

Are Incentives Good Enough To Achieve (Info)Social Order?

Rosaria Conte

Cristiano Castelfranchi

Division of AI, Cognitive & Interaction Modelling
IP-CNR, V.le Marx 15 - Rome - I-00137 - Italy
email: rosaria&pscs2.irmkant.rm.cnr.it

Abstract

In this paper, the role of incentives in social order is questioned, based on a notion of incentive as additional individual utility, provided by an external entity, to actions achieving global utility. Two notions of norms are compared: (1) inputs which modify agents' decisions through incentives (sanctions) and (2) prescriptions to execute obligatory action for intrinsic motivations. Two types of agents which reason upon norms are also compared: (1) incentive based rational deciders, and (2) normative agents which are prescribed to execute norms for intrinsic reasons. The two types of agents are expected to have a different impact on norm compliance. Under suboptimal conditions of application of sanctions (uncertain punishment), transgression is expected to propagate more easily and rapidly among incentive-based agents than among normative agents. In particular, incentive-based agents are expected to show a fast decline and even a collapse in compliance with the norms. Normative agents are expected to exhibit an oscillating behaviour, or at least a graceful degradation of compliance. Finally, the role of incentives is shown to have a lesser impact on natural social agents than expected by a model of rational decision. What is worse, incentives have been shown to produce even negative effects on several aspects of social learning and norm compliance.

1 The Problem of Social Order in Agent Mediated Interaction

The problem of social order in natural societies is a traditional concern of social philosophers and social scientists. With some right, rational action theory defines it as a dilemma, a yet unsolved (and insoluble?) incompatibility between individual and global utility. If individual agents are rational, that is, if they act to maximise their individual utility, they will inevitably achieve a globally undesirable state of affairs (some will gain at the expense of others). Moreover, individual utility maximisation has long-term self-defeating effects, since all agents (supposedly rational), will exploit and be exploited at once. Not surprisingly, the traditional (analytical) solution to this dilemma proposed by rational action scientists is a forward-looking agent, which calculates the short and long-term effects of an action. But with bounded rationality, a forward-looking agent cannot accomplish a thoroughly rational action. Hence, the necessity for means, social and institutional, designed to regulate societies and achieve a socially acceptable state of affairs. Rational agents' decisions must be modified through positive or negative incentives (sanctions). Indeed, sanctions and incentives provide *additional individual utility for self-interested agents to act according to some global utility*.

To achieve social order has become an urgent problem in agent mediated interaction, both in software agent-human agent, and in software agent-software agent interaction. This should not come as a surprise, since human agents are self-interested and software agents are designed to act in the interest of and on behalf of their users (Rao, 1998; Crabtree, 1998). In info societies as well as in natural societies, local and global utility are often incompatible, and individual utility maximisation is found to produce long-term self-

defeating effects (Crabtree, 1998). We will then speak of the problem of *infosocial* order as a new version of the old problem, which calls for much the same means and solutions already experienced in natural societies. Software agents scientists and designers are well aware of this necessity, as is documented by the recent studies in agent mediated electronic commerce (Dignum, 2000). Of late, the problem of infosocial order gave rise to a new field of investigation, i.e. the field of electronic institutions (see again, Dignum, 2000). Given the impact of rational action and game theory on the multiagent systems field, one could expect that means implemented to achieve infosocial order are inspired by the same principle mentioned above, that is, to provide additional utility for software agents to act according to existing institutions. Hence, the efficiency of electronic institutions is expected to rely upon the efficiency of sanctions and incentives as inputs to software agents' rational decisions.

In this paper, we intend to investigate different ways of implementing agents which reason upon norms. As will be shown in the next section, norm-based reasoning is not necessary to obtain a norm-corresponding behaviour. Other mechanisms have been implemented at the agent level, which do not allow for reasoning and decision based upon norms. These will be shortly examined in the next section. The focus of this paper is on intelligent norm-driven action, that is, on action based upon a decision to comply with some norm. In particular, we intend to question the efficacy of sanctions and incentives in the achievement of social order. In natural societies, the efficacy of sanctions and incentives is far from granted. Moreover, human agents do not always act upon rational decision, and normative action is not executed only when compliance is convenient in terms of individual utility. In natural societies, norms are even expected to be observed for intrinsic reasons, as ends in themselves. But what about infosocieties? Which impact can incentives be expected to have on the achievement of infosocial order? After a short review of earlier work on the implementation of social laws and conventions, we will compare two different views of agents reasoning about norms and other institutions:

- Incentive-based rational deciders
- Normative agents, which are prescribed to be intrinsically motivated to comply with the norms.

In the following section, we will formulate specific consequences that can be expected from either type of agents. Thereafter, some evidence from natural societies will be shown to match the expectations relative to normative agents, rather than those relative to rational deciders. Some speculations about the relative desirability of normative Vs. rational agents will conclude the paper.

2 Related Work

Attempts to implement laws and conventions at the level of the agent go back to the early 90s. The necessity to achieve coordination in motion has inspired the implementation of social laws in multiagent systems (Shoham and Tennenholtz, 1992). Analogously, the necessity for robust performance in joint action and teamwork inspired the implementation of commitment (Cohen and Levesque, 1990a; Kinny & Georgeff, 1994) and conventions (Jennings and Mandami, 1992; Jennings, 1995), and other norm-like mechanisms (such as responsibility, cf. Jennings, 1992). These models and the corresponding systems

present a twofold problem. On one hand, norms and laws are implemented as action constraints, which ensure optimal efficiency of the system, but grants no agent autonomy: agents cannot violate the norms. On the other, no innovation is allowed online: agents are not enabled to acquire norms. These can be modified and updated by the programmer when the system is online.

Impulse to the implementation of norms as input to agents' reasoning and decision-making comes from the rational action theory (Boman, 1999), which has hegemonial influence in multiagent systems (for a critique, see Castelfranchi and Conte, 1998). Based on the assumption that decisions are guided by agents' subjective preferences, norms are seen as external inputs to agents' decisions, because they modify the agents' preference order and therefore their utility function through sanctions (if you violate the norm, you will get a penalty) or positive incentives (if you observe the norm, you will get a reward). This conceptualisation of norms does justice to the autonomy of agents, which are not simply constrained by norms, but decide whether to execute them or not on the grounds of their criteria. At the same time, it allows norms to be updated and acquired online. Agents will receive (through communication) new normative inputs and the associated incentives, which they will take into account when acting. In this view, a norm-driven action is an autonomous action which implies norm-based reasoning and decision.

As will be shown in the next section, rational decision is one of the two major conceivable ways to implement intelligent and autonomous norm-driven action. Let us see the other and compare them.

2 How to Implement Norm-Based Reasoning?

There are two main approaches to implement agents that reason and decide upon norms. These approaches depend upon two different notions of a norm, norms as prescriptions to execute an action based upon incentive (we will call this incentive-based norms); norms as prescriptions to execute an action for intrinsic reasons.

In the case of incentive-based norms, a norm provides additional individual utility (usually through sanctions) for socially desirable action. Agents which are built upon this notion of norm will be here called rational deciders; rational decision is based upon a set of ordered preferences which gives rise to a subjective utility function: given a choice, the agent will be rational if it performs the action which maximises its utility. Sanctions and incentives must be provided in such a degree that the individual utility of socially desirable action is higher than the individual utility of socially undesirable action. Later on in the paper we will examine a number of specific effects that can be expected from this modality of norm implementation.

According to the second conceptualisation, norms are seen as prescriptions to execute actions based upon *intrinsic motivations* (we will call them motivating norms). It has been shown (cf. Conte & Castelfranchi, 1999) that one important aspect of normative prescriptions lies in the reasons why they ought to be adopted. In this conceptualisation, a norm is more than a simple prescription about a given action or decision. A norm prescribes not only *what* must (not) be done, but also *why* it must (not) be done. Sanctions are not inherent to norms. They are consequences of transgression, rather than reasons for obedience.

Norms prescribe that they be adopted just because they are norms, as intrinsic motivations. Indeed, this is the first and most important criterion for telling whether a given command is a norm or not. There are two orders of evidence that this is the case. First, agents (at least natural agents) are not (necessarily) informed about the entity of sanctions, nor will they feel entitled to ask questions about it, and still they can tell if the command is a norm or not. When you get on a flight, you do not ask the staff members what is the sanction for smoking, although probably you have never known it. Agents which are built upon this notion of norm are here called normative agents.

The main difference between these two notions of norms is that sanctions (or incentives) are inherent to the former but not to the latter. Consequently, the main difference between rational deciders and normative agents is that the former will execute the norm only in presence of sanctions or incentives, while the latter are requested to have an intrinsic motivation to obey the norm. This difference needs further consideration.

First, normative agents do often adopt the norm for utilitarian reasons. But this is only a sub-ideal (in the sense logically defined by Jones and Pšrn, 1991) state of affairs. With rational deciders, the utilitarian calculation is not sub-ideal: rational deciders are not prescribed specific reasons for obedience.

Secondly, normative agents are expected to decide whether to adopt a norm even if sanctions are not specified. Such a condition, conversely, is undecidable for rational deciders: they will have no sufficient elements to decide.

Third, and consequently, there can always be a subset, however small, of normative agents which will adopt the norm for ideal reasons (intrinsic motivations). But rational deciders cannot accept norms for intrinsic reasons, unless sanctions are also intended as internal. In such a case, of course, no significant difference holds between normative agents and rational deciders. This leads us to precisely what is here meant by sanctions and incentives.

2.1 Incentives and Sanctions

Here, we will provide an operational notion of incentive. We speak about a positive incentive as an additional expected benefit of an action. More precisely, ws_i is an incentive for agent ag_i to perform a given action a_i when

- agent ag_i does action a_i for any given goal g_i , and
- a_i brings about ws_i , a state of the world which achieves a further goal of ag_i , say, goal g_j , and
- ws_i increases the value (or the utility) of a_i , so that ag_i is more likely to perform it as a means to achieve g_i .

The world state ws_i , then, is a positive side-effect of a_i , an incentive. A sanction is a negative side-effect. Agent ag_i may be informed about the side-effects of a_i , and still this action is not initialised by any of them, but by the agent's original goal g_i . Nonetheless, actions' side-effects obviously interfere with agents' planning and decision-making. Suppose you want to get warm. Your plan library suggests several alternative plans: to turn on the heater, to wear warm clothes, or to make a fire. Suppose that the wooden

fire has a nice effluvium. Nexttime you want to get warm, you'll probably choose again the wooden fire, because the effluvium acted upon you as an incentive.

We will speak of a *social incentive* (positive or negative), when an incentive is controlled (provided or not) by another entity ag_j , where $ag_i \neq ag_j$. More precisely, agent ag_i has a social incentive¹ to execute a_i , when

- ag_j has the power to bring about (or to obstacle) ws_i
- ag_j has the goal to influence ag_i to execute a_i , that is has the goal that ag_i decide to execute a_i
- ag_j believes that goal g_i of agent ag_i is insufficient for ag_i to put a_i to execution
- ag_j believes that ws_i will increase the value of a_i for ag_i , and therefore the probability that ag_i will execute it
- ag_j gets ag_i to know that if ag_i will perform a_i , ag_j will bring about (or prevent) ws_i .

A social incentive is therefore an additional value or utility, provided by an external entity, which modifies the agent's decision. Such an external entity must have the power or capacity of bringing about a world state relevant for a_i 's goals. This will turn into a social power of ag_j 's: thanks to the power of bringing about ws_i , ag_j has also power over ag_i . Agents may control and influence other agents also by providing incentives to them.

2.2 Incentive-Based Rational Deciders

Rational deciders calculate the subjective expected value of actions according to their utility function. According to a classical strategy of choice, given an agent ag_i and a set of alternatives for action $A_i = a_1, \dots, a_n$, the value of each alternative (taking into account its costs) per its relative probability of occurrence will be compared. That which yields the maximum utility (including the alternative "don't act") will be put to execution.

How is it possible to have rational deciders to observe a norm? First, they must be informed about the norm. Agents must be provided with criteria to recognise norms. For example, a norm may be a command imposed by a given authority and associated with given incentives (usually, negative). With norms, a rational decider will perform the same calculation which is applied to any other decision: the utility of norm compliance is computed in the usual way:

$$v_c p_i + v_t (1 - p_i)$$

where v_c is the value of compliance; to simplify matters, we assume this value to be equal to the incentive (or sanction); p_i is the probability of its occurrence, and v_t stands for the value of transgression, which is always positive since a normative action is by default inconvenient. From this, it can be easily drawn the conclusion that if an incentive is lower than the value of transgression, a rational decider will not comply with the norm, unless the probability of incentive is lower than the complementary probability (to not receive incentive, or, which is the same, to undergo sanctions).

2.3 Normative Agents

¹>From now on, we will speak about social incentives, but will call them incentives for short.

Normative agents are cognitive agents which are prescribed to adopt norms as ends in themselves. Normative agents are here seen as BDI-like agents, characterised by mental states, namely goals and beliefs, and the capacity to reason and act upon them. Normative agents form beliefs about a norm, may decide to adopt it by forming a corresponding goal, and to achieve it, by executing a norm-driven action.

Ideally, norms not only prescribe a given action a_i , but also a given motivation for executing it, i.e. the goal to comply with the norm because it is a norm.

In the following two sub-sections, we will resume our model of norms presented elsewhere (cf. Conte & Castelfranchi, 1999), and will show how such a model accounts for the motivations prescribed by the norms.

2.3.1 Our formalism

The formalism used is a simplified version of Cohen and Levesque's (1990b) language for describing their theory of rational action. The language appears as a first-order language with operators for mental attitudes and action. Two modalities for beliefs and goals ($BEL\ x\ p$) and ($GOAL\ x\ p$) are defined according to the possible worlds semantics, and therefore through accessibility relations. Two modalities for action ($HAPPENS\ e$) and ($DONE\ a$) express, respectively, events taking place in the world independent of the agents' actions and occurrence of actions. Finally, time is represented as an infinite sequence of events. Beliefs and goals are given the usual possible world interpretation. As for consistency, the Hintikka axioms for beliefs apply to this model (see Halpern & Moses 1985). As for realism, goals are a subset of beliefs. (The accessibility relation G , which defines the set of worlds in which goals are achieved is a subset of the accessibility relation B , which defines the set of worlds belief-accessible to a given agent.). In such a model, in fact, a goal is defined as a belief-compatible desire. (In other words, agents cannot have goals which they believe to be unachievable.)

Many notions can be constructed on the ground of these primitive modalities plus the operators \Diamond for "later", $;$ for "sequence" and $?$ for the procedure to test whether a given proposition is true.

$(HAPPENS\ a)$	an action will happen next
$(DONE\ a)$	an action has just happened;
$(BEL\ x\ p)$	x has p as a belief
$(GOAL\ x\ p)$	x has p as a goal;
$(OUGHT\ p)$	there is an obligation whatsoever on proposition p ;
$(AGT\ x\ e)$	x is the only agent of this sequence e ;
$e_1 \leq e_2$	e_1 occurs before e_2
$p?$	test action
$\Diamond p$	p will be true at some point in the future

A number of definitions, grounded upon the above atomic predicates, are necessary to understand the formulae provided throughout the paper. Most of them are drawn from Cohen and Levesque's model, and we present them here for the convenience of the reader unacquainted with that model. Some have

been introduced by the authors and other collaborators in preceding works (Conte et al. 1991; Castelfranchi et al. 1992).

$$(DOES\ x\ a) \stackrel{def}{=} (HAPPENS\ a) \wedge (AGT\ x\ a) \quad (1)$$

This says that *x is the only agent of action a, which will happen next*. We need an analogous predicate for past actions,

$$(DONE - BY\ x\ a) \stackrel{def}{=} (DONE\ a) \wedge (AGT\ x\ a) \quad (2)$$

saying that, *x is the only agent of action a, which has just happened*.

Cohen and Levesque have also introduced the following predicate to refer to sequences of world states,

$$(BEFORE\ q\ p) \stackrel{def}{=} \forall c (HAPPENS\ c; p?) \supset \exists a (a \leq c) \wedge (HAPPENS\ a; q?) \quad (3)$$

In words, *q comes before p when, for all events c after which p is true, there has been at least one event a preceding c, after which q was true*.

As for goals, Cohen and Levesque have introduced the notion of achievement goal, which is defined as follows:

$$(A - GOAL\ x\ p) \stackrel{def}{=} (BEL\ x\ \neg p) \wedge (GOAL\ x\ \Diamond p) \quad (4)$$

that is, *x has an achievement goal p if x believes that p is not true now but wants it to eventually become true*. Throughout the paper, whenever the notion of goal is used, it will be meant as an achievement goal in the above sense, unless otherwise specified. Indeed, in our model (as well as in Cohen and Levesque's), an achievement goal is not yet an intention.

Cohen and Levesque's theory includes a notion of relativised goal:

$$(R - GOAL\ x\ p\ q) \stackrel{def}{=} (A - GOAL\ x\ p) \wedge \\ (BEFORE\ ((BEL\ x\ \neg q) \vee (BEL\ x\ p) \vee (BEL\ x\ \neg \Diamond p)) \\ \neg (A - GOAL\ x\ p)) \quad (5)$$

x has a goal p relativised to q, when x has an achievement goal p, and before ceasing to have p as an achievement goal, x believes either that p is realised or unachievable or that the escape condition q does not hold. Essentially, this means that x has p as long as and because he believes that q.

Our notion of a goal (Conte and Castelfranchi 1995a) is slightly weaker than that allowed by Cohen and Levesque. We propose to treat goals as *realistic* desires, rather than *chosen* ones. In our terms, a goal is but a regulatory mental attitude which calls for a series of operations, including some preliminaries, involved in planned action. In other words, along the lines of classical AI planning systems, we define a goal as a device which activates planning and *action*. In our terms, a goal may be abandoned not only when it is believed to be fulfilled or unachievable, but also when it is found incompatible with another more important goal.

2.3.2 Normative Beliefs and Goals

The pred *OUGHT* intuitively means that there is some sort of *obligation* on proposition p. For the time being, we take it as an atomic one-place predicate, although it seems possible to further analyse it as some sort of external reason which forces a given goal, namely the adoption of a given goal. However, we will

assume obligation as a primitive, which defines a set of worlds in which p follows from obligations. The relation of accessibility \mathbf{O} is a subset of \mathbf{B} .

In our model, agents have normative beliefs when they think there is an obligation on a given set of agents to do some action.

In the following, x and y denote agent variables with $x - y$ always implicitly stated, and a denotes an action variable.

We express the general form of a normative belief as follows:

$$(N - BEL \ x \ y_i \ a) \stackrel{def}{=} (\bigwedge_{i=1,n} (BEL \ x (OUGHT (DOES \ y_i \ a)))) \quad (6)$$

in words, x has a normative belief about action a relative to a set of agents y_i if and only if x believes that it is obligatory for y_i to do action a . The predicate *OUGHT* stands for an obligation for a set of agents y_i to do action a . A few words are needed to elucidate the semantics of our predicate *OUGHT*. This stands for an operator of obligation about any given state of the world. However, it should be taken in a somewhat weaker sense than what is usually meant by obligation in traditional deontic logic. In fact, while in traditional deontic systems, p necessarily follows from obligation (that is to say, it is not possible that at the same time p is false and obligatory), in other systems (Jones and Pšrn 1991), two concepts need to be distinguished, one referring to deontic necessity and the other to another type of obligation. The latter is defined as the circumstance in which a given proposition is both obligatory and possibly false in some sub-ideal world.

In order to express normative goals, another belief is needed, namely a pertinence belief: for x to believe that he is addressed by a given norm, he needs to believe that he is a member of the class of agents addressed by that norm:

$$(P - N - BEL \ x \ a) \stackrel{def}{=} (\bigwedge_{i=1,n} (N - BEL \ x \ y_i \ a)) \wedge (\bigvee_{k=1,n} (BEL \ x (x = y_k))) \quad (7)$$

where *P-N-BEL* stands for normative belief of pertinence; in words, x has a normative belief of pertinence when he has a normative belief relative to a set y_i and an action a , and believes that he is included in y_i .

Now, x 's beliefs tell him not only that there is an obligation to do action a , but also that the obligation concerns precisely himself.

We have not seen any normative goal yet. A normative goal is defined here as a goal always associated with and generated by a normative belief. Let us express a normative goal as follows:

$$(N - GOAL \ x \ a) \stackrel{def}{=} (R - GOAL \ x (DOES \ x \ a) (P - N - BEL \ x \ a)) \quad (8)$$

or, x has a normative goal concerning action a when he has the goal to do a relativised to his pertinence normative belief concerning a . A normative goal of a given agent x about action a is therefore a goal that x has as long as he has a pertinence normative belief about a : x has a normative goal in so far as he believes to be subject to a norm.

2.3.3 The Paradox of Normative Requests

What is the relationship between a normative belief and a normative goal? This question should be examined from two different perspectives.

From the point of view of the agent, a normative belief is *necessary* but not sufficient for a normative goal to be formed, and *a fortiori*, a normative action to be executed. Elsewhere (Conte and Castelfranchi, 1995a), we have examined several mechanisms of norm adoption, including instrumental and cooperative adoption. In other words, there may be several reasons for agents to adopt a norm: to avoid sanctions, to achieve positive side-effects (incentives), or even to achieve a goal which the norm is able to instore. In the latter case, the agents have one goal in common with the norms, or, better, with the system which has issued the norm.

From the point of view of the norm itself, a normative belief is not only *necessary* but ought to be also *sufficient* for a normative goal to be formed. Agents must know that action is obligatory (N-belief) to have a normative goal concerning that action. On the other hand, if they have a normative belief, they ought to want to execute it.

$$(N - BEL\ x\ y_i\ a) \supset (BEL\ x(OUGHT((P - N - BEL\ y_i\ a) \supset (N - GOAL\ y_i\ a)))) \quad (9)$$

Sub-ideally, this may not be the case. *ought* to be the case; this is what the norm expects. Indeed, this is how a norm can be distinguished from other, coercive, requests or commands. All that a norm says is what must be done: provided the agent is dutifully informed about it, it will have a normative will corresponding to it. Sanctions are consequent to action violations, and reasonably effects which agents learn to associate to it. In real matters, negative or positive incentives have a strongly motivating role in norm compliance. But on the one hand, this is not always and necessary the case: norms may and sometimes are observed for intrinsic reasons. On the other hand, this is a *sub-ideal*, however frequent, state of affairs (Jones and Porn, 1991), meaning that only in a subset of the worlds in which the norm is in force, a normative belief is sufficient for a normative goal to arise and the corresponding action to happen. This subset is that of ideal worlds. In sub-ideal worlds, that is the complementary subset, a normative belief is only a necessary but insufficient condition for a normative goal, and the latter is a necessary but insufficient condition for a normative action.

3 What Can Be Expected?

Which expectations can be made with regard to the effects of the two architectures? Both types of agents can violate the norm, since both types of agents are autonomous. Rational deciders will violate a norm when it is inconvenient for them to comply with it. Normative agents can violate a norm for a number of reasons, which include but are not reduced to utilitarian reasons, for example to solve a normative conflict. In case of a conflict between two norms, rational deciders are expected to choose that which is most convenient, or least inconvenient to them. On the contrary, normative agents are expected to apply the most important one, irrespective of their own convenience. Furthermore, normative agents can violate a norm which they consider unfair.

More explicitly, we can formulate two general expectations:

- incentive-based deciders will comply with the norms to the extent that the (positive or negative) incentive is such that the utility of obedience is higher than the utility of transgression (sanction is higher than the convenience of transgression);
- normative agents will comply with a norm as long as either ideal conditions apply (intrinsic motivations) or sub-ideal conditions apply (in this case they will behave as rational deciders) or ideal conditions apply and the norm is not unfair or contrary to duty.

3.1 Rational Deciders' Impact

More specifically, incentive-based deciders will violate a norm n_i as soon as one or more of the following conditions applies:

- Sanctions are not imposed: an incentive-based decider will certainly violate a norm if no sanction is expected to follow from violation, since by definition in absence of incentives norm compliance is individually irrational.
- Sanctions are expected but are not specified: in such a condition a rational decider will either infer the specification of sanctions, or will not take any decision.
- The sanction² for violating n_i is lower than the value of transgression with equal probability of application of the sanction (1/2).
- The sanction (negative incentive) for violating an incompatible norm n_j , where $(n_i \rightarrow n_j) \wedge (n_i \supset \neg n_j)$ is higher. This aspect of norm-based decision-making is important especially in societies of growing complexity, where the set of norms tends to increase, and conflicts among norms become more likely.
- The sanction (negative incentive) for violating the norm n_i is not or rarely applied: p_i tends to 0. Since the utility of norm compliance, as seen above, is equal to the value of incentive (or sanction) per its relative probability of occurrence (taking into account the utility of transgression), obviously with a probability proximate to zero, the utility of incentive is also nullified. Therefore, even with a moderately convenient value of transgression, a rational decider is likely to violate the norm. Consider that both the probability and entity of sanctions may be inferred by observing others' behaviour: the more others violate, the less likely and/or severe the sanction is expected to be. This has further consequences which we will examine in the following section.

With a homogeneous society of incentive-based deciders, any of the above conditions is followed by a fast decline or even a collapse in compliance with a given norm. The inconvenience of norm compliance will be detected sooner or later by all members of the society. Consequently, their decisions will rapidly converge on norm violation, unless the external source of sanctions monitors the behavioural effects of agents' decisions and takes efficient measures of norm enforcement, by either intensifying the application of sanctions or augmenting their entity.

²>From now on, we will speak of sanctions rather than incentives, because norms are enforced by sanctions more than by positive incentives. However, the formal reasoning can easily be extended to the other factor of enforcement.

3.2 Normative Agents' Impact

On the other hand, normative agents are expected

- To comply with a norm even if sanctions are not imposed, or are not imposed explicitly. *A fortiori*, normative agents may comply with norms when they know that sanctions are imposed but their entity and probability of application is uncertain.
- To execute norms even though sanctions are such that the utility of norm compliance is lower than the utility of transgression. A heterogeneous population of normative agents, where ideal and sub-ideal agents co-exist, ensures that even a small subset of agents will still apply the norm for intrinsic reasons.
- To comply with the norm n_i even when sanction is not or rarely applied. This is but a special case of the previous point. Of course, sub-ideal agents will converge on norm transgression. However, an even small number of stubborn agents will comply with a norm even when the sanctions are not or rarely applied.
- To comply with the norms when others violate. A persistent execution of the norm in a small share of the population (ideal agents) is expected. This has interesting further effects at the global level: since sub-ideal agents, as well as rational deciders, are able to infer the entity and probability of incentives by observing others' behaviours, some persistence in norm execution will have the consequence to limit or counteract this inference. Some oscillatory effects can be expected: agents which perceive ideal agents' behaviours will draw different conclusions on the entity/application of sanctions than others and will therefore be more likely to execute the norm. But as they perceive the behaviours of other sub-ideal agents, who were not exposed to the influence of ideal ones, they will go back to violation. Indeed, even ideal normative agents may be affected by others' decisions. Frequent transgressions may be perceived as "demotivating": the more a given norm is violated, the more it is perceived as unfair or inadequate or ineffective. This perception may reduce an intrinsic motivation to comply with that norm. However, no collapse in norm compliance is expected with normative agents but rather a "graceful" and non-linear degradation³.
- To solve norm conflicts even independent of the respective sanctions: with $(n_i - n_j) \wedge (n_i \supset \neg n_j)$
normative agents are not necessarily expected to choose the norm which grants them the higher individual utility. Again, an even small number of ideal normative agents will still choose the norm which is more important than the other according to some plausible criteria (entity of the injury consequent to norm transgression, reparability of norm transgression, etc.).

In short, the general expectation that incentives are a good solution to the problem of (info)social order should be reconsidered and mitigated. Incentives should be seen as useful means to enforce the norms, rather than as sufficient mechanisms for modelling and implementing them. Social order cannot primarily

³Simulation studies should be carried out to confirm this expectation.

rely upon incentives and sanctions, unless sanctions are always severe and certain so as to lower the utility of transgression compared to the utility of compliance.

4 Evidence from Natural Societies

Things work much better if norms are executed for their own sake, that is, if at least a share of the whole society accepts and complies with the norms for intrinsic motivations. But how is it possible that such type of norm exists at all? Or, better, how is it possible that autonomous agents have intrinsic reasons to comply with a norm? Does this type of agent really exist, or is it conceivable only in a morally ideal society?

A look at human societies shows some important phenomena. First, real (social or legal) norms are *not* primarily defined as incentive-based prescriptions, but rather as prescriptions which ought to be accepted for their own sake. Secondly, incentives have a lower effect on norm compliance than should be expected if a model of rational decision is accepted: natural agents take into account sanctions less than rational deciders are expected to do. Third, incentives may bear negative consequences on norm compliance. Let us examine each phenomenon with some detail.

4.1 Incentives and the Concept of a Norm

What are real norms? Which roles do incentives play in their definition and recognition? As said before, sanctions are neither necessary nor sufficient for norms. People can tell and accept a prescription as a norm, even if they do not know and are not informed about the respective sanctions. Indeed, this is quite often the case: agents take decisions in absence of "official" information about the entity and probability of sanctions. Moreover, agents may try to infer such information, but they will neither expect that such information be provided by the source of sanctions, nor are they allowed to exact it. Indeed, agents may take it into account "privately". To calculate the entity and probability of sanctions is (considered) an aggravation of crime, because to observe the norm ought to be a sufficient motivation. On the other hand, people may accept a command under threat even if they do not perceive it as a norm: agents may yield to intimidation even if they are perfectly aware that it is illegal (people may surrender to an armed criminal but denounce her as soon as possible).

4.2 Incentives in Norm Enforcement

Incentives do not enforce compliance as much as expected. Humans are rather heterogeneous with regard to normative decisions, although their decisions are often perceived as utilitarian. Statistics about crimes do not confirm the expectations allowed by the model of incentive-based decision. First, the average application of sanctions for certain crimes (burglary and robbery) is very moderate, and in some countries is close to 1%. Consequently, the utility of compliance should be close to 0, and compliance should collapse. Nonetheless, the majority of humans has never committed this type of crimes. Secondly, and moreover, the entity and probability of sanctions are not equivalent in decision-making: it is well-known that, with equal probability of application, compliance does not increase with the severity of sanctions.

Third, frequent transgressions certainly contribute to encourage transgression. But this is not only because the perception of frequent transgressions affects the computation of the utility of norm compliance. Other mental processes occur: either the formation of a normative belief is obstructed by the assumption that a disregarded norm is bad or unfair or inadequate and the legislator is weak and ineffective; or the normative goal is abandoned, because the control system is ineffective and unfair, and does not deserve obedience.

4.3 Bad Effects of Incentives with Human Agents

Good experimental evidence indicates that incentives may render a bad service to norms. Not only positive incentives have been found to reduce or inhibit intrinsic motivation (what is called overjustification; for a recent work, cf. Lepper, forthcoming): when agents receive a reward for an activity which they were intrinsically motivated to accomplish, their intrinsic motivation will decrease. What is worse, negative incentives may reduce the unpleasantness of transgression (Greene et al., 1976): the lower the sanction, and the more the agents which comply with the norm will be attracted to violation. Social psychologists explain these findings in terms of self-perception (Bem, 1972): the less my action (compliance) is justified by some external factor, the more I need to find an internal reason for it. I will therefore be led to develop some good feeling or positive attitude with regard to it. If I complied with a norm which is not enforced by severe sanctions, I must have had a good reason to do so. The norm must be an important one, or else, I may start to think that to comply with that norm is good for its own sake. I develop an intrinsic motivation towards that norm, or towards the norms in general.

But why are incentives applied, then? We all know that they are applied rather frequently. Rewards are used in education and learning with good results. The same is true for sanctions: parents keep punishing children when they do something wrong. Delinquents are imprisoned, although less often than desirable. Fraud and deception are castigated by the community. Social psychologists suggest some answers to this question. First, the smaller the incentives the better (Greene et al., 1976). Secondly, they work much better in improving the quality of performance than in motivating action (Tang and Hall, forthcoming), which is why rewards work better in physical and mental learning than in moral and social education. Thirdly, they work when no intrinsic motivation has developed yet (Tang and Hall, forthcoming). Once the desirable behaviour has appeared, incentives cease to be useful and may even demolish the good job done. Fourth, and moreover, they work at their best *if agents perceive them as side-benefits, or additional motivations*, rather than as unique or primary reasons for action (Hennessey and Zbikowski, 1993).

5 Final Remarks. Normative Agents Vs Rational Deciders: Which One would You Prefer to Deal with?

In this paper, the role of incentives in social order has been questioned, based on a notion of incentive as additional individual utility, provided by an external entity, to actions achieving global utility.

Two notions of norms have been defined and compared: (1) inputs which modify agents' decisions through incentives (sanctions) and (2) prescriptions to execute obligatory action for intrinsic motivations. Two types of agents which reason upon norms were also compared: (1) incentive based rational deciders, and (2)

normative agents which are prescribed to execute norms for intrinsic reasons. Expectations about the effects of these two types of agents on norm compliance have been formulated. With relatively inefficient application of sanctions (punishment), transgression propagates more easily and rapidly among incentive-based agents than among normative agents. Under suboptimal conditions of application of sanctions (uncertain punishment), normative agents are expected to exhibit an oscillatory or at least a graceful degradation of compliance, while incentive-based agents are expected to show a fast decline and even a collapse. Finally, the role of incentives in natural societies has been discussed. This role is shown to have lesser impact on natural social agents than expected by a model of rational decision. What is worse, incentives have been shown to produce even negative effects on several aspects of social learning and norm compliance.

However, which lesson can be drawn from observation of natural societies and extended to info societies? Is the observation of natural societies any relevant for software agent engineering? Our answer is, yes, if applications to agent-mediated interaction are considered. In this context, agent scientists and designers face an important pragmatic task: to design systems which can interact with one another or with humans in a useful, reliable, and trustworthy way *from the point of view of the human user*. The good question then is, with whom does a human agent prefer to interact with? More specifically, when it comes to execution of norms, which one is preferable, a rational decider or a normative agent. Here, it is necessary to distinguish the two main roles that a software agent is expected to play in agent mediated interaction: that of *user representative* and that of *partner*. In e-commerce, for example, a system represents a given user in finding good partners for bargain, giving assistance in negotiation, etc.. But interesting applications under development see software agents as partners of negotiation (cf.,). As to the role of representative, a rational decider which is benevolent to its user, has her preferences as its ultimate goals and applies strategies to maximise her utility is probably the best choice. But as to the second role, that of *partner*, it is not so clear what should be preferred. Ultimately, one prefers to deal with trustworthy agents. But are incentive-based rational decider trustworthy partners? Is it preferable to deal with a system which respects the norms only in the interests of its own user (and therefore to the extent that this is convenient to her), or with a system which, so to speak, takes norms seriously and respects them for their own sake? More specifically, which condition is more encouraging from the human agent point of view, an efficient and severe sanctioning system, or a society of trustworthy partners? If the system is not efficient enough, it is certainly preferable to have a chance to meet agents which respect the norms independent of sanctions. But even if the sanctioning system were efficient enough, wouldn't it be more appealing to have at least a chance to deal with "good guys", meet nice partners? Isn't it better from the human point of view to know that your partner behaved correctly not because it was more convenient to do so, but because of its good will? At this stage, these questions do not allow for a conclusive answer. But we think that we should be prepared to provide a pondering answer in the near future.

References

- Bem, D.J. 1972. Self-perception theory. In L. Berkovitz (ed) *Advances in Experimental Social Psychology*, New York, Academic Press.
- Boman, M. 1999. . In Conte, R., Falcone, R., e Sartor, G. s.i. of Artificial Intelligence and Law on "Agents and Norms" 7, 1999.
- Castelfranchi, C. e Conte, R. Limits of economic rationality for agents and MA systems. *Robotics and Autonomous Systems*, Special issue on Multi-Agent Rationality, Elsevier Editor
- Castelfranchi, C., Miceli, M., Cesta, A., 1992, *Dependence Relations among Autonomous Agents*. In Y. Demazeau, E. Werner (eds), *Decentralized AI - 3*, 215-31. Amsterdam: Elsevier.
- Cohen, P. R. & Levesque, H. J. 1990b. Persistence, Intention, and Commitment. In *Intentions in Communication*, ed. by P.R Cohen, J. Morgan & M.A. Pollack, 33-71. Cambridge, MA: The MIT Press.
- Cohen, Ph. & Levesque, H. (1990a). Intention is choice with commitment. *Artificial Intelligence*, 42(3), 213-261.
- Conte, R. e Castelfranchi, C. >From conventions to prescriptions. Towards an integrated view of norms. *Artificial Intelligence and Law* 7: 323-340, 1999.
- Conte, R. & Castelfranchi, C. 1995a. *Cognitive and social action*. London: UCL Press.
- Conte, R., Miceli, M., Castelfranchi, C., 1991, *Limits and Levels of Cooperation. Disentangling Various Types of Prosocial Interaction*. In Demazeau, J.P. Mueller (eds), *Decentralized AI-2*, Y. , 147-157. Amsterdam: Elsevier.
- Crabtree, B. What Chance Software Agents, *The Knowledge Engineering Review*, 13, 1998, 131-137.
- Greene, D., Sternberg, B., and Lepper, M.R., 1976. Overjustification in a Token Economy. *Journal of Personality and Social Psychology*, 57: 41-54.
- Halpern, J.Y., Moses, Y.O., 1985, *A Guide to the Modal Logics of Knowledge and belief. Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 480-91. Los Altos, CA: Kaufmann.
- Hennessey, B.A. and Zbikowski, S.M. 1993. Immunizing children against the negative effects of reward: A further examination of intrinsic motivation focus sessions, *Creativity Research Journal*, 6: 297-307.
- Jennings N. (1995). Commitment and Conventions: the foundation of coordination in multi-agent systems. *The Knowledge Engineering Review* , 8,.
- Jennings, N. (1992). On being responsible, in Y. Demazeau & E. Werner (eds) *Decentralized Artificial Intelligence 3*, Amsterdam: Elsevier Science Publisher, 93-102.
- Jennings, N. R. & Mandami, E. H. (1992). Using joint responsibility to coordinate collaborative problem solving in dynamic environments. In *Proceedings of the 10th National Conference on Artificial Intelligence*, San Mateo, CA: Kaufmann, 269-275.
- Jones A.J.I, Pšrn, I., 1991, *On the Logic of Deontic Conditionals*. In J.J.C. Meyer, R.J. Wieringa (eds), *First International Workshop on Deontic Logic in Computer Science*, 232-47.
- Kinny, D. & Georgeff, M. (1994). Commitment and effectiveness of situated agents. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, IJCAI-93, Sydney, 82-88.
- Lepper, M.R. (forthcoming) Theory by numbers? Some concerns about meta-analysis, *Applied Cognitive Psychology*
- Rao, A.S. A Report on Expert Assistants at the Autonomous Agents Conference. *The Knowledge Engineering Review*, 13, 1998, 175-1179.
- Shoham, Y. & Tennenholtz M. (1992). On the synthesis of useful social laws in artificial societies. *Proceedings of the 10th National Conference on Artificial Intelligence*, San Mateo, CA: Kaufmann, 276-282.
- Tang, S. and Hall, V.C. (forthcoming) The overjustification effect: A meta-analysis, *Applied Cognitive Psychology*