

A Model of Normative Multi-Agent Systems and Dynamic Relationships

Fabiola López y López¹ and Michael Luck²

¹ Facultad de Ciencias de la Computación
Benemérita Universidad Autónoma de Puebla, México
fabiola@cs.buap.mx

² Department of Electronics and Computer Science
University of Southampton, UK
mml@ecs.soton.ac.uk

Abstract. For agents, one of the advantages of being part of a society is the satisfaction of those goals whose success depends on the abilities of other agents. In turn, societies are controlled by norms and, consequently, agents must be able first to model the society in which they exist, and then to identify the different relationships, due to norms, in which they might be involved in order to act appropriately. Both of these could mean the difference between the success or failure of their goals. To this end, this paper focuses on the identification of the basic components of norm-based systems, and on representing and analysing the *dynamic* relationships between member agents which result from the processing of norms.

1 Introduction

For agents, one of the advantages of being part of a society is the satisfaction of those goals whose success depends on the abilities of other agents. Since societies are controlled by norms, agents must be able to model their society, and they must be able to identify the different relationships, due to norms, in which they might be involved, in order to act appropriately. We argue that the correct identification of such relationships may be the difference between the success or failure of an agent's goals. For example, to select a plan, agents take into account not only their own obligations and prohibitions but also those of other agents, and to adopt a norm, agents must recognise an issuer's authority. Much research has been undertaken on how to incorporate norms into agents and multi-agent systems in taking steps toward the computational implementation of societies, institutions and organisations. This research has ranged from fundamental work on the importance of norms in agent behaviour [5, 21], to proposing internal representations of norms [3, 4, 13], analysing the different types of norms [7, 19], considering their emergence in groups of agents [22], proposing logics for their formalisation [18, 23], and describing how agents manage norm adoption and compliance [2, 8, 14]. In the field of agents, research has primarily been focused at the level of multi-agent systems where norms represent the means to achieve coordination among their members. Current models of multi-agent systems regulated by norms assume not only that agents are able to comply with norms but also that they are able to obey the authorities of the system, mostly as an end in itself [6, 9, 11, 16]. This means that authorities, norms and the relationships that arise from norms are all fixed at the start, so that the authority

of agents can neither be objected to nor constrained (because the relationship is fixed). In this way, if an authority decides to apply punishments, an agent must accept those punishments even though it may consider that there is no motive to do so. This also constrains the flexibility of the system and loses the advantages that the autonomy of agents might provide. We argue that a model of multi-agent systems that considers the dynamism that arises from norms must be proposed. Moreover, since normative relationships that exist at one time may not last until another moment, agents must be provided with the means to identify changes in these relationships.

To date, existing models of system regulated by norms have not considered dynamic relationships between agents, yet these relationships can determine whether an agent will comply with norms. This paper address this limitation by proposing a model of multi-agent systems regulated by norms, and by describing the *dynamics* of norms and how, from the different stages in which norms are processed, different relations among agents can emerge. Besides the informal description, formal specifications of the main concepts and processes are given in order to avoid any ambiguity arising through the use of informal natural language. In particular, this avoids inconsistencies which might complicate the use and correct implementation of the theoretical framework provided.

In this document, first a general structure of a norm, and the basic characteristics of normative agents are discussed in Section 2. Then, in Section 3 a model of a multi-agent system regulated by norms is proposed. In the same section, some roles for agents that arise from norma, are identified. After that, the changes that occur in a system when norms are issued, complied with, or violated are described (Section 4). Finally, in Section 5, a set of normative relationships between agents is provided, before presenting our conclusions.

2 Norms and Normative Agents

In this section we describe the basic blocks from which to build up our theory of normative multi-agent systems. This conceptual infrastructure provides the basis for a broad theory, and underpins several aspects not included in this paper, but described elsewhere [14, 15]. As a means to develop a formal model of a normative agent without repeating earlier work, we adopt the SMART agent framework described in [10]. In addition, in what follows, we also adopt the Z specification language to construct the formal model, because Z schemas allow, among other things, an easy transition from specifications to programs, there are tools that allow type checking, and so on. A Z schema contains two parts: the declaration part which declares local variables, and the predicate part which expresses some properties of the values of these variables. Z is based on set-theory and first order logic, with details available in [20]. For brevity, however, we will not elaborate the use of Z further.

2.1 Agents

In the SMART agent framework, an *attribute* represents a perceivable feature of the agent's environment which, here, is represented as a predicate or its negation. Then, a particular *state* in the environment is described by a set of attributes, a *goal* represents situations that an agent wishes to bring about, *motivations* are desires or preferences that affect the outcome of the reasoning intended to satisfy an agent's goals, and *actions* are discrete events that change the state of the environment when performed. For the

purposes of this paper, we formally describe attributes, environmental states, goals and actions. Details of the remaining elements are not needed, so we simply consider them as given sets.

$[Predicate, Motivation]$

$Attribute ::= pos\langle Predicate \rangle \mid not\langle Predicate \rangle$

$EnvState == \mathbb{P}_1 Attribute \quad Goal == \mathbb{P}_1 Attribute$
 $Action == EnvState \rightarrow EnvState$

An autonomous agent is described by a non-empty set of attributes representing its permanent features, a set of goals that it wants to bring about, a set of capabilities that it is able to perform, and a non-empty set of motivations representing its preferences.

<i>AutonomousAgent</i>	
<i>capabilities</i> : $\mathbb{P} Action$;	<i>goals</i> : $\mathbb{P} Goal$
<i>motivations</i> : $\mathbb{P} Motivation$;	<i>beliefs</i> : $\mathbb{P}_1 Attribute$
<i>goals</i> $\neq \emptyset$;	<i>motivations</i> $\neq \emptyset$

2.2 Norms

An agent may have access to certain norms that are represented as data structures relating to social rules. These may be common to all agents (such as with a mutually understood social law) or only available to some. *Norms* are the mechanisms that a society uses in order to influence the behaviour of agents within it. Norms can be created from different sources, varying from built-in norms to simple agreements between agents, or more complex legal systems. They may persist during different periods of time; for example until an agent dies, as long as an agent remains in the society for which the norms were issued, or just for a short period of time until normative goals become satisfied. There are also different aspects that can be used to characterise them. First, norms are always prescribed to be complied with by a set of *addressee* agents in order to *benefit* another set of agents (possibly empty). They specify something that ought to be done, and consequently they include *normative goals* that must be satisfied by addressees. Sometimes, these normative goals must be directly intended, whereas at other times their role is to inhibit specific goals (as in the case of prohibitions). Second, norms are not always applicable, and their activation depends on the *context* in which agents are situated; there may be *exceptions* when agents are not obliged to comply with the norm. Finally, in some cases, norms suggest the existence of a set of *punishments* to be imposed when agents do not satisfy the normative goals, and a set of *rewards* to be received when agents do. Both punishment and rewards are also represented as sets of goals (which can be empty) that must be satisfied by someone else. Thus, the general structure of a norm can be formalised as follows. (Note that we specify normative goals as a set, to allow for the possibility of multiple goals in a norm, though we recognise that this will typically be a singleton set.)

<i>Norm</i> <i>addressees, beneficiaries</i> : $\mathbb{P} Agent$ <i>normativegoals, rewards, punishments</i> : $\mathbb{P} Goal$ <i>context, exceptions</i> : <i>EnvState</i>
<i>normativegoals</i> $\neq \emptyset$; <i>addressees</i> $\neq \emptyset$; <i>context</i> $\neq \emptyset$ <i>context</i> \cap <i>exceptions</i> = \emptyset ; <i>rewards</i> \cap <i>punishments</i> = \emptyset

In the formalisation above, some constraints are imposed on the elements of a norm to eliminate the possibility of having norms that prescribe nothing, norms that are addressed to no one, norms that do not specify the situations in which they must be complied with, or norms that have inconsistencies in describing either the states in which agents are immune or between the rewards and punishments associated with a norm.

Norms can be divided, without eliminating the possibility of having further categories, into four types: *obligations*, *prohibitions*, *social commitments* and *social codes*. Broadly, we can say that *obligations* and *prohibitions* are norms adopted once an agent becomes a member of a society, *social commitments* are norms derived from agreements or negotiations between two or more agents, and *social codes* are norms motivated by feelings such as love, pity, friendship, or social conformity. It is not the purpose of this paper to discuss the different categories of norms but these can be found elsewhere [13]. In the remainder of this paper we will use the term *norm* as an umbrella term to cover every type of norm, even those that do not include punishments; although they can be created in different ways and with different purposes, we argue that all of them share the same structure. An important consideration at this point is that we understand *prohibitions* as norms whose normative goals must be avoided by addressee agents.

2.3 Normative Agents

We define a normative agent as an autonomous agent whose behaviour is shaped by the obligations it must comply with, the prohibitions that limit the kinds of goals that it can pursue, the social commitments that have been created during its social life, and social codes that may not carry punishments, but whose fulfillment is a means of being identified as part of a community.

<i>NormativeAgent</i> <i>AutonomousAgent</i> <i>norms</i> : $\mathbb{P} Norm$

2.4 Permitted and Forbidden Actions

Sometimes, it is useful to observe norms not through the normative goals that ought to be achieved, but through the actions that can lead to the satisfaction of such goals. Then, we can consider actions that are either *permitted* or *forbidden* by a norm as follows. If there is a situation that makes a norm active, and the results of an action benefit the achievement of the associated normative goal, then such an action is *permitted* by the respective norm. For example, if the normative goal of a norm is to have taxes paid, then the action *paying taxes* is a permitted action if it changes an agent's situation of having taxes unpaid into a situation where taxes are paid. By analogy, we can define *forbidden*

actions as those actions leading to a situation that contradicts or hinders the normative goal. For example, the action *illegal parking* is a forbidden action by a norm whose normative goal is to avoid parking in front of a hospital entrance. In general, it is not trivial to observe how the results of an action might benefit or hinder the achievement of normative goals. For instance, if we spend all our money and then try to pay our taxes, it might be not obvious that spending money may hinder the normative goal of paying taxes. To avoid drilling down into the intricate details of this important but somewhat secondary concern in relation to the focus of this paper, the associations between situation states that might *benefit* or *hinder* goals are taken for granted and formalised as follows.

$$\left| \begin{array}{l} benefited_ : \mathbb{P}(EnvState \times Goal); \quad hindered_ : \mathbb{P}(EnvState \times Goal) \end{array} \right.$$

Now, we define two relations that hold among an action and a norm, which either permit or forbid the action, as follows.

$$\left| \begin{array}{l} permitted_ , forbidden_ : \mathbb{P}(Action \times Norm) \\ \hline \forall a : Action; n : Norm; \bullet \\ permitted(a, n) \Leftrightarrow (\exists g : n.normativegoals \bullet benefited(a \ n.context, g)) \wedge \\ forbidden(a, n) \Leftrightarrow (\exists g : n.normativegoals \bullet hindered(a \ n.context, g)) \end{array} \right.$$

In other words, if an action is applied in the context of a norm, and the results of this action benefit the normative goals, then the action is permitted, otherwise the action is forbidden. In this way, norms act as action-filtering norms.

3 Normative Multi-Agent Systems

Norms cannot be studied independently of the system for which they were created because they relate two or more members in a society, and it is the *social structure* that enforces norm compliance. Consequently, before describing how many processes due to norms are triggered, an analysis of the main components of a social system regulated by norms must be provided. A *normative multi-agent system* can be defined as a set of *normative* agents, which are controlled by a set of common *norms* ranging from obligations and social commitments to social codes. This *control* can be observed in three different aspects. First, member agents must recognise themselves as part of the society. Second, complete control cannot be exerted if sanctions or incentives are not applied when norms are either violated or complied with. Third, changes in the current normativity must be allowed as a way to solve unpredictable conflicts between agents and norms, or both. Each of these aspects is discussed in the subsections below.

3.1 Membership of Normative Societies

The performance of every structure of control relies on the capabilities of its members to recognise and follow its norms. However, since agents are autonomous, the fulfillment of norms can never be taken for granted. In fact, autonomous agents decide whether to comply with norms based on their own current goals and motivations [14]. It is also possible that not all the norms that one agent has adopted belong to just one system, because agents can be part of more than one society at the same time. In addition, due

to agent limitations, it is possible that not all the norms of the system are known by any agent. These characteristics can be formally expressed by saying that the set of norms adopted by any member is not necessarily a subset of the norms of the system, and also that the intersection of both sets of norms is not empty. Now, part of being a member of a society means that agents are subject to some of the norms in the system. In other words, the set of addressee agents of every norm must be included in the set of member agents, because it does not make any sense to have norms addressed to non-existent agents.

3.2 Interlocking Norms

The norms of a system are not isolated from each other; sometimes, compliance with them is a condition to trigger (or activate) other norms. That is, there are norms that prescribe how some agents must behave in situations in which other agents either comply with a norm or do not comply with it [17]. For example, when employees comply with their obligations in an office, paying their salary becomes an obligation of the employer; or when a plane cannot take-off, providing accommodation to passengers becomes a responsibility of the airline. Norms related in this way can make a complete chain of norms, because the newly activated norms can, in turn, activate new ones. Now, since triggering a norm depends on past compliance with another norm, we call these kinds of norms *interlocking norms*. The norm that gives rise to another norm is called the *primary* norm, whereas the norm activated as a result of either the fulfillment or violation of the first is called the *secondary* norm. In terms of the norm model mentioned earlier, the *context* is a state that must hold for a norm to be complied with. Since the fulfillment of a norm is assessed through its normative goals, the context of the secondary norm must include the satisfaction (or non-satisfaction) of all the primary norm's normative goals. Figure 1 illustrates the structure of both the primary and the secondary norms and how they are interlocked through the primary norm's normative goals and the secondary norm's context. To formalise this kind of norm, some definitions are needed. We say that a norm can be considered as *fulfilled* in a specific state of the system if its corresponding normative goals are a logical consequence of such a state.

$$\frac{\begin{array}{l} \text{logicalconsequence}_- : \mathbb{P}(\text{EnvState} \times \text{EnvState}) \\ \text{fulfilled}_- : \mathbb{P}(\text{Norm} \times \text{EnvState}) \end{array}}{\forall n : \text{Norm}; st : \text{EnvState} \bullet \\ \text{fulfilled}(n, st) \Leftrightarrow (\forall g : n.\text{normativegoals} \bullet \text{logicalconsequence}(st, g))}$$

Formally, a norm is interlocked with another norm *by non-compliance* if, in the context of the secondary norm, an instance of the primary norm can be considered as violated. This means that when any addressee of a norm does not fulfill the norm, the corresponding interlocking norm will be triggered. The formal specification of this is given below. There, n_1 represents the primary norm, whereas n_2 is the secondary norm.

$$\frac{\text{lockedbynoncompliance}_- : \mathbb{P}(\text{Norm} \times \text{Norm})}{\forall n_1, n_2 : \text{Norm} \bullet \text{lockedbynoncompliance}(n_1, n_2) \Leftrightarrow \\ (\exists n_1 : \text{Norm} \mid \bullet \neg \text{fulfilled}(n_1, n_2.\text{context}))}$$




Fig. 1. Interlocking Norm Structure

Similarly, a norm is interlocked with another norm *by compliance* if, in the context of the secondary norm, an instance of the primary norm can be considered as fulfilled. Thus, any addressee of the norm that fulfills it will trigger the interlocking norm. The specification of this is given as follows.

$$\left| \begin{array}{l} \textit{lockedbycompliance}_- : \mathbb{P}(\textit{Norm} \times \textit{Norm}) \\ \hline \forall n_1, n_2 : \textit{Norm} \bullet \textit{lockedbycompliance} (n_1, n_2) \Leftrightarrow \\ (\exists n_2 : \textit{Norm} \mid \textit{fulfilled} (n_1, n_2.\textit{context})) \end{array} \right.$$

3.3 Enforcement and Encouragement of Norm Compliance

Complete control cannot be exerted if, for each norm in the system, there is no other norm that prescribes how some agents have to react when the original norm is violated [17]. For example, if there is an obligation to pay accommodation fees for all students in a university, there must also be a norm stating what hall managers must do when a student does not pay. These norms are addressed to a specific group of agents responsible for punishing non-compliance. It is only through these norms that some agents are entitled to punish other agents. Chaos might emerge in a society if such responsibility is given either to no one or to anyone. Addressee agents of this kind of norm are frequently called the *defenders* of a norm.

To describe these norms, we observe that the violation of a norm can be detected by an agent when it realises that the associated normative goals were not satisfied. Once this event becomes identified by defenders, their duty is then to start a process in which rebellious agents can be punished. This suggests that these norms can be modelled as interlocking norms with the additional restriction that every punishment included in the violated norm must appear in the normative goals of the secondary norm. That is, defenders of norms must have the goal of punishing every offender of a norm. Figure 2 shows how both the structure of a norm and the norm which enforces it, are related. Formally, a relationship between a norm directed to control the behaviour of agents and a secondary norm can be defined as follows. A norm *enforces* another norm through punishments if they are interlocked by non-compliance, and the punishments associated with the unfulfilled norm are part of the normative goals of the enforcement norm. We call these kinds of norms *enforcement* norms.

$$\left| \begin{array}{l} \textit{enforce}_- : \mathbb{P}(\textit{Norm} \times \textit{Norm}) \\ \hline \forall n_1, n_2 : \textit{Norm} \bullet \textit{enforce} (n_1, n_2) \Leftrightarrow \\ \textit{lockedbynoncompliance}(n_1, n_2) \wedge n_2.\textit{punishments} \subseteq n_1.\textit{normativegoals} \end{array} \right.$$




Fig. 2. Enforcement Norm Structure

So far, we have described secondary norms in term of punishments because punishments are one of the more commonly used mechanisms to enforce compliance with norms. However, a similar analysis can be undertaken for secondary norms corresponding to the process of rewarding members fulfilling their duties. The relations between norms and norms to reward their compliance are shown in Figure 3. Formally, we say that a norm *encourages* compliance with another norm through rewards if they are locked by compliance, and the rewards associated with the fulfilled norm are part of the normative goals of the *encourage* norm. We call these kinds of norms *reward* norms.

$$\begin{array}{|l}
 \text{reward}_- : \mathbb{P}(\text{Norm} \times \text{Norm}) \\
 \hline
 \forall n_1, n_2 : \text{Norm} \bullet \text{reward}(n_1, n_2) \Leftrightarrow \\
 \text{lockedbynoncompliance}(n_1, n_2) \wedge n_2.\text{rewards} \subseteq n_1.\text{normativegoals}
 \end{array}$$




Fig. 3. Reward Norm Structure

Now, it is important to state that this way of representing enforcement norms can create an infinite chain of norms because we would also have to define norms to use when authorities or defenders do not comply with their obligations to either punish those agents breaking rules or reward those agents who fulfill their responsibilities [17]. To avoid this chain of norms, and by taking the risk of being considered as absolutist, in what follows we consider that no punishments are applied when an enforcement norm is not fulfilled. This means that neither authorities nor defenders can be judged (at least in this normative system) by dismissing their responsibilities. (A similar analysis can be undertaken for reward norms.) Nevertheless, if required, our model and formalisation for enforcing and encouraging norms can be used recursively as necessary. There is nothing in the definition of the model itself to prevent this.

3.4 Dynamic Normativity and Legislation

In general, norms are introduced into a society as a means to achieve social order. Some norms are intended to avoid conflicts between agents, others to allow the establishment

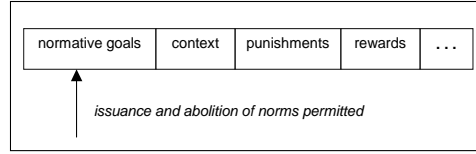


Fig. 4. Legislation Norm Structure

of commitments, and others still to unify the behaviour of agents as a means of social identification. However, neither all conflicts nor all commitments can be anticipated and controlled by norms. Consequently, in a *dynamic* multi-agent system there must exist the possibility of creating new norms, modifying existing norms, or even abolishing those that become obsolete. Now, although it is possible that many of the members of a society have capabilities to do this, these actions must be restricted to be carried out by a particular set of agents in a particular situation in order to avoid anyone imposing its norms, because some conflicts of interest might emerge. In other words, norms stating when actions to legislate are permitted must be also included [12]. These norms are called *legislation* norms, and they must specify that actions to issue and abolish norms are only permitted by a particular set of agents represented in its addressees (see Fig. 4). These constraints are specified below.

$$\begin{array}{|l}
 \hline
 \text{legislate}_- : \mathbb{P} \text{ Norm} \\
 \hline
 \forall n : \text{ Norm} \bullet \text{ legislate}(n) \Leftrightarrow (\exists \text{ issuingnorms}, \text{ abolishnorms} : \text{ Action} \bullet \\
 \text{permitted}(\text{issuingnorms}, n) \vee \text{permitted}(\text{abolishnorms}, n)) \\
 \hline
 \end{array}$$

3.5 Formal Model

All the elements discussed above are now incorporated into the formal representation of a *normative multi-agent system*, given below.

$$\begin{array}{|l}
 \hline
 \text{NormativeMAS} \\
 \hline
 \text{members} : \mathbb{P} \text{ NormativeAgent}; \quad \text{normsNMAS} : \mathbb{P} \text{ Norm} \\
 \text{enforcenorms} : \mathbb{P} \text{ Norm}; \quad \text{rewardnorms} : \mathbb{P} \text{ Norm} \\
 \text{legislationnorms} : \mathbb{P} \text{ Norm} \\
 \hline
 \forall ag : \text{ members} \bullet ag.\text{norms} \cap \text{normsNMAS} \neq \emptyset \\
 \forall rg : \text{ normsNMAS} \bullet rg.\text{addressees} \subseteq \text{members} \\
 \forall en : \text{ enforcenorms} \bullet (\exists n : \text{ normsNMAS} \bullet \text{enforce}(en, n)) \\
 \forall rn : \text{ rewardnorms} \bullet (\exists n : \text{ normsNMAS} \bullet \text{reward}(rn, n)) \\
 \forall ln : \text{ legislationnorms} \bullet \text{legislate}(ln) \\
 \hline
 \end{array}$$

That is, a normative multi-agent system comprises the following elements: a set of member agents able to reason about norms, a set of norms directed to regulate the behaviour of these agents (represented here by the variable *normsNMAS*), a set of norms directed to enforce and judge the latter set of norms (*enforcenorms*), the set of norms directed to encourage compliance with norms through rewards (*rewardnorms*), and the norms issued to allow the creation and abolition of norms (*legislationnorms*).

In the schema, the first predicate states that all members must have adopted some of the norms of the normative multi-agent system, and the second makes explicit that addressee agents of this set of norms must be members of the system. The last three predicates respectively describe the structure of enforcement, reward and legislation norms. Notice that whereas every enforcement norm must have a norm to enforce, not every norm may have a corresponding enforcement norm, which means that no one in that society is legally entitled to punish an agent that does not fulfill such a norm.

3.6 Normative Roles

Defining a normative multi-agent system in this way allows the identification of different roles for agents that depend on the kinds of norms agents are responsible for. Specifically, are agents entitled to create, modify, or abolish the set of norms of a society. No other members of the society are endowed with the power and authority to do so. These kinds of agents can, in turn, be either elected or decreed, and we call them *legislators*. An agent is an *defender* if it is directly responsible for either applying punishments or giving rewards. *Addressee* agents are directly responsible for the achievement of normative goals, and *beneficiaries* are agents whose goals can benefit when a normative goal becomes satisfied. Both addressees and beneficiaries can be directly identified from the model of norms. To identify defenders and legislators we need the following relations. The first and second relations state which agents are entitled to punish or reward a norm in a normative multi-agent system. The third relation specifies which agents can be considered the defenders of a particular norm. Finally, the fourth relation states who is a legislator.

$$\begin{aligned} \text{canpunish}_- &: \mathbb{P}(\text{NormativeAgent} \times \text{Norm} \times \text{NormativeMAS}) \\ \text{canreward}_- &: \mathbb{P}(\text{NormativeAgent} \times \text{Norm} \times \text{NormativeMAS}) \\ \text{isdefender}_- &: \mathbb{P}(\text{NormativeAgent} \times \text{Norm} \times \text{NormativeMAS}) \\ \text{islegislator}_- &: \mathbb{P}(\text{NormativeAgent} \times \text{NormativeMAS}) \end{aligned}$$

$$\begin{aligned} \forall ag : \text{NormativeAgent}; n : \text{Norm}; nmas : \text{NormativeMAS} \bullet \\ \text{canpunish}(ag, n, nmas) &\Leftrightarrow (n \in nmas.\text{normsNMAS} \wedge \\ &(\exists en : \text{Norm} \bullet (en \in nmas.\text{enforcenorms} \wedge \\ &ag \in en.\text{addressees} \wedge \text{enforce}(en, n)))) \wedge \\ \text{canreward}(ag, n, nmas) &\Leftrightarrow (n \in nmas.\text{normsNMAS} \wedge \\ &(\exists en : \text{Norm} \bullet (en \in nmas.\text{enforcenorms} \wedge \\ &ag \in en.\text{addressees} \wedge \text{reward}(en, n)))) \wedge \\ \text{isdefender}(ag, n, nmas) &\Leftrightarrow (\text{canpunish}(ag, n, nmas) \vee \\ &\text{canreward}(ag, n, nmas)) \wedge \\ \text{islegislator}(ag, nmas) &\Leftrightarrow (\exists ln : \text{Norm} \bullet \\ &ln \in nmas.\text{legislationnorms} \wedge ag \in ln.\text{addressees}) \end{aligned}$$

These *normative roles* for agents are not mutually exclusive. Agents are able to have more than one normative role at the same time, depending on the kind of norm being considered. For example, in a social commitment, the beneficiary agent can be a defender and consequently encourage the fulfillment of a norm, and either apply punishments or give rewards. In an office, the manager can be both a legislator and then

impose his own norms, and a defender entitled to punish his employees. The more complex a society, the more elaborate these normative roles become, and in some cases all legislators, authorities (judges), and police make a complex structure of control generally named *government*, with its own legal norms directed at controlling the rest of the society. Thus, being a defender is a relationship that holds between an agent and the enforcement norm that entitles it to defend the norm. Similarly, being a legislator means that there is a norm that entitles an agent to modify the current legislation by creating new norms and abolishing some of the norms already created. Considering defenders and legislators in this way allows us to represent the fact that all these elements cannot be taken independently of each other, but are complementary.

4 Dynamics of Norms

Norms are not a static concept. Once they are included in a system, they cause certain behaviour in each of the members. In Figure 5, the different processes through which a norm passes from its creation to its abolition can be observed. Arrows represent the transitions from one stage of norms to another. That is, first a legislator issues a norm. After that, the norm is spread by either indirect or direct communication. Then, adoption of norms takes place. Through this process, an agent expresses its willingness to comply with the norm as a way of being part of the society. Once a norm is adopted, it remains inactive, or in *latency*, until the applicability conditions are satisfied. In exception states, agents are not obliged to comply with norms, and consequently norms can be dismissed. However, in the majority of the cases, two different situations might occur after a norm becomes activated: a norm is either fulfilled or violated by addressee agents. After a norm is complied with, a reward can be offered. By contrast, if the norm is violated there are two possibilities: either punishments are applied or they are not. Finally, as time progresses, some norms become abolished or modified.

Considering the *dynamics* that result from norms is an important issue that deserves our attention, because interesting relations among agents can be identified in each (see Section 5). In turn, according to these relationships, different reactions of agents are expected. For instance, when a norm is activated, defenders are just entitled to require its fulfillment. However, if the norm is violated, defenders are entitled to apply punishments. Consequently, we argue that the correct identification of the different stages of a norm is key to modelling the normative behaviour of agents. In the following subsections, the transitions between these different stages are described and formalised from the perspective of an external observer.

4.1 Changing Legislation

Legislation of norms is a responsibility only attributed to legislator agents. Such a responsibility comprises at least three functions, namely issuance, abolition, and modification of norms. Unfortunately, due to their complexity, details of how such functions are carried out cannot be given here. Determining how, when and why a norm must be created requires a complex analysis of the current conditions of a system. However, we can still introduce two functions to identify the recently created norms (*newnorms*),

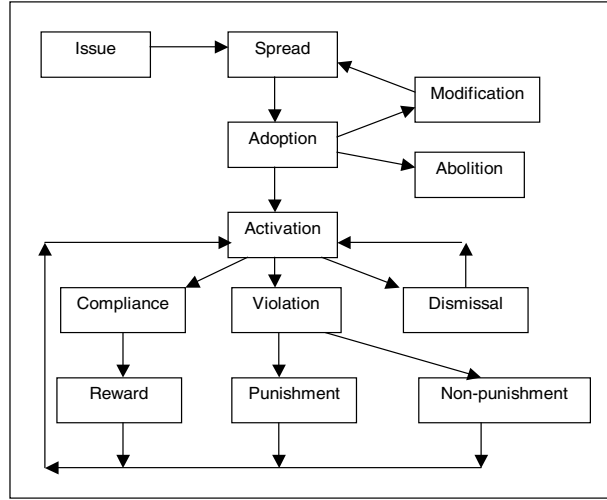


Fig. 5. Norm Dynamics

and the norms that must be abolished (*obsoletenorms*). These functions might be equivalent to asking legislators about the results of their assigned tasks. In turn, modification of norms can be seen as the abolition of a subset of norms together with the issuance of another subset of norms with the same name, so a specific function to do that is not included here. Now, after new norms are created and others are abolished, the spreading and updating of norms is needed. As a result of these changes at global level, the set of member agents must also change. That is, some of these norms become internally adopted or abolished by addressee agents. This is represented by the functions *spreadnorms* and *abolishnorms*, which can be seen as the processes through which agents are notified of the creation of new norms and the abolition of norms that become obsolete. The *NormLegislation* schema formalises the functions associated with the legislation of norms, in which the variable *nmas* represents the normative multi-agent system in which changes in legislation occur.

<i>NormLegislation</i>	
$nmas$	$NormativeMAS$;
$legislators$	$\mathbb{P} NormativeAgent$
$newnorms, obsoletenorms$	$\mathbb{P} NormativeAgent \rightarrow \mathbb{P} Norm$
$spreadnorms$	$(\mathbb{P} NormativeAgent \times \mathbb{P} Norm) \rightarrow \mathbb{P} NormativeAgent$
$abolishnorms$	$(\mathbb{P} NormativeAgent \times \mathbb{P} Norm) \rightarrow \mathbb{P} NormativeAgent$
$\forall ag : legislators \bullet islegislator (ag, nmas)$	
$dom newnorms$	$= \mathbb{P} legislators$;
$dom obsoletenorms$	$= \mathbb{P} legislators$

Now, the process that changes norms in both the system and its members can be represented as follows.

$\begin{aligned} & \text{ChangeLegislation} \\ & \Delta \text{NormLegislation} \\ & \text{mmas}'.\text{normsNMAS} = \text{mmas}.\text{normsNMAS} \setminus \\ & \quad \text{obsoletenorms legislators} \cup \text{newnorms legislators} \\ & \text{mmas}'.\text{members} = \text{spreadnorms} (\text{abolishnorms} (\text{mmas}.\text{members}, \\ & \quad \text{obsoletenorms legislators}), \text{newnorms legislators}) \end{aligned}$

The first predicate states that the set of norms, after a change in legislation, is composed of all the old norms except those recently abolished, combined with the recently created norms. The second predicate represents how all members are informed of legislation changes through a composition of functions. That is, first members are informed about norms that must be abolished because they are now considered obsolete, and then they receive information about the recently created norms.

4.2 Normative Multi-Agent System State

After norms are issued, spread, and then adopted, they enter in a cycle in which different agents intervene. To capture the different stages in which a norm is processed, we specify the *state* of a normative multi-agent system as follows.

$\begin{aligned} & \text{NMASState} \\ & \text{NormativeMAS} \\ & \text{currentsituation} : \text{EnvState} \\ & \text{formeractivenorms}, \text{activenorms}, \text{fulfillednorms} : \mathbb{P} \text{Norm} \\ & \text{unfulfillednorms}, \text{punishednorms}, \text{rewardednorms} : \mathbb{P} \text{Norm} \\ & \text{activenorms} \subseteq \text{normsNMAS} \end{aligned}$

At a particular instant of time, some norms become *activated*. This means that the conditions under which a norm must be fulfilled are satisfied. Moreover, other previously activated norms become either *fulfilled* or *unfulfilled*. Furthermore, some of the unfulfilled norms become *punished*, and some of the fulfilled ones become *rewarded*. Identifying these stages of norms is important because any change in them can cause reactions in other agents. For example, addressee agents acquire new responsibilities because of active norms, and they deserve to be rewarded or punished due to fulfilled or unfulfilled norms respectively. In addition, some agents might require compliance with active norms, or apply punishments to addressees of unfulfilled norms, etc. In the schema for the state of a normative multi-agent system, the *formeractivenorms* variable represents the norms that were activated previously, and the *currentsituation* represents the state of the general environment.

4.3 Assessing Compliance with Norms

Although not all norms change their stage at the same time, we take a particular point in the time to assess them all. Now, as mentioned before, the easy way to determine if a

norm has been fulfilled is by observing the current state of the system and then verifying if the associated normative goals are satisfied. This form of verifying compliance with norms can be used for any kind of norm, ranging from the norms of the normative system to the norms to enforce compliance with. These changes are represented in the schema below.

<p style="text-align: center;"><i>AssessNorm</i></p> <hr/> <p><i>NormativeMAS</i>; $\Delta NMASState$ <i>observedchanges</i> : <i>EnvState</i> \rightarrow <i>EnvState</i> <i>newactive, newfulfilled</i> : $\mathbb{P} Norm$ <i>newpunished, newrewarded</i> : $\mathbb{P} Norm$</p> <hr/> <p><i>currentsituation'</i> = <i>observedchanges</i> <i>currentsituation</i> let <i>newactive</i> == {<i>n</i> : <i>normsNMAS</i> <i>logicalconsequence</i> (<i>currentsituation'</i>, <i>n.context</i>)} • let <i>newfulfilled</i> == {<i>n</i> : <i>activenorms</i> <i>fulfilled</i> (<i>n</i>, <i>currentsituation'</i>)} • let <i>newpunished</i> == {<i>n</i> : <i>unfulfillednorms</i> ($\exists en$: <i>enforcenorms</i> • (<i>enforcepunish</i> (<i>en</i>, <i>n</i>) \wedge <i>fulfilled</i> (<i>en</i>, <i>currentsituation'</i>)))} • let <i>newrewarded</i> == {<i>n</i> : <i>fulfillednorms</i> ($\exists en$: <i>enforcenorms</i> • (<i>enforcereward</i> (<i>en</i>, <i>n</i>) \wedge <i>fulfilled</i> (<i>en</i>, <i>currentsituation'</i>)))} • (<i>formeractivenorms'</i> = <i>formeractivenorms</i> \cup <i>activenorms</i> \ <i>newactive</i> \wedge <i>activenorms'</i> = <i>newactive</i> \wedge <i>fulfillednorms'</i> = <i>fulfillednorms</i> \cup <i>newfulfilled</i> \wedge <i>unfulfillednorms'</i> = <i>unfulfillednorms</i> \cup (<i>activenorms</i> \ <i>newfulfilled</i>) \wedge <i>punishednorms'</i> = <i>punishednorms</i> \cup <i>newpunished</i> \wedge <i>rewardednorms'</i> = <i>rewardednorms</i> \cup <i>newrewarded</i>)</p>
--

In this schema, *observedchanges* is a function that reports the observed changes in the social environment, and can be used to update the sets of norms. First, the set of new *active* norms is calculated by analysing if the context to trigger a norm, is a *logical consequence* of the current situation of the system. After that, the set of active norms that were *fulfilled* by their corresponding addressee agents is calculated by verifying the satisfaction of the corresponding normative goals. Next, unfulfilled norms that were *punished* are found by verifying if the norm that enforces it has already been satisfied. Verifying if fulfilled norms were *rewarded* is done similarly. After all these steps, the states of norms are updated accordingly.

5 Normative Relationships

As stated earlier, norms at different states create different kinds of relationships among agents. We identify four sets. The first is created due to the authority of certain agents in the system. The next is created once norms become activated. Norms that have been complied with also generate relations among agents through offered rewards. Finally,

violated norms, and their associated punishments, cause agents to be related in a different way. These relationships are used by agents when reasoning about norms is needed, and a decision must be taken. Then, by using the proposed structure of the norm, the definition of a normative multi-agent system, and the different normative roles that agents might have be performing, we describe the set of relationships we are interested in. These relations are illustrated in Figure 6 in which rounded boxes represent the state of a norm, and hexagons symbolize the relationships created by them.

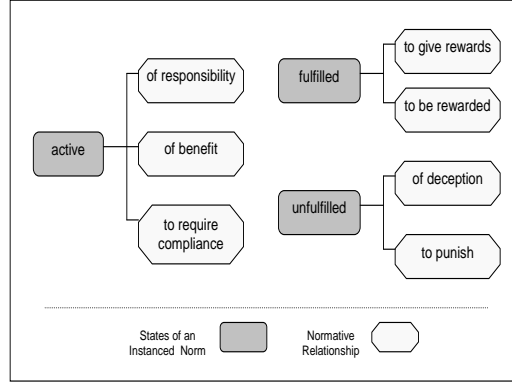


Fig. 6. Normative Relationships

5.1 Legislation Relations

As stated earlier, not all agents in a normative multi-agent system are entitled to legislate and, therefore, before a norm is adopted, agents must recognise the authority of the issuer, otherwise the validity of the norm could be questioned, and then rejected. Formally, we say that an agent is a *legal authority* for another agent if it is a legislator in the normative multi-agent system to which the second agent belongs.

$$\begin{array}{|l}
 \hline
 \text{legalauthority}_- : \mathbb{P}(\text{NormativeAgent} \times \text{NormativeAgent} \times \text{NormativeMAS}) \\
 \hline
 \forall ag_1, ag_2 : \text{NormativeAgent}; nmas : \text{NormativeMAS} \bullet \\
 \text{legalauthority}(ag_1, ag_2, nmas) \Leftrightarrow \\
 \text{islegislator}(ag_1, nmas) \wedge ag_2 \in nmas.\text{members}
 \end{array}$$

5.2 Active Norm Relations

Norms become *activated* when the current situations of an agent (or a group of agents) match the context in which a norm must be fulfilled. For example, if a driver wants to park its car in front of an entrance, the norm that forbids such an action is applied, otherwise agents do not need to be concerned with them. From this situation, four relations among agents can be inferred as follows.

It can be observed that some norms include exception states in which an agent is not obliged to respect those norm. An exception state could be treated as a state not included in the context of a norm, because in that case the norm would not be activated and, consequently, agents would not be obliged to comply. Although the results are similar, we prefer to make them explicit because it allows an agent to explain why it is not obliged to comply with that norm. This latter aspect can be useful if the norm is addressed to a set of agents, only some of which are excepted from their responsibilities. Formally, we say that an agent *can dismiss* a norm in a particular state of the system if that agent is an addressee of the norm, and the exception states of the norm are a logical consequence of the current state.

$$\frac{\text{candismiss}_- : \mathbb{P}(\text{NormativeAgent} \times \text{Norm} \times \text{EnvState})}{\forall ag_1, ag_2 : \text{NormativeAgent}; n : \text{Norm}; nmas : \text{NormativeMAS}; \\ st : \text{EnvState} \bullet \text{candismiss}(ag_1, n, st) \Leftrightarrow \\ (ag_1 \in n.\text{addressees} \wedge \text{logicalconsequence}(st, n.\text{exceptions}))}$$

Another important relationship that can be observed here is the relation between an addressee agent, a norm, and its defender. In this situation, it can be said that an agent is entitled to require compliance with norms either by threatening agents with future punishments, or by offering future rewards. Formally, it can be said that an agent *can require* another agent to fulfill a norm if it is a designated defender in the system, and the second agent is an addressee of the norm.

$$\frac{\text{canrequire}_- : \mathbb{P}(\text{NormativeAgent} \times \text{NormativeAgent} \times \text{Norm} \\ \times \text{NormativeMAS})}{\forall ag_1, ag_2 : \text{NormativeAgent}; n : \text{Norm}; nmas : \text{NormativeMAS}; \\ st : \text{EnvState} \bullet \\ \text{canrequire}(ag_1, ag_2, n, nmas) \Leftrightarrow (\text{isdefender}(ag_1, n, nmas) \wedge \\ ag_2 \in nmas.\text{members} \wedge ag_2 \in n.\text{addressees})}$$

Finally, there are two further important relationships between agents. The first is the responsibility that an addressee agent has as soon as a norm becomes activated. Note that although an agent has a responsibility to fulfill, it does not means that it is going to do so. The decision is only made by the agent itself. Formally, we say that an agent *has a responsibility* to another if there is a norm already addressed to the first agent, and the benefits may be enjoyed by the second.

$$\frac{\text{hasresponsibility}_- : \mathbb{P}(\text{NormativeAgent} \times \text{NormativeAgent} \times \text{Norm} \\ \times \text{NormativeMAS})}{\forall ag_1, ag_2 : \text{NormativeAgent}; n : \text{Norm}; nmas : \text{NormativeMAS}; \\ st : \text{EnvState} \bullet \\ \text{hasresponsibility}(ag_1, ag_2, n, nmas) \Leftrightarrow (n \in nmas.\text{normsNMAS} \wedge \\ ag_1 \in nmas.\text{members} \wedge ag_2 \in nmas.\text{members} \wedge \\ ag_1 \in n.\text{addressees} \wedge ag_2 \in n.\text{beneficiaries})}$$

The second relationship is its counterpart which relates to the expectations of a beneficiary agent to receive something from the responsibilities of others. Formally we

say that an agent *expects benefits* from the responsibility of another agent if the former is the beneficiary of a norm addressed to the second agent.

$$\begin{array}{|l}
\hline
\textit{expectsbenefit}_- : \mathbb{P}(\textit{NormativeAgent} \times \textit{NormativeAgent} \times \textit{Norm} \times \\
\textit{NormativeMAS}) \\
\hline
\forall ag_1, ag_2 : \textit{NormativeAgent}; n : \textit{Norm}; nmas : \textit{NormativeMAS}; \\
st : \textit{EnvState} \bullet \\
\textit{expectsbenefit}(ag_1, ag_2, n, nmas) \Leftrightarrow (n \in nmas.\textit{normsNMAS} \wedge \\
ag_1 \in nmas.\textit{members} \wedge ag_2 \in nmas.\textit{members} \wedge \\
ag_2 \in n.\textit{addressees} \wedge ag_1 \in n.\textit{beneficiaries})
\end{array}$$

5.3 Fulfilled Norm Relations

Once a norm is *fulfilled*, no further action is necessary except maybe by addressee agents claiming rewards from a defender. Then, two complementary relationships are identified as follows. First, we say that an agent has the responsibility of *rewarding* another agent if the first agent is a defender of the norm and the second is an agent who has fulfilled it. In addition, an agent has the right to be *rewarded* by a defender of a norm if the first agent has already complied with it.

$$\begin{array}{|l}
\hline
\textit{rewards}_- : \mathbb{P}(\textit{NormativeAgent} \times \textit{NormativeAgent} \times \textit{Norm} \times \textit{EnvState} \\
\times \textit{NormativeMAS}) \\
\textit{rewarded}_- : \mathbb{P}(\textit{NormativeAgent} \times \textit{NormativeAgent} \times \textit{Norm} \times \textit{EnvState} \\
\times \textit{NormativeMAS}) \\
\hline
\forall ag_1, ag_2 : \textit{NormativeAgent}; n : \textit{Norm}; nmas : \textit{NormativeMAS}; \\
st : \textit{EnvState} \bullet \\
\textit{rewards}(ag_1, ag_2, n, st, nmas) \Leftrightarrow (ag_2 \in n.\textit{addressees} \wedge \\
ag_2 \in nmas.\textit{members} \wedge \textit{canreward}(ag_1, n, nmas) \wedge \textit{fulfilled}(n, st)) \wedge \\
\textit{rewarded}(ag_1, ag_2, n, st, nmas) \Leftrightarrow (ag_1 \in n.\textit{addressees} \wedge \\
ag_1 \in nmas.\textit{members} \wedge \textit{canreward}(ag_2, n, nmas) \wedge \textit{fulfilled}(n, st))
\end{array}$$

5.4 Unfulfilled Norm Relations

By contrast, when a norm is *violated*, several events take place and other kinds of relationships hold. Obviously, addressees of an unfulfilled norm will do nothing, and would prefer that their failure remains hidden in order to avoid facing the consequences of their actions. However, a *deception* situation emerges in which the interests of third agents (the beneficiaries) might be badly affected by the irresponsibility of offenders. Agents in this situation could claim compensation. Formally, it can be said that an agent is *deceived* by another agent if a norm was violated by the second agent, and the benefits could have been enjoyed by the first.

$$\begin{array}{|l}
\hline
\textit{deceived}_- : \mathbb{P}(\textit{NormativeAgent} \times \textit{NormativeAgent} \times \textit{Norm} \times \textit{EnvState}) \\
\hline
\forall ag_1, ag_2 : \textit{NormativeAgent}; n : \textit{Norm}; nmas : \textit{NormativeMAS}; \\
st : \textit{EnvState} \bullet \textit{deceived}(ag_1, ag_2, n, st) \Leftrightarrow \\
(ag_1 \in n.\textit{beneficiaries} \wedge ag_2 \in n.\textit{addressees} \wedge \neg \textit{fulfilled}(n, st))
\end{array}$$

In addition, defenders also have a different relation with addressees. When a norm is activated, defenders are entitled only to enforce a norm, but when the norm is violated they have the responsibility to start a sequence of events leading to punish rebellious agents. Nevertheless, it could be possible that none of the defenders realises the occurrence of these events, and consequently the rebellious agent never becomes punished. Then it can be said that an agent must *punish* another agent if the first is a defender of the norm and the second is an agent who has violated it.

$$\begin{array}{|l}
 \hline
 \textit{punishes}_- : \mathbb{P}(\textit{NormativeAgent} \times \textit{NormativeAgent} \times \textit{Norm} \times \textit{EnvState} \\
 \times \textit{NormativeMAS}) \\
 \hline
 \forall ag_1, ag_2 : \textit{NormativeAgent}; n : \textit{Norm}; nmas : \textit{NormativeMAS}; \\
 st : \textit{EnvState} \bullet \textit{punishes} (ag_1, ag_2, n, st, nmas) \Leftrightarrow \\
 (\textit{isdefender} (ag_1, n, nmas) \wedge ag_2 \in nmas.\textit{members} \wedge \\
 ag_2 \in n.\textit{addressees} \wedge \neg \textit{fulfilled} (n, st))
 \end{array}$$

As we can observe, all these relationships are relativised both to a normative multi-agent system to which agents belong, and to the prevailing situation of agents. That is, no relationships due to norms can be created when agents do not belong to the same system, or when the conditions to activate a norm do not hold. We say that in a normative multi-agent system where social control has been defined through norms, some relations can be identified. That is, at a particular time, there are responsibilities that agents acquire through norms, situations in which addressee agents can be excepted from such responsibilities, enforcement mechanisms that might be applied to agents with duties, rewards that must be given to respectful agents, punishments that must be applied to norm offenders, and deceived agents expecting compensations. All these relationships change as soon as new norms become activated, fulfilled or violated.

6 Conclusion

So far in our work, the basic components of a system controlled by norms have been identified. We call these kinds of systems *normative multi-agent systems*, and we describe them as consisting of: a set of member agents whose compliance with norms is neither always enforced nor always expected, a set of norms directed at controlling the behaviour of all members, a set of legal norms to enforce compliance with regulations through punishment, a set of legal norms to reward agents who fulfill norms, and a set of norms to entitle some agents to change regulations. In general, current models of multi-agent systems regulated by norms include norms as obligatory actions that might otherwise be penalised [6, 9, 11, 16]. They typically do not make any distinction among norms. By contrast, our model divides these norms into three different classes which allow agents not only to identify the roles of other agents in a society but also to identify the limits of their responsibilities (given by the normative goals of the norms they have to comply with). In this way, agents' authority can be constrained.

Our model of norms facilitates the modelling of norms that must be complied with depending on compliance with other related norms (*contrary-to-duty* norms). By using interlocking norms, mechanisms to enforce compliance with norms are given through

enforcement and reward norms. In addition, the *dynamism* that occurs in a system due to norms has been analysed and, according to the different stages in the processing of norms, some *normative relationships* have been identified. The key concept here is the normative behaviour of agents caused not only by the existence of norms, but also by their issuance, fulfillment or violation, which in turn must be the result of the decisions of each of the members.

By studying the characteristics of normative multi-agent systems, we have set up the basis of a framework to represent different kinds of social systems regulated by norms that include elements that allow agents to reason about norms. In addition, the set of normative relationships identified in this paper might enable agents to take more effective decisions in situations where norms are involved. For example, agents who have benefited from a fulfilled norm might decide to reciprocate with the addressees of such a norm in their subsequent interactions. Normative relationships are also useful for identifying situations in which a subset of agents is legally empowered, and informing about the decision of when a new norm can be adopted or complied with. This is the focus of the next stage in our work. We also aim to extend our work on norm compliance [14] to introduce strategies in which agents are externally influenced to comply with a norm. Additionally, we must provide an analysis of those situations in which agents might adopt new norms. We believe that the normative roles that we have defined here can be used by agents to identify empowered agents, and therefore to identify from whom an order can be received.

Our analysis builds on much important work on norms. Ross, for example [17], describes some of the norms and relationships presented in this paper. In turn, Conte and Castelfranchi [3] have already mentioned some of the normative roles we present, and some of the processes involved in reasoning about norms. Jones and Sergot [12] also mention the characteristics of agents entitled to manage an institution. The closest work is by Balzer and Tuomela [1], who present the formalisation of an institution controlled by norms. However, no work considers the dynamics of norms nor the relationships that emerge from them, for use by agents to reason about norms.

Acknowledgments The first author is supported by the Faculty Enhancement Program (PROMEP) of the Mexican Ministry of Public Education(SEP) and the Benemérita Universidad Autónoma de Puebla (México). Authors also express their gratitude to the valuable comments provided by the anonymous reviewers of this paper.

References

1. W. Balzer and R. Tuomela. Social institutions, norms and practices. In C. Dellarocas and R. Conte, editors, *Social Order in Multi-Agent Systems*, pages 161–180. Kluwer Academic Publishers, 2001.
2. G. Boella and L. Lesmo. Deliberative normative agents. In C. Dellarocas and R. Conte, editors, *Social Order in Multi-Agent Systems*, pages 85–110. Kluwer Academic Publishers, 2001.
3. R. Conte and C. Castelfranchi. *Cognitive and Social Action*. UCL Press, 1995.
4. R. Conte and C. Castelfranchi. Norms as mental objects. From normative beliefs to normative goals. In C. Castelfranchi and J. P. Müller, editors, *From Reaction to Cognition (MAAMAW'93)*, LNAI 957, pages 186–196. Springer-Verlag, 1995.

5. R. Conte, R. Falcone, and G. Sartor. Agents and norms: How to fill the gap? *Artificial Intelligence and Law*, 7(1):1–15, 1999.
6. C. Dellarocas and M. Klein. Contractual agent societies: Negotiated shared context and social control in open multi-agent systems. In C. Dellarocas and R. Conte, editors, *Social Order in Multi-Agent Systems*, pages 113–133. Kluwer Academic Publishers, 2001.
7. F. Dignum. Autonomous agents with norms. *Artificial Intelligence and Law*, 7(1):69–79, 1999.
8. F. Dignum, D. Morley, E. Sonenberg, and L. Cavendon. Towards socially sophisticated BDI agents. In E. Durfee, editor, *Proceedings on the Fourth International Conference on Multi-Agent Systems (ICMAS'00)*, pages 111–118. IEEE Computer Society, 2000.
9. V. Dignum and F. Dignum. Modelling agent societies: Coordination frameworks and institutions. In P. Brazdil and A. Jorge, editors, *Progress in Artificial Intelligence Knowledge Extraction, Multi-agent Systems, Logic Programming, and Constraint Solving*, LNAI 2258, pages 191–204. Springer-Verlag, 2001.
10. M. d’Inverno and M. Luck. *Understanding Agent Systems*. Springer-Verlag, 2001.
11. M. Esteva, J. Padget, and C. Sierra. Formalizing a language for institutions and norms. In J. Meyer and M. Tambe, editors, *Intelligent Agents VIII (ATAL'01)*, LNAI 2333, pages 348–366. Springer-Verlag, 2001.
12. A. Jones and M. Sergot. A formal characterisation of institutionalised power. *Logic Journal of the IGPL*, 4(3):429–445, 1996.
13. F. López y López and M. Luck. Modelling norms for autonomous agents. In *Proceedings of the Fourth Mexican International Conference on Computer Science (ENC'03) (to appear)*. IEEE Computer Society Press, 2003.
14. F. López y López, M. Luck, and M. d’Inverno. Constraining autonomy through norms. In C. Castelfranchi and W. Johnson, editors, *Proceedings of The First International Joint Conference on Autonomous Agents and Multi Agent Systems AAMAS'02*, pages 674–681. ACM Press, 2002.
15. M. Luck and M. d’Inverno. A formal framework for agency and autonomy. In V. Lesser and L. Gasser, editors, *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95)*, pages 254–260. AAAI Press/MIT Press, 1995.
16. Y. Moses and M. Tennenholtz. Artificial social systems. Technical report CS90-12, Weizmann Institute, Israel, 1990.
17. A. Ross. *Directives and Norms*. Routledge and Kegan Paul Ltd., 1968.
18. M. Sergot. Normative positions. In P. Mc Namara and H. Prakken, editors, *Norms, Logics and Information Systems*, pages 289–308. IOS Press, 1999.
19. M. Singh. An ontology for commitments in multiagent systems: Toward a unification of normative concepts. *Artificial Intelligence and Law*, 7(1):97–113, 1999.
20. J. M. Spivey. *The Z Notation: A Reference Manual*. Prentice-Hall, 1992.
21. R. Tuomela. *The Importance of Us: A Philosophical Study of Basic Social Norms*. Stanford University Press, 1995.
22. A. Walker and M. Wooldridge. Understanding the emergence of conventions in multi-agent systems. In V. Lesser and L. Gasser, editors, *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95)*, pages 384–389. AAAI Press/MIT Press, 1995.
23. R. Wieringa, F. Dignum, J. Meyer, and R. Kuiper. A modal approach to intentions, commitments and obligations: Intention plus commitment yields obligation. In M. Brown and J. Carmo, editors, *Deontic Logic, Agency and Normative Systems*, pages 80–97. Springer-Verlag, 1996.