LAFORIA Research Report 95/34, December 1995. Older version under the title "Beyond cooperation-ism and competition-ism (Exploring social phenomena with a generalized tit-for-tat model)", p. 464, Proceedings First International Conference on Multi-Agent Systems (ICMAS-95), San Francisco, CA, June 1995.

Beyond Cooperation and Competition : Explorations with a Quantitative Tit-For-Tat Model [*]

Elpida Tzafestas
LAFORIA-IBP, Université Pierre et Marie Curie,
4, Place Jussieu, 75252 Paris Cedex 05, FRANCE.
brensham@laforia.ibp.fr

Abstract. The modeling of cooperative processes has up to now relied almost exclusively on the traditional cognitivist paradigm, employing explicit representations of goals, beliefs and actions, as seen from an observer's "objective" viewpoint. We present here a quantitative tit-for-tat agent model that is parameterisable, adaptive and scalable. Unlike traditional game theoretic contexts, the kind of social behavioral phenomena we intend to explore shows *relativity* or *subjectivity* (the same social situation may be perceived differently by different agents) and *dynamicity* (if agents are adaptive, an external observer will think the rules of the game change dynamically). This is achieved through the definition of "social" objects or properties, which can be abstract or have a material form, and that are perceivable, accessible and manipulable by the agents. Agents have internal, idiosyncratic motivations that they try to satisfy according to the social feedback. The quantification of the basic tit-for-tat model concerns a threshold parameter (for detection of cooperation) and the perception/action parameters and structures. Diversity in parameters and structures shows as relativity of the game and adaptivity of structures shows as dynamicity. To deploy the potential of the approach, we present three examples where the differential expression of those matching motivations of individual agents gives rise to a variety of social phenomena with intricate dynamics. The implications for the study of emergence in artificial life are finally briefly sketched.

Résumé. Jusqu'à maintenant, la modélisation des processus coopératifs a reposé presque exclusivement sur le paradigme cognitiviste traditionnel, c'est-à-dire le paradigme qui utilise des représentations explicites de buts, de croyances et d'actions, du point de vue d'un observateur "objectif". Nous présentons dans cet article un modèle tit-for-tat quantitatif d'agent, paramétrable, adaptatif et dimensionnable. Contrairement aux contextes traditionnels de la théorie des jeux, les phénomènes sociaux que nous avons l'intention d'explorer avec ce modèle montrent des propriétés de *relativité* ou de *subjectivité* (la même situation sociale peut être perçue différemment par les différents agents) et de *dynamicité* (si les agents sont adaptatifs, un observateur externe verra un jeu dont les règles varient dynamiquement). Pour ce faire, nous définissons des "objets sociaux" matériels ou abstraits, qui soient perceptibles, accessibles et manipulables par les agents. Les agents ont des motivations internes idiosyncratiques qu'ils cherchent à satisfaire en accord avec le feedback social. La quantification du modèle tit-for-tat de base s'applique sur un paramètre de seuil, servant pour la détection de coopération, et sur les paramètres et structures de perception/action. La diversité des paramètres et des structures se manifeste comme la relativité du jeu et l'adaptativité des structures se manifeste comme la dynamicité. Le potentiel de cette approche est démontré sur trois exemples, dans lesquels l'expression differenciée des motivations des agents donne naissance à une multitude des phénomènes sociaux ayant des dynamiques intéressantes. Finalement, nous esquissons brièvement l'intérêt de ce modèle pour l'étude de l'émergence dans le domaine de la vie artificielle.

[*] This work has been the object of a paper (Tzafestas 1995*a*) that was revised to yield this report. Additional applications of the quantitative tit-for-tat model appear in the fifth and seventh chapter of my Ph.D. thesis (Tzafestas 1995*b*).

"In a sense, cooperation could be older than life itself."
(Nowak et al. 1995, p. 55)

1 Introduction

We have been seeking a bottom-up social behavior model in an attempt to integrate social with non-social behavior when designing autonomous agents and especially mobile behavior-based robots or animats. This model of social behavior should fulfill the following requirements : i) *reactivity*, in order to be integrated with other reactive, but non-social behaviors (however, we are going to rediscuss and relax this constraint in section 6), ii) *adaptivity*, in the sense of possibility to modify the internal parameters or structures of the behavior, and iii) *scalability*, which is abstractly the possibility to ascend representational levels so as to explore the potential for higher-level "cognitive" functions.

A parallel motivation for these explorations has been the curiosity to investigate the implications of a "closed-world hypothesis" : the kind of social behavior we wish to investigate is by no means limited to game-theoretic settings, where agents are by definition competitive, but involves more general *participant relations* where agents measure degrees of individual satisfaction that depend on the social context[1]. In such participant contexts, there is no external "policeman" or "manipulator" agent that has set the rules of the game once and for all according to his own interests, but rather the agent themselves "create" this game dynamically. The key idea is that the basic behavior of all agents participating in a social group is the same, but significant variations in internal parameters, degrees of satisfaction etc. exist -and these variations may give rise to complex phenomena, as will be shown in sections 3 to 5. The kind of social behavioral phenomena we intend to explore shows therefore *relativity* or *subjectivity* (the same social situation may be perceived differently by different agents) and *dynamicity* (if agents are adaptive, an external observer will think the rules of the game change dynamically). The role of such a diversity for the evolution of cooperation is by now widely accepted. Diversity, which was initially attributed to stochasticities or errors (Boyd 1989, Nowak & Sigmund 1989), is now being investigated as an intrinsic element and the motor of the evolutionary process (May 1987, Nowak & Sigmund 1992, Glance & Huberman 1994).

The longer-term ambition is the study of the relation between social autonomy and emergence of higher-order organizations. The main hypothesis of this study is that higher-order organizations emerge as a result of cooperation between lower-level entities, and this process is irreversible (see also (Tzafestas 1995*c*)). The lower-level social system has to be autonomous for higher-order organizations to be truly emergent, that is for emergent organizations to be relatively unpredictable. Such a social autonomy will show as robustness to perturbations and can only be achieved whenever there is enough diversity.

[1] In this sense, participation corresponds to what Axelrod and Dion (1988) called a behavior-dependent context of play. It is also consistent with Becker's (1976) operational analysis which suggested that altruism (cooperation) has selective advantage in contexts involving generalized physical or social interaction. De Sousa (1990) argued that such sociobiological models are valid candidates for the study of social phenomena.

Intuitively, a multi-agent system is supposed to have an operational advantage in comparison with a mono-agent system, that is, sociality means that agents "cooperate" toward the achievement of a common goal. The modeling of cooperative and social processes by the distributed artificial intelligence community has relied on this top-down view and has followed the traditional cognitivist paradigm, that employs explicit representations of goals, beliefs and actions, from un "objective" observer's viewpoint : a good review of applications to problem solving may be found in (Durfee et al. 1989). This approach has focused on the subproblems of goal identification and communication, conflict resolution, negotiation etc. (see several papers in (Bond & Gasser 1988), (Huhns 1987), (Gasser & Huhns 1989), (Demazeau & Müller 1990) and (Demazeau & Müller 1991)). The same top-down cognitivist paradigm also served as a foundation of research on communication and social structure (Werner 1989) (Werner 1990) (Findler & Malyankar 1993).

However, as Castelfranchi has pointed out (Castelfranchi 1990) (Conte et al. 1991), the fundamental question should not be "how to get multiple agents achieve a social goal", but instead "why does an autonomous agent enter into social interactions" (*the sociality or dependency problem*) and "how does an agent get his problem to become social, i.e. get it adopted by other agents" (*the goal adoption problem*). Castelfranchi further proceeded to declare that cooperation is just a special case of goal adoption, which in most approaches takes place as if by magic, and that for adoption to occur, communication or request are not compulsory. He advocated therefore a selfish view of generic social exchange, based on reciprocal self-interest of autonomous agents, rather than on instruction and persuasion. Castelfranchi's view is essentially ***bottom-up*** and ***selectivist*** : what appears as agent sociality is nothing else but the sharing of otherwise individualistic, selfish goals of the agents.[2]

We call ***sociality*** *the goal sharing (or need sharing) of agents, that is, sociality intervenes to the agents' "motivational" system and modifies those motivations, or else it induces* ***social perturbation****.*

Of course, goal sharing does not imply that all agents have the *same* goal ; much more complex social relations and structures emerge if the agents' goals form a multi-level dependency network.

In this sense, cooperation is just positive sociality, i.e. the kind of sociality that induces an improvement of individual agents' performance, always passing through their motivational system. The difference with the top-down view is that now cooperation becomes relative : ***we call cooperation the type of sociality that leads to an individual performance improvement of an agent according to its*** **own** ***criteria***. That is, agents decide themselves about what is cooperative and what is not, and a situation that is perceived as cooperative by an agent may not be considered so by another one. A population of agents having this relative and differentiated form of sociality, hence a society of agents, will give birth to emergent structures, among which are social standards. Relative sociality will result in an agent deciding whether he will follow the standards or not, whether he will "cooperate" or not, so that those emergent social structures will be dynamic. We may say that this society is

[2] This didn't prevent him from remaining within the cognitivist paradigm during his analysis of social commitments (Castelfranchi 1995), even if his reasoning does not necessarily rely on a mental-states-based description.

autonomous, in the sense that it evolves mainly due to internal forces, rather than in response to some sort of external "instruction".

Intrinsic reciprocity in social relations, that is *motivated reciprocity* stimulated by *"kin" recognition* (i.e., recognition of agents that share goals), parallels the two hypotheses formed by evolutionary biologists in order to explain the cooperative living systems (see Trivers (1971) for reciprocity and (Hamilton 1964) for kinship) and on which a whole branch of research is founded, the study of the evolution of cooperation (Axelrod & Hamilton 1981). Note that this type of sociality resembles the vertebrate "explicit" sociality, rather than the insect eusociality, where societies are formed by differentiation (Wilson 1975). As already said, this choice has been made with the longer-term aim to investigate whether higher-level cognitive abilities may emerge out of reactive, pro-social ones (point made by McFarland (1994)). With these in mind, we quantified the original tit-for-tat model (Axelrod & Hamilton 1981) in a way to be able to apply it in the context of participant social relations : instead of a static game, such as the prisoner's dilemma, we wish to investigate situations where agents create themselves this "game" dynamically and agents' history plays a more direct role than in a game. On top of that, our explorations focus on agents' operationality, that is on their satisfaction during their "life" rather than on their reproductive value (fitness).

2 The quantitative tit-for-tat model (or, why my cooperation is not the same as yours)

Social behaviors are just like other agent behaviors in that they depend on some internal motivational drive as well as on external stimuli. We may assume that the motivation of the agent to engage in the social behavior/task is genetic and is a decreasing function of the degree of satisfaction of this task ; however, in the following examples, we will only consider agents with a single, social task, so we will not need to compare this motivation to others, nor perform any arbitration between motivations. The external stimuli, upon which such a behavior depends, express measures of the corresponding social activity which may be interpreted as friendly (cooperative) or hostile (non-cooperative). We expect an agent to be friendly in friendly environments and hostile in hostile ones. Furthermore, by coupling the stimuli perception to the internal degree of satisfaction, we may say that whenever that degree of satisfaction -itself being a function of the value of the perceived stimulus- exceeds a given threshold, the social environment is considered cooperative, otherwise it is considered hostile. The agent will then modify its degree of participation to its social environment according to this perception. The perceived stimulus is some sort of "social" object (hint, in terms of Hogg & Huberman (1992)), abstract or material, directly visible and manipulable by the agents. Finally, the actions of the agents influence directly the *same* social objects that are perceived, so that social interaction is a continuous, self-catalyzing process, relying on the presence of such *matching agent motivations*. The above are summarized in figure 1.

This model is actually ***a quantitative tit-for-tat model***, where the agent actions as well as the perception and satisfaction functions are continuous, rather than binary. This allows us to include *adaptive components* into the agents' behavior, that will modify these functions in order to achieve higher satisfactions -an illustrative example is presented in the next section. The *a()* and *b()* functions are in essence *metabolic functions* and we can complexify them at will, as is the case in sections 4 and 5. The perception function might include substantial *filtering*, equivalent to distortion of objective reality (if we assume that such a thing as objective reality exists!). Filtering may also be seen as an interpretation of reality in terms of what the agent understands. It will be later argued that those *a()*, *b()*, satisfaction and filtering (from now on denoted *f()*) functions are actually implicit predictive functions of the behavior of other agents -however, there is no explicit logically or symbolically represented belief, expectation, intention etc. The output functions *a()* and *b()* may also include side-effects, such as mobility, aggression and flight in the example of section 5. Unlike Hamilton's model, the perception function of the agent, which gives a measure of cooperation, reasons on immediate satisfaction, rather than on ultimate fitness.[3] The immediate satisfaction is, as already said, a linear function of the perceived stimulus value, which in the subsequent sections is computed by averaging the outputs of agents. This measure of participation/cooperation can be then considered as an *inclusive reward*, by analogy with Hamilton's inclusive fitness.

Social behavior

If $perception(social\_stimulus) \geq T$,
then "cooperative or friendly environment, hence cooperate"
output *a(perception)*
else "non-cooperative or hostile environment, hence defect"
output *b(perception)*
where :
social_stimulus is the social object involved,
T is the cooperation threshold
a() and *b()* are the action functions of the agent in friendly or hostile environment, respectively.

Figure 1 The quantitative tit-for-tat model [4]

Agents are generally heterogeneous : thresholds, as well as functions *f()*, *a()* and *b()*, may vary widely from agent to agent, still the basic model remains intact. Agents are all born with the same needs qualitatively, but not quantitatively (we wouldn't call it a society if the agents were not supposed to share something). Moreover, different agents may be defined to have different individual needs or participate indirectly in the process that triggers them (by outputting a different kind of object than the one perceived) etc. -so that a variety of complex situations may be modeled. In sum, what one agent perceives as cooperative may not be so for a second one : cooperation criteria are relative and agent actions have a differential impact upon other agents.

The above quantified tit-for-tat model was adopted for a number of reasons. First, the basic tit-for-tat model has been proven to be an optimal and evolutionary stable strategy, since it is nice, retaliatory and forgiving (Dawkins 1976) (Axelrod 1984). Hence, as a starting point, the tit-for-tat model seems a plausible model of generic social interaction. However, and due to its binary nature, it looks insufficient for modeling complex social phenomena, such as the ones presented in subsequent sections. Defining continuous satisfaction, *a()*, *b()* and *f()* functions, multiplies its

[3] Cognition might just be the process that learns to associate the two measures.

[4] In a game-theoretic context, such as the prisoner's dilemma, a tit-for-tat agent is an agent who has the simple strategy of cooperating in the beginning of the game and returning its opponent's move thereafter (Axelrod & Hamilton 1981).

modeling potential, as has been already exposed, and allows great variations that may give rise to intricate dynamics. However, this is not a game-theoretic setting like that exemplified by the prisoner's dilemma (as in Axelrod & Hamilton (1981)), since there is no a priori temptation for agents to defect. Actually, the agents don't have any particular "bias" or "will" toward either cooperation or competition. Cooperation, deceit and a plethora of other social phenomena are all in the eye of the observer. What the agents are doing is trying to maximize their individual satisfaction and so act according to perceived social feedback -which isn't any different from other types of feedback they receive. Agents in pursuit of multiple satisfactions, or else governed by multiple hedonistic motivations, will unavoidably seek compromises between those motivations. Returning to the prisoner's dilemma and games theory -that have benefited from a recent resurrection within the sciences of complexity and artificial life (Nowak & May 1992), (Sigmund 1992), (Angeline 1994), (Batali & Kitcher 1994)-, the above presented quantitative tit-for-tat model may certainly be instantiated in such contexts by simply adding *a special play-maker agent* who will set the rules of the interaction and will "manipulate" the rest of the agents by modifying the rewards they receive. Spontaneous emergence of social phenomena will then be pruned and canalized to particular sub-spaces according to the play-maker's own motivations and intentions.[5]

All the simulations whose results are reported in the following paragraphs have been performed using synchronous simulation mode (i.e. all agents perceive the world and act at the same discrete time steps) ; this was meant to reduce complexity of the emergent phenomena and thus enhance observability, while subsequent phases of the study will certainly involve continuous time-scales (or, equivalently, delayed action models) which have been shown to have drastic impact on the dynamic behavior and especially the stability of distributed systems (Huberman & Glance (1993), Bersini & Detours (1994)). The results presented next in increasing order of example complexity do not purport to be significant or novel by themselves, but to demonstrate through seemingly simple examples the potential of the model to account for complex "emergent" phenomena. This is also the reason for analyzing them only qualitatively. Finally, and unlike work on the original tit-for-tat model (Axelrod & Hamilton 1981), we don't study evolutionary phenomena, but emergent social ones during the agents' life-time.

3 Cooperation, environment and theology : Diversity in parameters

The first example instantiated after the quantitative tit-for-tat model involves a community of agents each one seeking to maximize the received "affection" by other agents. As already said, the affection measure each agent perceives is the average of all the affection rates output by the agents. We assume that all values are normalized between 0 and 1 and that *f()* is the identity function, *f(x)=x*, so that there is no filtering of reality. We also assume that there are variations in the agents threshold, *a()* and *b()* functions -because, if all agents were born the same, they would be happy from the beginning and there would be no problem to solve : all tit-for-tat agents would start by cooperating and the system would be stable. Finally, we are only considering the case

[5] Sometimes, nature itself may be regarded as the play-maker, who sets resources and other constraints ; real biological examples generally fall into this category (Wilson 1975).

where $a()$ and $b()$ are linear functions : $a(x)=a*x$ and $b(x)=b*x$, though we could easily generalize. For agents to be considered as rational, $a(T) \geq b(T)$ should hold, whence $a \geq b$. It follows that the three parameters of our affection model are T, a and b.

Experiment 1 : Variations and perturbations. We simulated several variation schemes for interaction periods from 40 to 100 cycles. Fixing T to 0.5, defining a as random between 0 and 1 and $b=a/2$ or $b=a/3$; fixing T to 0.5, defining b as random between 0 and 1 and $a=max(2*b,1)$; having a and T be randoms, $b=a/2$ or $b=a/3$, etc. In all cases, the agents performed only a few (from 0 to 5) cooperative moves in successive cycles in the beginning of the interaction and the system then stabilized to all agents acting non-cooperatively. We then introduced perturbation to the system, where with a given probability p ($0<p<1$) the global affection value was reduced by a factor of f ($0<f<1$). Again, simulating for a number of pairs of values p-f in a number of schemes of variations yielded the same qualitative results, as might also be intuitively expected, inducing however additional variation in the satisfaction state of the agents. From this it might be deduced that *the present rigid interaction scheme by itself causes already very bad behavior for perturbation to make any difference*. Note that the variations may be seen as the genetic cause of misimplementation and misperception, to use Axelrod & Dion's (1988) terminology, while the perturbations correspond to changes in the "rules of the game".

Experiment 2 : Adaptive agents. Given the low satisfaction rates of agents in the previous experiment, we added adaptive mechanisms that modify those internal T, a and b parameters.[6] We differentiate between two such mechanisms, the *passive* mechanism which makes the agent more indifferent to the absence of affection by reducing T, and the *active* mechanism which makes the agent give out more by increasing a and b up to the maximum value of 1. If we wish to draw a real-life analogy that might help in better evaluating our results, we might associate the active and passive mechanism to christian and buddhist behavior respectively[7]. In both cases, we defined a learning horizon w (during which an agent measures its satisfaction) and a learning *rate* which serves to update T or a respectively according to the adaptation law : *T(or a)=T(or a)-rate*diff*, where *diff=T-average_input_level_during_w*. Update (adaptation) only takes place when *diff* is positive. Simulating the community for various parameter sets (T, a, w, *rate*) in all three modes, i.e. without adaptation, in active mode and in passive mode, no significant changes were observed. This was apparently counter-intuitive ; closer examination and inspection of the simulations revealed that what was wrong in the initial model was the method of experience accumulation : actually the agents should be satisfied (or happy inside) *by introspection*, i.e. when they *act* cooperatively and not when they *perceive* a cooperative environment (*diff=T-average_output_level_during_w*). For this to work, however, we needed an additional parameter : *generosity*, i.e. a probability g ($0 \leq g \leq 1$) with which an agent does *not* return a bad move. Simulation with different degrees of *generosity*, learning rates etc. showed that the new adaptive agents' average

---

6 Another simple example of adaptivity in tit-for-tat agents is the reinforcement model "Pavlov" of Nowak and Sigmund (1993).

7 We do not mean however to give here any *explanation* about theological positions and beliefs, we are just using this analogy as a basis for better interpretation and more profound understanding of the results obtained.

satisfaction were an order of magnitude larger than that of the non-adaptive, non-generous agents (around 30% as opposed to 2-3% in the non-adaptive case). As in the previous experiment, the uniform case (with fixed *g*=0.2 and *r*=0.3) was half as performant as the case of extreme variation (*g* and *r* random) where average satisfaction could rise up to 50%.

Due to the statistical nature of experiments and the large variations in random settings of individual runs, we compared cases by randomly breeding a population and then initializing it in different contexts : all active, all passive, mixes, etc. Qualitatively, the results have been the same over a large set of populations ; to see the comparative quantitative differences, we present in the following tables some typical cases (however, and due to the random variations introduced, the quantitative results may seem inconsistent among different tables). Table 1 gives the comparative results for a population of agents when defined as non-adaptive, as active or as passive. As can be seen from those results, the active agent does not differ from the non-adaptive one (only a little bit in the beginning, but not in the long term), while the passive mode shows slow, but steady improvement. This may be attributed to the fact that the "active" agent actually does not adapt its cooperation detection mechanism, which is unlike the "passive" case (where the threshold changes).

	t=200	*t=400*	*t=600*
10 non-adaptive	34.71	35.78	35.42
10 active	35.72	35.16	35.44
10 passive	36.51	39.87	46.59

Table 1 Comparative results for 10 agents (columns give the average satisfaction percentages at different moments during each run).

Experiment 3 : Adaptation and environment. We then proceeded to examine whether the two adaptation mechanisms would perform differently in environments with different degrees of perturbation (where perturbation corresponds to reducing the average output of the agents by a given factor with a given probability). We simulated homogeneous and heterogeneous populations, consisting respectively of only passive or active agents or of a mix of the two. In all cases the average satisfaction of the agents has been found insensitive to the degree of perturbation -which at first glance is not expected, since perturbation in this context is nothing else than an invitation to learn. This is an indication that *the variations between agents are much more critical than the environmental perturbation* (in this sense, learning is true social learning).

	No perturbation			*Low perturbation (0.1, 0.1)*			*High perturbation (0.8, 0.8)*		
	t=200	*t=400*	*t=600*	*t=200*	*t=400*	*t=600*	*t=200*	*t=400*	*t=600*
10 active	57.61	56.18	56.43	57.41	56.51	55.89	56.07	55.61	55.96
10 passive	56.76	61.10	64.64	55.62	59.98	63.86	54.78	60.55	63.91
5 ac., 5 pas.	56.91	61.87	65.31	55.22	60.57	63.91	56.52	56.11	56.57

Table 2 Comparative results for 10 agents (columns give the average satisfaction percentages at different moments during each run)

Experiment 4 : Meta-adaptive or multi-modal agents. We then went on to investigate agents that possess both adaptation modes and have to learn to apply the one that proves to be most effective, i.e. the one that leads to higher satisfaction. To this end, we introduced continuously updated success measures in the two adaptation modes with probabilities relative to those success measures. All agents at birth have *p(active)=p(passive)=0.5*. Results from comparative runs for 40 such multi-modal agents are given in table 3. It may be observed that although the passive adaptation mode has been proven more effective than its active counterpart, the multi-modal agents are not capable of recognizing this (the relative probabilities of the two modes are both and always around 50%). Why is this so ? By closely inspecting the simulations it was found that actually the active adaptation mode has been "exploiting" the passive one's superiority, by being successful not thanks to its own competence, but because the passive mode has been improving the default behavior of the agent (by reducing the threshold).

	No perturbation			*Low perturbation (0.1, 0.1)*			*High perturbation (0.8, 0.8)*		
	t=200	*t=400*	*t=600*	*t=200*	*t=400*	*t=600*	*t=200*	*t=400*	*t=600*
satisfaction	49.73	50.50	51.92	49.54	50.27	51.88	49.96	51.14	53.14
p(passive)	50.50	49.62	50.23	49.11	50.57	51.82	50.33	50.51	49.91

Table 3 Comparative results for 40 multi-modal agents at different moments and contexts.

Experiment 5 : Injecting vicious agents. Next, we introduced apparently vicious agents (but only apparently so!), that is, agents that defect when you don't expect it, because their perception system is noisy ; as with global perturbation, we defined a perturbation probability (constant across different agents and equal to 0.5) and a perturbation factor (again constant and equal to 0.5). Furthermore, we banned them from any adaptation possibility. We run several mono-type or multi-type populations of agents, always mixing with a proportion of vicious ones. In all cases, *the vicious agents showed inferior satisfactions than their adaptive counterparts, while the adaptive agents showed substantially superior satisfaction when vicious agents were present* -because then they learned to become better. Actually, presence of vicious agents is nothing else but another kind of more elaborate perturbation for "normal" adaptive agents. Finally, the average satisfaction of a population without vicious agents was found to be between that of a population with a high percentage of vicious agents (around 50%) and one with a low percentage (around 10%) -this is a consequence of the contribution of each sub-population to the overall satisfaction.

Experiment 6 : Doing the unthinkable. In this last experiment, we modified the active adaptation mechanism so as to scale generosity (*"learn to love your enemy"*). The same adaptation formula previously applied to a was now applied to a *and* g. Not unsurprisingly, various simulation runs showed that the active agents quickly developed generosity close to the unity and arrived at satisfaction rates as high as 95%. Of course, the same is true for a passive mechanism that scales generosity (although generosity is an active parameter, not very compatible with the passive mechanism). The remarkable difference between the two, as is shown in table 4, is that now the active mode is much more "efficient" than the buddhist one, both in terms of absolute satisfaction and learning speed.

	No perturbation		*Low perturbation (0.1, 0.1)*		*High perturbation (0.8, 0.8)*	
	t=200	t=400	t=200	t=400	t=200	t=400
40 active	90.52	93.11	90.54	93.15	90.29	93.01
40 passive	81.76	84.58	80.98	84.43	80.36	84.14

Table 4 Comparative results for 40 adaptive agents that scale generosity as well

Evaluation. The motivation behind these experiments has been to show that the social learning (or cooperation) problem owes its existence to the presence of large (genetic) variations in the agents' program. I have also tried to understand if there is a relation between the degree of environmental perturbation and the adaptation mode. To tell the truth, I was expecting the active approach (the christian one) to prevail in more predictable, less perturbed and hence more controllable, environments, and the passive approach (the buddhist one) to prevail in more unpredictable environments, where agents' actions might all too often be in vain. Clearly, this is not the case : the passive approach is *always* more attractive, regardless of the specificities of the environment (Borges (1976) traces "buddhist" trends throughout the history of western culture, from Heraclitus to Schopenhauer and to Bergson). Moreover, for agents that try to learn the optimal adaptation mode, it has been impossible to stabilize to such a mode, because we cannot find a safe success criterion for a mode. This could suggest that if christianism and buddhism have prevailed respectively in different environments, this is surely a result of meme spreading (Heylighen (1992*a*,*b*), Gabora (1993)) ; how these memes made it there and so successfully, remains an enigma. Another observation concerns the temporal behavior of the two adaptation modes : the active one stabilizes quickly to a large value, while the passive one goes on improving for a long time. Perhaps this means that with a given life-time the active approach is ensured to lead to some more-or-less satisfactory state, while the passive approach has to resort to the evolutionist idea of perpetual reincarnation and chance renewal [8] (assuming that all the rest is equal between the two approaches). Finally, the extreme active mechanism which includes generosity scaling, seems so unnatural (since it invites exploitation) and so difficult to develop by itself (albeit, admittedly, very effective!) that christianism's claim looks logical, that it took a major event to discover it, the incarnation of the god's son himself.

4 Security and stability in an artificial economy : Diversity in structures

The second example that made use of the quantitative tit-for-tat model was an artificial economy. An artificial economy is a population of agents that receive salaries and make expenses, i.e. consume. The dynamics of the system become interesting once we define heterogeneous agents that react differently to the social situation (the "social object" used for interaction is the consumption of different agents). The satisfaction criterion for agents is to be able to spend as much as they intend to. An agent is therefore satisfied if $intended\_consumption \leq safe * threshold$, where the *threshold* in this case expresses a security factor and *safe* stands for the agent's cumulative surplus over time. Usual "rational" agents are supposed to have high satisfaction values according to the above criterion, since their consumption intentions are not irrational. It should be stressed that we are not simulating the evolution of prices or any other financial-type phenomena, like in Beltratti & Margarita (1992), Nottola et al. (1992), de la Maza & Yuret (1994), Benos & Tzafestas (1995). We purport, instead, to show how individual motivations and predispositions toward social structures give rise to intricate social phenomena.

At each cycle an agent receives its salary (*s*) and computes its intended consumption level as *c*random_factor*, where *c* is the associated cost and *random_factor* is a random between 0.9 and 1.1 (the actual cost won't be exactly the same at each cycle). On top of this regular consumption, the agents make investments periodically (every T_i) that are computed as $(s\text{-}c)*T_i*threshold*a\_random\_factor$, where *a_random_*

[8] It is a pity that the presence of the word "Zen" in the title of Ray (1994) was just a joke.

factor is as before (this corresponds to a "preplanned" investment which is consistent with the agent's rationality). We have defined three types of consumer agents :

- ***Rational consumers.*** Agents that act as above. This appears as a conservative regulation of their expenses and that allows them to make investments occasionally. They show some kind of social adaptation by "mimicry" since they drift their expenses toward the social average (i.e. the input) to a degree of 10%. The *a()*, *b()* and *f()* functions are *a(x)=x*, *b(x)=safe*threshold* (= the maximum consumption allowed) and *f(x)=intended_consumption+0.1*(x-intended_consumption)*.
- ***Aggressive consumers.*** They don't do any investments. Instead, they want to show off at times, by spending more than everybody else (actually, more than the average during the previous cycle). Aggressive consumers have a threshold of 1 (they might spend all they have at once). Their aggressivity is again expressed periodically (every T_i) and their *f()* function is then *f(x)=max(x*(1+aggressivity), intended_consumption)*, where *aggressivity* is an additional parameter that takes a value between 0.5 and 1.
- ***Scrooge consumers.*** They appear as very insecure. They differ from the previous two types of consumers in their perception filter. While all the others are satisfied if their *intended expense* does not exceed a security level, those scrooge agents are satisfied if the *external perceived average expense* does not exceed this level (*f(x)=x*). Every time this condition is not met, they (absurdly to an observer) decrease that security level by using the formula *threshold=threshold-(x*insecurity/safe)*, where *insecurity* is an additional behavioral parameter ranging from 0.5 to 1. They make no investments either and they demonstrate no social drift.

Experiment 1 : Rationality versus variability versus aggressivity. We simulated populations of 15-20 rational agents with cost and salary variabilities (from 1 to 10, with *cost<salary*) and random thresholds between 0 and 1. The average degree of satisfaction in the case of *2<salary<3*, *salary/10<cost<salary* was found to be around 95% and to have a slight decreasing tendency with the degree of variability. Experiments for long periods (around 500 cycles) and a lot of agents (around 100) with high variability showed that the average satisfaction would never fall below 90-92% in any case. This is an indication that the agents' satisfaction is not sensitive to social variability, as long as the agents are rational. We then simulated mixes of populations : 20 rational agents with 1 to 10 aggressive ones. As was expected, the average agent satisfaction dropped immediately and was found to be around 85-90%. By inspecting individual agents we observed some marginal unsatisfactory agent cases : moderately aggressive agents that have a *cost/salary* ratio close to 1, rational agents with high *insecurity* and highly aggressive agents independently of *cost-salary* relation. Moreover, while a rational agent's *safe* value tends to increase slowly, an aggressive agent's satisfaction state becomes worse and worse as time goes by. Of course, not all aggressive agents with low *cost/salary* ratio end up unhappy : slightly aggressive agents with low *cost/salary* ratio tend to behave only slightly worse than rational ones. Next, we simulated mixes of populations of rational, aggressive and scrooge agents. The qualitative behavior of rational and aggressive agents did not change, while scrooge agents converged to one of the following two behavioral profiles : scrooge agents with moderate or low *insecurity* and low *cost/salary* ratio

tended to be as happy as rational ones, while those with high *insecurity* or high *cost/salary* ratio tended to stabilize to a constantly unhappy state. As a general observation it can be said that, as was probably expected, the scrooge agents are marginal -they don't really participate to the social context- while the rational ones are not too much affected by the presence of aggressive agents, unless they are really too poor. Those latter agents may become very unhappy, if the social environment allows it.

Experiment 2 : An uneven economy. Next, we proceeded to simulate economies with two classes of agents, the rich ones (*salary* from 30 to 40) and the poor ones (*salary* from 2 to 10), with *cost* from *salary/10* to *salary*. We simulated a population of 80 rational agents with the probability of an agent being poor 50% and 90%. The average satisfaction has been found inferior to that of the classless society and the less satisfied agents have been found to be the extremely poor (low *salary* and high *cost/salary* ratio) and rather insecure agents. On rare runs, some rich agents were found unhappy due to high *insecurity*. We repeated the simulations with a mix of 50 rational and 50 aggressive agents and the average satisfaction has been found even lower : the reason is the presence of unhappy aggressive agents ; aggressive agents tend to be more unhappy if they are poor and the higher their *aggressivity*, the less significant the role of the *cost/salary* ratio.

Experiment 3 : Exogenous perturbations. To investigate the stability of the economy and its resistance to exogenous perturbations, we ran twice a mixed 85-agent population -with all three types of agents and salaries between 2 and 10 (for every agent, *cost* was between *salary/10* and *salary*)- for a period of 160 cycles. During the first run and at t=80, we increased each agent's *cost* to a value between *(salary+cost)/2* and *salary*, leaving *salary* intact. During the second run and again at t=80, we increased all agents' salaries by a factor from 1.5 to 3. In the first run we observed that rational agents' satisfaction was not much affected by the perturbation, especially for those of high or moderate *security*, while aggressive ones deteriorated proportionally to their *aggressivity*. Scrooge agents satisfaction state remained fairly stable as well -because after rising costs the other agents were spending less on average. In the second run we observed the inverse phenomena, of rational agents that stabilized after the salary increase, highly aggressive agents that suffered less during the second phase -because they could then satisfy their passion (genetic defect)- and more or less degrading scrooge agents.

Experiment 4 : What does it take to kill an agent ? Once certain about rational agents' resistance to perturbations, we tried to examine under what conditions aggressive or scrooge agents might become very unhappy. Again, we simulated the above 85-agent mixed population for a period of 160 cycles. At t=80, we introduced 20 rational but "alien" agents (with salaries between 20 and 30 and for each agent *cost* between *salary/10* and *salary/2*, i.e. sufficiently small). We assumed that those agents had been brought up and had stabilized in a different, much more wealthy, environment and that they were injected into this economy somehow by force. As was expected, poor but highly aggressive agents (*cost/salary* ratio close to 1) were driven to despair : aggressive agents, previously stabilized to a satisfaction of 50-70%, might fall down to 20-30% after the aliens' intrusion. We also simulated an economy of 60 heterogeneous agents for a period of 160 cycles. At t=80 we introduced as above 20 alien agents, but this time the agents were aggressive. By inspecting the agents throughout the run, we observed that the highly insecure scrooge agents' satisfaction degraded significantly during the second phase : agents whose satisfaction had fallen from 35% to 25% during the first phase were found to degrade up to 10% during the second phase (actually, the degradation is greater than those figures show, since they correspond to average satisfaction from the beginning of the experiment).

Evaluation. Our motivation for running those experiments was to investigate how stable such an artificial economy is to different kinds of perturbation and how robust each of the individual agent types is. It can be deduced from the experiments that the rational agents are the most robust to all kinds of perturbations that have been studied, hence they are less liable to manipulation ; they show some drift toward the social standards, but this enforces, rather than breaking rationality. Both aggressive and scrooge agents are unstable and not robust because, unlike rational ones, they don't reason in terms of their internal needs (relation between *cost* and *salary*), but their behavior is more context-dependent, i.e. they can be easily manipulated. Furthermore, if given the chance, the aggressive agents have the potential to cause instability in the population. The difference between aggressive and scrooge agents is that the former ones *are irrational in acting*, while the latter ones are *irrational in sensing*. This is why the aggressive agents may become really dangerous for the other agents, whereas the scrooge agents are only dangerous to themselves.

5 Productivity, insulationism and social groups : Diversity and mobility

In this last series of experiments, we have been trying to see to what extent the model supports side-effect phenomena of cooperation and non-cooperation, like the formation of spatial structures, distributions etc. We assume a population of agents situated in a small-size grid-shaped world. Each of the agents may execute a number of operations (it could be, for instance, different types of assembly operations in an industrial assembly cell) and has a measure of its productivity (it could be execution time). The goal is for agents to maximize their participation in the production process by working at maximum speed ; we assume that agents may detect average social productivity in their position and compare it to their own. The satisfaction criterion for agents is to participate in a social (territorial) group whose average productivity is not too small compared to their own. If this criterion is not satisfied, the agents will tend to flee by moving elsewhere (the assembly robots would shift to a neighbouring cell), so that mobility is a side-effect of defection. According to the above, the *f()*, *a()* and *b()* functions are : *f(x)=x*, *a(x)=speed* and *b(x)=0* (with flight as a side-effect). The threshold is defined as *threshold=speed*tolerance*, where *tolerance* takes a value between 0 and 1. It is important to note that, in this context which involves territoriality (clustering), there is no need for a special kin recognition mechanism, like it is put forward by Hamilton (1964), even if the environment is a competitive medium. The results of actions themselves are an indication of the proportion of kin in the group ; this appears as an indirect kin recognition mechanism. Of course, this makes sense only in long interaction sequences, as the analysis of Axelrod & Hamilton (1981) reveals.

Experiment 1 : Variations and the need for aggressivity. We simulated a population of 15 agents distributed over a 3x3 grid for a period of 40 cycles. We performed four runs : with only four agent speeds (0.25, 0.5, 0.75, 1) or with random speeds between

0 and 1, and with *tolerance=0* or *tolerance* between 0 and 1. The average population satisfaction was found larger in the case of four speeds than in the case of random speeds and higher in the case of random tolerance than in the case of zero *tolerance*. By closely inspecting individual agents, we found that the most happy agents were the least productive ones independently of *tolerance*, because, although they had to flee quite often, they were happy in almost any group. Inversely, the most unhappy agents were the highly productive and a little tolerant ones because they couldn't *actively* seek satisfaction by chasing other significantly slower agents ; instead, they had to rely on the slower agents' possibility to detect social un-satisfaction so as to flee. Clearly, the behavior of the population looks sub-optimal, since slow, but highly tolerant, agents never flee. What appears necessary is a mechanism for direct aggressivity and an aggressivity that would be perceivable by all agents as an additional cue for social un-fitness.

Experiment 2 : What aggressivity really is. The new modified model for the aggressive agent is shown in figure 2. Now the agent's satisfaction is defined as detection of cooperation *and* absence of aggression. We repeated the four previous runs, this time with aggressive agents and the average satisfaction was found to be better in all cases by a factor of 1.5 to 2. Agents' behavior appeared more natural and more balanced : an agent tended to chase slower agents and flee away from faster ones. Obviously, faster agents usually chased others more often than they fled. To better observe the comparative advantage of aggressivity, we compared the dynamics of the same 15-agent population with random *tolerance* and *speed* when defined as kind and when defined as aggressive. The improvement of average satisfaction with aggressivity has been found dramatic : the 20-30% satisfaction of kind agents would rise up to 80-90% in the aggressive case. Inspection of individual agents also revealed that aggressivity allows to highly fast and only a little tolerant agents to be maximally satisfied. It seems therefore that in this context *aggressivity is not a particular motivation of the agent to chase others but a deadlock avoidance and conflict resolution mechanism*. Observations similar to this one were drawn by Galliers (1990). It is also noteworthy that in this case the most unhappy agents are found to be -somewhat counter-intuitively- the stronger/faster ones, who may be manipulated in this sense by external forces. This is no longer strange, once we observe that an agent's internal need is in fact proportional to its productivity/strength/speed.

Aggressive behavior

If $perception(social\_ stimulus) \geq T$,
then
    if there is aggression in the environment
        if I am angry (from the previous cycle)
            "previously angry, I assume everything is alright now,
            I return to normal state and I give a new chance"
            output a(perception)
        else "hostile environment,
            since I am not angry, I am probably the cause for
            others' aggression, so I defect and flee"
            output b(perception) and flee as a side-effect
    else "everything is fine"
        output a(perception)
else "I am stronger than the others, so I aggress, but output normally"
    output a(perception) and aggress as a side-effect

Figure 2 The aggressive agent (with 1-step memory)

Experiment 3 : Aggressivity and insulationism. Next, we tried to investigate to what extent the model favors insulationism of agents. To this end, we run the same population of 6 (aggressive) agents in a 3x3 and in a 2x2 grid for a period of 40 cycles. In the beginning of the run, all the agents were placed in the same, randomly chosen, position in the grid. In the end of the run we found that, when in the 3x3 grid, the agents distribute themselves over the grid ending up alone in a position, while in the 2x2 case there were two pairs of agents and two loners. In both runs, the average agent satisfaction was the same (95%) and the same was true of the average number of moves (0.83) and aggressions (0.5). Furthermore, in the 2x2 case, both pairs of agents comprised a rapid, yet tolerant, agent and a slow one. As a result of the above, it may be said that a sufficiently large world allows for agents to isolate themselves, while a smaller world forces them to seek compromises with other agents. Clearly, the important factor in such a costrained world is *tolerance*.[9] Those compromises might be perceived by external observers as cases of what Trivers (1971) called "subtle cheating" (where the deceived agents persist to this sub-optimal "choice", because it is less costly than the complete absence of interaction).

Experiment 4 : It's not crazy, it's different. Another phenomenon that deserved study was the presence of speed variations non-uniformly distributed over the space of possible values (from 0 to 1). To this end, we performed two runs with a majority of 15 similar agents and a single very different agent in a 3x4 world. The first run involved a majority of rapid agents (with speeds from 0.8 to 1) with low *tolerance* (0.2) and a unique slow agent (initialized with a random *speed* from 0.1 to 0.2) with random *tolerance*, while the second one involved a majority of slow, moderately tolerant agents (with speeds from 0.1 to 0.2 and *tolerance=0.5*) and a single rapid agent with low *tolerance*. In the first run we observed that the rapid agents formed clusters quickly (by t=10 they had stabilized), while the slow agent "wandered" over the grid until it finally ended up alone somewhere. Of course, putting a large population of rapid agents that distribute over the whole grid leads to the slow agent wandering indefinitely. In the second run the dynamics were the inverse : the rapid agent quickly chased all those in its position and ended up alone ; the chased agents stabilized immediately elsewhere in the grid.

Experiment 5 : Sociality cannot be forced. To make up for this potential for insulationism, we added a sociality bias (during flight the agent would not choose an adjacent position randomly, but it would choose the highest populated one). We then repeated the two above runs twice by using aggressive agents during the first and socializing ones during the second. It was found that in the case of a highly productive majority a cluster of highly tolerant agents quickly emerged that accomodated the slow one, while in the case of an unproductive majority there was no difference with before : once more, the highly productive agent found himself alone soon. This is an indication that sociality should be more explicit to favor social clusters : the agents should have another, more direct, kind of (again selfish!) benefit by sticking together, such as another social task whose satisfaction will depend on the mere presence or even on the number of agents around (alliances of primates are an example of such a

[9] It might be interesting to investigate whether scaling tolerance is feasible, by analogy to section 3 above.

process), so that the agents would try to compromise between apparently conflicting motivations.[10]

Experiment 6 : Self-organization. To investigate the potential of the model for self-organization, we performed a run where, after the system had stabilized spatially, we re-initialized the agents speeds to different values, in order to inspect the new dynamic behavior. To accentuate the observable phenomena, we run the experiment with 15 slow, moderately tolerant and a single rapid and little tolerant agent for a period of 80 cycles. Speed re-initialization occurred at t=40, after the initial population had stabilized. Given this distribution of agents the system stabilized in comparable times in both phases of the run, as was expected. The important observation to draw is that the population's self-organization potential is such that it will try to find a new stable state if there is one. Statistically, and if given enough time, it will. If the environment is changing more rapidly than the average transition between stable states, the agents will appear as being after a Sisyphus task, where once close to a stable state (nirvana), all the work accomplished will "magically" disappear and they will have to start over anew.

Experiment 7 : Conflicting motivations and social groups. The last experiment consisted in testing the potential for compromise between conflicting motivations. We then defined motivations that were triggered positively by some values of the social stimulus and negatively by others ; the two sets of positive and negative trigger values were disjoint and their union was the whole set of possible values for the stimulus. Now the agents' motivations receive both positive and negative feedback and the *f()* function becomes : *f(x)=max(f(+)-f(-),0)*, where *f(+)* and *f(-)* denote respectively the averages of positive and negative values of the social stimulus involved. A run of 15 agents in a 3x3 grid for 400 cycles revealed, as was expected, that the agents formed clusters according to matching positive triggers. A final run with agents having two quantitative tit-for-tat tasks showed that the complexity of the "emergent" phenomena may rise unexpectedly on the possibility of variants, as long as the agents seek individually to maximize satisfaction by compromising different, possibly conflicting motivations.

Evaluation. The potential of simple, local rules to account for complex spatial structures is by now a commonplace (see, for example, work in cellular automata, Wolfram 1986).[11] What we have attempted to demonstrate with these experiments is the versatility of social phenomena that a parameterizable and adaptive reciprocal agent model, like the quantitative tit-for-tat model, may account for. Phenomena such as cluster formation and territoriality, self-organization, stability, but also the role of aggressivity, direct sociality and conflicting motivations, have become apparent in the simple case studies presented.

[10] Humans are fortunately unlike animals in that they are able to scale this sociality, to learn different things and somehow "choose" to become social -or the opposite...

[11] The combined role of diversity and spatiality for the *evolution* of cooperation has been stressed in (Nowak & May 1992), (Nowak et al. 1995).

6 Lessons

We have presented a model of social behavior based on reciprocity and participation (implicit kinship). The model is a generalization of the well-known tit-for-tat model : generalization consists in having continuous satisfaction functions and cooperation criteria and giving room to parameter variations. After conducting those above case studies in social phenomena, we may now summarize the model's interesting features and its potential for further research :

- All agents are trying to solve their own instantiation of the one and only problem there is : individual satisfaction maximization (often called hedonism). Individual instantiation shows in the presence of different filtering mechanisms, satisfaction functions etc. This genetic variation may lead to complex phenomena such as formation and destruction of spatial clusters. The model also lends itself to adaptivity, as is demonstrated in section 3. This adaptivity and the possibility to scale the adaptation mechanism by continuously ascending levels is a good foundation for evolvability of the agents toward higher-level, more "cognitive" possibilities. Other extensions may be envisaged, too : combination of territoriality with adaptation (which could lead to speciation), integration with non-social tasks etc.
- Stability is a consequence of variations and interactions *("The real cause of stability in a distributed system is sufficient diversity"*, Hogg & Huberman (1992)). The dynamics of the system may become much more intricate when perturbations occur and the robustness of the system is not ensured when those perturbations are frequent and acute enough. Potential for self-organization is a consequence of the dynamic implicit kinship recognition mechanism ; of course, we always run the risk of having this self-organizing process turn into blind opportunism (free riders situation). These observations also imply that societies of quantitative tit-for-tat agents may be studied and analyzed using dynamical systems theory, as in (Huberman & Hogg 1988), (Hogg & Huberman 1991), (Huberman & Glance 1994).

We have emphasized on simulation rather than on analysis in an attempt to show the versatility of potential "emergent" phenomena. To give an idea of the complexity of the analysis of such a system, let us take the basic model of figure 1 for the case of two agents. The social interactions between two agents may be classified to one of four types (Hamilton 1964), as table 5 suggests.

Sender	*Receiver*	
	+	-
+	mutuality	deceit
-	eavesdropping	spite

Table 5 The four types of agent interactions according to Hamilton (1964) (+ and - denote the change in an agent's fitness after the interaction)

Agent i (i=1,2) may be completely specified by the quadruplet (T_i, $a_i()$, $b_i()$, $f_i()$). *Mutuality* emerges in the case of $y(1)=(a_1(f_1(1))+a_2(f_2(1)))/2>max(T_1,T_2)$, where $y(t)$ is the social stimulus quantity after t cycles. *Spite* emerges when $y(1)<min(T_1,T_2)$. In between lie a number of cases of *eavesdropping/deceit* (the two phenomena are complementary : if one agent plays first and gets eavesdropped, it is as though the

other one was playing first and was deceiving[12]). Eavesdropping/deceit will be continuous, i.e. the deceived agent will never "understand" it is being deceived, if (assuming agent-1 is the eavesdropper) :

$$f_1(y(1)) \geq T_1,\ f_2(y(1)) < T_2, \text{ with } y(1) = (a_1(f_1(1)) + a_2(f_2(1)))/2,$$
$$f_1(((a_1(y(t)) + b_2(y(t)))/2) \geq T_1,\ \forall t=1,2...,$$
$$f_2(((a_1(y(t)) + b_2(y(t)))/2) < T_2, \forall t=1,2...,$$
$$\text{with } y(t)=a_1(f_1(y(t-1)))+ b_2(f_2(y(t-1)))/2.$$

If there exists a *t* such that at least one of the above inequalities does not hold, this is the moment of time that the system will fall back on spite. Assuming more than just two agents specified by different quadruplets, we can easily see that the complexity of the analytical equations that describe the system rises tremendously.

Looking deeper into the inequalities, we may also observe that agent-i may be defined as "rational" if

$$f_i(x) \geq T_i \Rightarrow f_i(a_i(f_i(x))) \geq T_i, \text{ and}$$
$$f_i(x) < T_i \Rightarrow f_i(b_i(f_i(x))) < T_i$$ [13].

These inequalities mean that an agent is rational if he treats the results of its own actions rationally, i.e. if he detects the outcome of a cooperative/non-cooperative action of its own as cooperative/non-cooperative, respectively ; otherwise, and depending on the interaction context, he may appear as either a sucker or a con man. Of course, very few agents can be thought of as obeying this rule : most of the times, we encounter agents that have sophisticated filters and are inclined to believe, for instance, that they are smarter or better informed than their neighbors, or that they have some other kind of relative advantage. The *a()*, *b()* and *f()* functions have therefore a ***predictive role*** and adapting these functions corresponds to ***belief update***. In this sense, reactivity is no more than an implementation principle and mechanism ; in essence, the reactive agents have reactions pre-planned by their designer (this point was also made by Drummond (1993) and McFarland (1991 & 1992)). Those reactions are therefore based on "predictions" that may be adaptive, so that reactivity is no more reactivity to an *action*, but reactivity to the *meaning* of an action or *cognitive reactivity*, as Castelfranchi (1993) defined it.

Returning to the issue of ***rationality***, a consequence of individualistic and local motivation and perception structures that vary genetically from agent to agent is that talking about globally rational agents does no more make much sense. Pollock (1993) argued that rationality, which he defined as *"the tool for living the good life"*, is context-dependent and is constrained by the agent's knowledge, abilities and environment. The aggressive and scrooge consumers that manage to "survive" (be happy) in the artificial economies of section 3, do so because their internal parameters allow them to interact with the "rational" consumers in a way that makes everybody happy. It follows then that the aggressive and scrooge agents are no less rational than the others, since they manage to live a good life as well ! Looking deeper into what the supposedly "rational" or "irrational" economic agents do, it has been shown that the

[12] Note, however, that in this case, the deceiving agent is the one that defects while the other cooperates, whereas in participation contexts the agent that defects is always a loser.

[13] If *a()*, *b()* and *f()* are linear functions, $a(x)=k_a x$, $b(x)=k_b x$ and $f(x)=k_f x$, these inequalities translate into $k_b<(T_i/k_f^2)<k_a$ -note how the presence of the filtering function *f()* changes the shape of the perception space.

"rational" ones are less manipulable than the others (and this *without* possessing any higher cognitive abilities). The reason is that their actions/decisions are based more on internal than on external factors, so that their responsiveness to unexpected events is low : in this sense, ***what appears as rationality is in fact conservativeness and resistance to change/innovation***.[14] Note, however, that the average consumption of the economy increases slowly with time and this thanks to those aggressive agents. The social innovation and "progress" is therefore due to the presence of more manipulable, self-catalyzing agents (whether this "progress" serves a particular purpose is yet another story).

A final consequence of selfishness and individualistic rationality is that stability in social groups is precisely the result of heterogeneity and complementarity, as in the case of agents that need to form large working groups. This is achieved through variation between agents, which might be genetic or a result of differentiation through adaptation during life-time ; additional examples and a more extensive discussion of variation and differentiation appears in (Tzafestas 1995*b*). For example, we have seen above that the "rational" consumers are stable to all kinds of perturbations. So, what are "irrational" agents for ? The answer seems to lie in the social "progress" as defined above. If there is any reason to look for progress -in the sense of innovation- then we would better introduce such "irrational" agents that would ensure the social drift ; on the other hand, to guarantee stability of the society, the proportion of irrational agents should be low. Such a heterogeneous stable society is nothing else but an autonomous super-agent and stability nothing else but self-consistency of this super-agent.

7 Other related work

Besides work in evolutionary games and cooperation, the proposed model has also been partly stimulated by the Eco-Problem-Solving paradigm (EPS) and work in autonomous agents action selection (Maes 1989) (Ribeiro et al. 1992). EPS (Ferber 1989) (Ferber & Jacopin 1991) relies on an agent model possessing three binary state variables (satisfaction, flight, liberty), defining a total of six states (there are three unreachable/impossible states) and a set of primitive actions that change the values of those variables (Ferber & Jacopin 1990). The model has been used successfully to solve classical AI problems (Ferber & Jacopin 1991) (Drogoul 1993) traditionally tackled by planning methods. However, this paradigm does not include a measure of cooperation that might be used to direct a multi-agent system toward a collective consensus, since satisfaction is binary ; therefore, it has been applied with success only to cases that involve not cooperation but (spatial) action coordination[15]. The need for continuous satisfaction functions has been identified by researchers in the autonomous agents field (for instance, Steels (1994*a*)) and traces its roots to work in action selection (see, for example, the spreading activation scheme of Maes (1989), and the hybrid approach of Ribeiro et al. (1992) that fuses spreading of activation with eco-problem-solving principles). The necessity for such explicit representations of a system's behavioral goal, although not in the traditional sense of the term, has been

[14] It looks to me no coincidence that as people get older they become both more rational and more conservative.

[15] This formulation was given by Eric Jacopin.

also stressed by Zeghal (1994) in the domain of motion coordination. Finally, an approach to the sociality or cooperation problem by action learning through different types of reinforcement appeared in Mataric (1994) ; this last approach differs from ours in that we assume genetically determined actions and only parameter learning, if any at all.

8 Conclusions and perspectives

The above discussion and experiments have shown that the presented model possesses some key properties (adaptivity, scalability, variability) that make it a valid candidate for pursuing research in social behavior themes (including cooperation). The target behavior of the model is not limited to game-theoretic contexts, but involves more general participant relations, where agents “create” their own game dynamically without being subject to an externally forced static reward scheme. Diversity in parameters and structures gives rise to some intricate emergent phenomena ; on top of these, mobility induces elaborate spatial structures, such as clusters, formations etc. Conservativeness, in the sense of resistance to change enforces social stability, but has to be coupled with a diversity generator (or irrationality) in order for social standards to evolve.[16] The overall result in view of our long-term goals is that this model may account for a number of dynamic phenomena that may serve as a basis for the emergence of higher-order organizations (such a promising phenomenon is, for example, the formation of agent clusters in section 5).

Two immediate applications envisaged are the modeling of autonomous mobile robot cooperation tasks, such as the one described by McFarland (1994) and Steels (1994*b*)[17], and the distributed or cooperative problem solving, in the same line that is followed by Hogg & Huberman (1992). Further work includes a study in speciation and a study of the impact of diversity in an evolutionary context. Finally, we started investigating the role of sociality and cooperation for the emergence of higher-order organizations in (Tzafestas 1995*c*).

References

Angeline, P.J. (1994). An alternate interpretation of the iterated prisoner’s dilemma and the evolution of non-mutual cooperation, pp. 353-358, in R. Brooks and P. Maes (Eds.), *Artificial Life IV, Proceedings of the Fourth Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, MIT Press, Cambridge, MA, 1994.

Atlan, H. (1979). *Entre le cristal et la fumée - Essai sur l’organisation du vivant*, Éditions du Seuil, Paris.

Axelrod, R., W.D.Hamilton (1981). The evolution of cooperation, *Science*, 211:1390-1396, 1981.

Axelrod, R. (1984). *The evolution of co-operation*, Basic Books, 1984, also Penguin Books, 1990.

[16] The importance of diversity from a physics and an information science viewpoint to theoretical biology has been stressed by many authors, for example Schrödinger (1944), Atlan (1979), Dyson (1985). Similar observations on the social level have been made by Lévi-Strauss (1962) and Lumsden & Wilson (1981).

[17] Numaoka (1995) characterized this situation as the blind hunger dilemma.

Axelrod, R., D. Dion (1988). The further evolution of cooperation, *Science*, 242:1385-1390, 1988.

Batali, J., P. Kitcher (1994). Evolutionary dynamics of altruistic behavior in optional and compulsory versions of the iterated prisoner's dilemma, pp. 343-348, in R. Brooks and P. Maes (Eds.), *Artificial Life IV, Proceedings of the Fourth Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, MIT Press, Cambridge, MA, 1994.

Becker, G.S. (1976). Altruism, egoism and genetic fitness : Economics and sociobiology, *Journal of Economic Literature*, 14:817-826.

Beltratti, A., S. Margarita (1992). Evolution of trading strategies among heterogeneous artificial economic agents, pp. 494-501, in J.-A. Meyer, H. L. Roitblat and S. W. Wilson (Eds.), *From Animals to Animats 2, Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, Hawaii, 1992, Bradford MIT Press, 1992.

Benos, A., E. Tzafestas (1995). Alternative distributed models for the comparative study of stock market phenomena, *Proceedings 1995 Joint Conference on Information Systems (JCIS'95)*, Durham, NC, October 1995, also *Cahier de Recherche du Groupe HEC 551/1995*.

Bersini, H., V. Detours (1994). Asynchrony induces stability in cellular automata based models, pp. 382-387, in R. Brooks and P. Maes (Eds.), *Artificial Life IV, Proceedings of the Fourth Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, MIT Press, Cambridge, MA, 1994.

Bond, A.H., L. Gasser, Eds. (1988). *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, San Mateo, California, 1988.

Borges, J.-L. (1976). *¿Qué es el budismo ?*, Columba S.A., Buenos Aires, 1976.

Boyd, R. (1989). Mistakes allow evolutionary stability in the repeated prisoner's dilemma game, *Journal of Theoretical Biology,* 136(1989):47-56.

Castelfranchi, C. (1990). Social power : A point missed in multi-agent, DAI and HCI, pp. 49-62, *in (Demazeau & Müller 1990).*

Castelfranchi, C. (1993). Reactivity at the social level : Some basic issues, *Proceedings of 1993 Italian Workshop on DAI*, D. D'Aloisi & M. Miceli (Eds.).

Castelfranchi, C. (1995). Commitments : From individual intentions to groups and organizations, pp. 41-48, *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS95)*, San Francisco, CA, June.

Conte, R., M. Miceli, C. Castelfranchi (1991). Limits and levels of cooperation : Disentangling various types of prosocial interaction, pp. 153-165, in *(Demazeau & Müller 1991).*

Dawkins, R. (1976). *The selfish gene*, Oxford Univ. Press, 1976.

de la Maza, M., D. Yuret (1994). A futures market simulation with non-rational participants, pp. 325-330, in R. Brooks and P. Maes (Eds.), *Artificial Life IV, Proceedings of the Fourth Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, MIT Press, Cambridge, MA, 1994.

Demazeau, Y., J.-P. Müller, Eds. (1990). *Decentralized A.I., Proceedings of the First European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW '89)*, Elsevier/North-Holland, 1990.

Demazeau, Y., J.-P. Müller, Eds. (1991). *Decentralized A.I. 2, Proceedings of the Second European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW '90)*, North-Holland, 1991.

De Sousa, R. (1990). The sociology of sociobiology, *International Studies in the Philosophy of Science*, 4(3):271-283.

Drogoul, A. (1993). De la simulation multi-agents à la résolution collective de problèmes - Une étude de l'émergence de structures d'organisation dans les systèmes multi-agents, *Thèse de Doctorat d'Université Paris VI*, November 1993.

Drummond, M. (1993). Planning : From prediction to reaction (Tutorial on planning*), Talk delivered in University Pierre et Marie Curie*, Paris, Spring 1993.

Durfee, E.H., V.R. Lesser, D.C. Corkill (1989). Trends in cooperative distributed problem solving, *IEEE Transactions on Knowledge and Data Engineering*, 1(1) :63-83, March 1989.

Dyson, F. (1985). *The origins of life*, Cambridge University Press, Cambridge.

Ferber, J. (1989). Eco-Problem-Solving : How to solve problems by interactions, *Proceedings of the 9th Workshop on Distributed Artificial Intelligence*, 1989, pp. 113-128.

Ferber, J., E. Jacopin (1990). A multi-agent satisfaction planner for building plans as side effects, *Rapport LAFORIA 07/90*, January 1990, 11p.

Ferber, J., E. Jacopin (1991). The Framework of Eco-Problem-Solving, pp. 103-114, in *(Demazeau & Müller 1991).*

Findler, N.V., R. Malyankar (1993). Alliances and social norms in societies of non-homogeneous, interacting agents, *Proceedings 2nd International Symposium on Simulation Societies '93 (Approaches to simulating social phenomena and social processes)*, July 1993.

Gabora, L.M. (1993). Meme and variations : A computational model of cultural evolution, *Proceedings of the Complex Systems Summer School*, 1993, 18p. Also Poster Presentation*, Artificial Life IV Workshop*, Cambridge, MA, 1994.

Galliers, J.R. (1990). The positive role of conflict in cooperative multi-agent systems, pp. 33-46, in *(Demazeau & Müller 1990).*

Gasser, L., M.N. Huhns, Eds. (1989). *Distributed Artificial Intelligence, Vol. 2*, Pitman Publishing/Morgan Kaufmann, London/San Mateo, California, 1989.

Glance, N.S., B.A. Huberman (1994). The dynamics of social dilemmas, *Scientific American*, March 1994, pp. 58-63.

Hamilton, W.D. (1964). The genetical evolution of social behaviour - I & II*, Journal of Theoretical Biology*, 7(1964):1-16 and 17-52.

Heylighen, F. (1992*a*). Evolution, selfishness and cooperation, *Journal of Ideas*, 2(4):70-76, 1992.

Heylighen, F. (1992*b*). "Selfish" memes and the evolution of cooperation, *Journal of Ideas*, 2(4):77-84, 1992.

Hogg, T., B.A. Huberman (1991). Controlling chaos in distributed systems, *IEEE Transactions on Systems, Man and Cybernetics, Special Issue on Distributed Artificial Intelligence*, 21(6):1325-1332, November/December 1991.

Hogg, T., B.A. Huberman (1992). Better than the best : The power of cooperation, in L. Nadel and D. Stein (Eds.), *SFI 1992 Lecture Notes in Complex Systems*, pp. 163-184, Addison-Wesley, 1993.

Huberman, B.A., T. Hogg (1988). The Behaviour of Computational Ecologies, pp. 77-115, in B.A. Huberman (Ed.), *The Ecology of Computation*, North-Holland, Amsterdam, 1988.

Huberman, B.A., N.S. Glance (1993). Evolutionary games and computer simulations, *Proceedings National Academy of Sciences (USA)*, Vol. 90, pp. 7716-7718, 1993.

Huberman, B.A., N.S. Glance (1994). Beliefs and cooperation, *Proceedings International Conference on "Chaos and Society"*, June 1994.

Huhns, M.N., Ed. (1987). *Distributed Artificial Intelligence*, Pitman Publishing/Morgan Kaufmann, London/San Mateo, California, 1987.

Lévi-Strauss, C. (1962). *La pensée sauvage*, Librairie Plon, Paris.

Lumsden, C.J., E.O. Wilson (1981). *Genes, mind and culture - The coevolutionary process*, Harvard University Press, Cambridge, MA.

Maes, P. (1989). The Dynamics of Action Selection, *Proceedings IJCAI-89*, Detroit, 1989, pp. 991-997.

Mataric, M.J. (1994). Learning to behave socially, pp. 454-462, in D. Cliff, P. Husbands, J.-A. Meyer and S.W. Wilson (Eds.), *From animals to animats 3, Proceedings of the 3rd International Conference on Simulation of Adaptive Behavior*, Bradford/MIT Press, Cambridge, MA, 1994.

May, R.M. (1987). More evolution of cooperation, *Nature*, 327:15-17.

McFarland, D. (1991). Defining motivation and cognition in animals, *International Studies in the Philosophy of Science*, 5(2):153-170, 1991.

McFarland, D. (1992). Animals as cost-based robots, *International Studies in the Philosophy of Science*, 6(2):133-153, 1992.

McFarland, D. (1994). Towards robot cooperation, pp. 440-444, in D. Cliff, P. Husbands, J.-A. Meyer and S.W. Wilson (Eds.), *From animals to animats 3, Proceedings of the 3rd International Conference on Simulation of Adaptive Behavior*, Bradford/MIT Press, Cambridge, MA, 1994.

Nottola, C., F. Leroy, F. Davalo (1992). Dynamics of artificial markets : Speculative markets and emerging "common sense" knowledge, pp. 185-194, in F.J. Varela and P. Bourgine (Eds.), *Toward a practice of autonomous systems, Proceedings of the First European Conference on Artificial Life*, Paris, December 1991, MIT Press/Bradford Books, 1992.

Nowak, M., K. Sigmund (1989). Oscillations in the evolution of reciprocity, *Journal of Theoretical Biology*, 137(1989):21-26, July.

Nowak, M.A., R.M. May (1992). Evolutionary games and spatial chaos, *Nature*, 359:826-829, 1992.

Nowak, M., K. Sigmund (1992). Tit-for-tat in heterogeneous populations, *Nature*, 355:250-253.

Nowak, M.A., K. Sigmund (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game, *Nature*, 364:56-58, July.

Nowak, M.A., R.M. May, K. Sigmund (1995). The arithmetics of mutual help, *Scientific American*, 262(6):50-55, June.

Numaoka, C (1995). Introducing the blind hunger dilemma: Agents' properties and performance, pp. 290-296, *Proceedings ICMAS'95*, San Francisco, CA.

Pollock, J.L. (1993). The phylogeny of rationality, *Cognitive Science*, 17:563-588.

Ray, T.S. (1994). An evolutionary approach to synthetic biology : Zen and the art of creating life, *Artificial Life*, 1(1&2):179-209, Fall 1993/Winter 1994.

Ribeiro, F., J.-P. Barthès, E. Oliveira (1992). Dynamic selection of action sequences, pp. 189-195, in J.-A. Meyer, H.L. Roitblat and S.W. Wilson (Eds.), *From Animals to Animats 2, Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, Hawaii, 1992, Bradford MIT Press, 1992.

Schrödinger, E. (1944). *What is life ?*, Cambridge University Press, Cambridge.

Sigmund, K. (1992). On prisoners and cells, *Nature*, 359:774, 1992.

Steels, L. (1994*a*). Emergent functionality of robot behavior through on-line evolution, pp. 8-14, in R. Brooks and P. Maes (Eds.), *Artificial Life IV, Proceedings of the Fourth Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, MIT Press, Cambridge, MA, 1994.

Steels, L. (1994*b*). A case-study in the behavior-oriented design of autonomous agents, pp. 445-452, in D. Cliff, P. Husbands, J.-A. Meyer and S.W. Wilson (Eds.), *From animals to animats 3, Proceedings of the 3rd International Conference on Simulation of Adaptive Behavior*, Bradford/MIT Press, Cambridge, MA, 1994.

Trivers, R. (1971). The evolution of reciprocal altruism, *Quarterly Review of Biology*, 46(4):35-57, 1971.

Tzafestas, E.S. (1995*a*). Beyond cooperation-ism and competition-ism (Exploring social phenomena with a generalized tit-for-tat model) p. 464*, Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95)*, San Francisco, CA, June.

Tzafestas, E.S. (1995*b*). Vers une systémique des agents autonomes : Des cellules, des motivations et des perturbations*, Thèse de Doctorat de l'Université Pierre et Marie Curie*, Paris, December.

Tzafestas, E.S. (1995*c*). Aging agents*, LAFORIA Research Report*, to appear.

Werner, E. (1989). Cooperating agents : A unified theory of communication and social structure, pp. 3-36, in *(Gasser & Huhns 1989)*.

Werner, E. (1990). Distributed cooperation algorithms, pp. 17-31, in *(Demazeau & Müller 1990).*

Wilson, E.O. (1975). *Sociobiology - The new synthesis*, Belknap Press/Harvard University Press, Cambridge, MA, 1975.

Wolfram, S., Ed. (1986). *Theory and applications of cellular automata*, World Scientific, Singapore, 1986.

Zeghal, K. (1994). Un modèle de coordination d'actions pour agents mobiles, *Rapport de recherche LAFORIA, 94/09*, June 1994, 24p.