning when he logs the patch. In every round: set payoff payoff - cost If he logs: set payoff payoff + [trees] of patch-here
old-payoff		Payoff of the previous round
cost	Slider: [1, 10]	Fixed cost that the agent pays at every round.
q	let q (payoff - old-payoff) / (abs payoff + abs old-payoff)	Probability of changing *minimal-cut* if the *payoff* of the current round is lower than the one of the previous round.
"Total Biomass" (in plots)	sum [trees] of patches	Total biomass in the initial period (sum of bxy) Total biomass in the current period
"Green Patches"(in plots)	count patches with [trees > 0]	

current-institution	At the beginning of the simulation is =0 After, is the average of the agents' *minimal-cut*	Minimum level of tree biomass that a patch should have in order to be logged. If it is low, it means you can cut all. If it is high, it means you can not cut.
current-institution	Slider: the maximum value of the slider is (max-tree-growth – 0.5). When you do setup, netlogo calculates it. After doing setup, you decide the value of the slider.	Exogenously imposed institutionMinimum level of tree biomass that a patch should have in order to be logged. If it is low, it means you can cut all. If it is high, it means you can not cut.
tolerance-threshold	ifelse high-tolerance = true [set tolerance-threshold (2 * max-tree-growth) / 3] [set tolerance-threshold max-tree-growth / 3]	Tolerance level
unsatisfied	count turtles with [abs (minimal-cut - current-institution) > tolerance-threshold or payoff-satisfaction = 0]	
payoff-satisfaction	ask turtles with [payoff < old-payoff] [let q (payoff - old-payoff) / (abs payoff + abs old-payoff) if (- random-float 1) > q [set payoff-satisfaction 0	At the end of each period each agent checks its payoff satisfaction. If the current payoff is lower than the previous one, he changes its *minimal-cut* with probability *q*. A random extraction determines whether he actually changes its belief.
initial-loggers	Slider: [0, 100]	Initial number of agents
enforcement-level	Slider: [0, 100]	Enforcement level
probability-to-be-caught	Random 100	The probability to be caught is random