

# Modeling Control Mechanisms with Normative Multiagent Systems: The Case of the Renewables Obligation

Guido Boella<sup>1</sup>, Joris Hulstijn<sup>2</sup>, Yao-Hua Tan<sup>2</sup>, and Leendert van der Torre<sup>3</sup>

<sup>1</sup> Universitá di Torino

<sup>2</sup> Vrije Universiteit, Amsterdam

<sup>3</sup> CWI Amsterdam and Delft University of Technology

**Abstract.** This paper is about control mechanisms for virtual organizations. As a case study, we discuss the Renewables Obligation (RO), a control mechanism that was introduced in the United Kingdom to stimulate the production of renewable energy. We apply a conceptual model based on normative multiagent systems (NMAS). We propose to model both the participants and the normative system as autonomous agents, having beliefs and goals. Norms, which can be internalized by the agents as obligations, are translated into conditional beliefs and goals of the normative system, which concern both detection and sanctioning measures. We show that the model can handle both the regulative and the evidential aspects of the case.

## 1 Introduction

Recent developments in the areas of computer supported collaborative work, distributed knowledge management and ‘grid’ architectures for sharing resources and computational services have lead to an increasing interest in what has been termed a *virtual organization*: a collection of enterprises or organizations that need to coordinate across organizational boundaries [23, 27, 26]. A crucial aspect of virtual organizations is that participants are autonomous: they can join and leave, and although some participants are more powerful than others, there is no central authority that can completely impose its will. Joining a virtual organization may provide benefits which participants could not achieve by themselves. On the other hand, participants must trust other participants not to behave opportunistically. To create a sustainable network, participants must therefore observe some general norms about what constitutes accepted behavior. In the case of computational coordination infrastructures, such norms will have to be enforced automatically, by means of electronic data interchange protocols or web services [21]. That means that norms will have to be encoded explicitly, in the form of some specific *control mechanism* [17, 6]. Ideally, a control mechanism is evaluated before it is implemented. To facilitate the specification and evaluation of control mechanism, we need a conceptual model that allows us to reason about the expected behavior of participants when they are subjected to norms in a virtual organization.

We discuss a case study of an actual control mechanism, the Renewables Obligation (RO), which was introduced in the United Kingdom to stimulate the production of

energy from renewable sources [20]. The ruling involves an obligation for energy suppliers to produce evidence of having distributed a certain minimal amount of renewable energy. In case energy producers do not comply, a buy-out fee must be paid. This is an example of a so called regulative rule. To present evidence of the amount of renewable energy produced, suppliers use so called Renewables Obligation Certificates (ROCs). This shows the use of a so called constitutive rule. The ROCs can be traded freely. The whole process is administered by a special agency, called OfGEM, with several tasks. It must monitor the amount of ROCs presented by suppliers, to detect whether the obligation has not been breached. If so, it must collect the penalty. OfGEM also accredits those renewable energy producers that are allowed to issue ROCs. This shows the use of a delegation mechanism.

Because of the inherent autonomy of participants and the lack of central control, the conceptual models used to design and reason about virtual organizations are likely to be similar to the kinds of models used in agent-oriented software engineering [9, 28, 8]. Typically, such models speak of an organizational structure involving agents that fulfill tasks on the basis of the organizational roles they play. Because agents may have conflicting tasks, or even conflicting individual goals, the possible behavior of agents is restricted by social norms. Although the RO case is not about computational agents as such, the network of energy producers does form a virtual organization. Energy producers share a network on which they have to coordinate energy distribution. There is a global objective, to produce more renewable energy, but the government is not in a position to force all energy producers to make the initial investments involved. Therefore a control mechanism is introduced, that should achieve the global objective indirectly. So in both cases we are dealing with a form of mechanism design: a control mechanism is designed in such a way, that the resulting agent society or virtual organization will uphold a general norm, and thus be able to sustain itself.

In this paper we therefore present a style of conceptual modeling based on *normative multiagent systems* (NMASs). Normative multiagent systems are "sets of agents (human or artificial) whose interactions can fruitfully be regarded as norm-governed; the norms prescribe how the agents ideally should and should not behave. [...] Importantly, the norms allow for the possibility that actual behavior may at times deviate from the ideal, i.e., that violations of obligations, or of agents' rights, may occur" [13]. The idea is to model all parties involved by *autonomous agents*: agents that are free to determine their courses of action, based on their interests (goals) and on their current information about the world (beliefs). Norms govern the behavior of agents, based on the roles they occupy in a virtual organization. A crucial aspect of our approach is that in principle any decision making entity can be modeled as an autonomous agent. In particular, the normative system itself can be viewed as an autonomous agent too, with specific beliefs and goals [5, 4]. The model has also been applied to other case studies. In particular, it has been used to explain issues of trust in electronic commerce [3].

The remainder of the paper is structured as follows. In section 2 we give a brief description of the RO case. In section 3 we define our version of normative multiagent systems, and illustrate it by two examples. In section 4 we then show how the RO case can be modeled using normative multiagent systems. The paper ends with lessons learned from this modeling exercise.

## 2 Case Description: Renewables Obligation

An example of an actual control mechanism is provided by the Renewables Obligation case [20, 14]. In order to comply with international environmental agreements, such as the Kyoto protocol, governments implement different incentive schemes, to stimulate the generation and supply of ‘renewable energy’. Research and development of renewable energy generating technologies such as wind turbines, photovoltaic panels, hydro-electric power generators, and others, require high initial investments by energy producers. Therefore the production costs of renewable energy are higher than those of the energy produced by conventional means. The incentive scheme implemented in Great Britain, starting from April 2002, is based on the Renewables Obligation (RO) [20]. This is a legal obligation on all licensed electricity suppliers to produce evidence that they have supplied a percentage target of their electricity from renewable energy sources to customers in Great Britain. A special organization, the Office of Gas and Electricity Markets (OfGEM), a branch of the official regulator of the British Gas and Electricity markets, has been set up to manage the scheme.

Suppliers are required to produce evidence to OfGEM of their compliance with the RO. An important evidence token is the so-called Renewables Obligation Certificate (ROC). A ROC is received by the supplier when it buys electricity from an accredited renewable producer. A ROC can also be traded. Because electricity can be added, bought and sold as a commodity, it does not matter which supplier actually produces the ROCs, and which supplier buys them. This has led to the development of a market for ROCs. If the target number of ROCs increases, the total number of renewable energy in the system will increase, which was the objective of the scheme.

So suppliers can meet their Renewables Obligation in three ways. They can produce ROCs corresponding to the target level, expressed as a percentage of all electricity supplied to customers in Great Britain; they can use a so called buy-out clause which allows them to pay £30.51/Mwh for any shortfall below the target level, or they can use a combination of ROCs and buy-out fees. If a supplier fails to meet its obligation, either through ROCs or buy-out payments, the supplier is likely to be in breach of the Electricity Act, and may be liable to enforcement action. In practice this means that an additional fine has to be paid. Note that the OfGEM can manipulate the level of the buy-out fees in such a way that it becomes more beneficial for a supplier to buy the relatively expensive renewable energy with corresponding ROCs, rather than to violate the Renewable Obligation.

OfGEM has the authority to accredit electricity producers that are capable of generating electricity from renewable sources. Such renewable energy producers are subsequently allowed to issue ROCs. By modifying the accreditation criteria, the British government can make adjustments. For instance, there are plans to extend eligibility to producers that make use of biomass [16].

The efficiency of the Renewables Obligation has been evaluated and the result is largely positive: “The large majority of respondents considered that the Obligation has provided a positive stimulus for investment in renewable technologies, particularly lower cost technologies such as onshore wind and landfill gas. Most considered that the Obligation is largely working as anticipated and would deliver a significant expansion in renewable electricity generation over the coming years.” [16]

### 3 Normative Multiagent Systems

We present the basic idea of our version of normative multiagent systems here. For a more technical exposition, please refer to [5, 4].

The model is based on mental attitudes of agents, like beliefs (information) and goals (internal motivation), that drive decisions to plan and execute actions. External motivations, such as social norms or laws, can be represented in the form of obligations, once they are known and accepted by the agent. Later we show how obligations can again be reduced to goals of the individual agent and the normative system, using violation conditions and sanctions.

Mental attitudes of agents are represented in a logical representation language. The logic allows us to derive what other mental attitudes can be inferred from the specification of the agents, and what not. When the rules of the model are implemented, we can run simulations of the decision making of various agents.

#### 3.1 Mental Attitudes

In our logic, the mental attitudes are not represented as sets of sentences as is customary, but as sets of conditionals or production rules. This expresses the fact that mental attitudes are context dependent [12]. So each attitude *Bel*, *Goal*, etc., is represented by a set of rules of the form  $A \rightarrow B$ , where both  $A$  and  $B$  are formulas, composed of facts by means of logical operators  $\wedge$  (and),  $\vee$  (or),  $\neg$  (not) and the constant  $\top$  (always true). Here  $A$  represents the conditions under which the facts represented by  $B$  may be inferred by the agent. Moreover,  $B$  may contain special decision variables, also called actions, that will alter the state of the world. The values of decision variables are under the control of the agent. For simplicity, both facts and decision variables are represented by boolean variables, being either true or false. The decision making process of an agent is represented by a forward reasoning loop, which runs roughly as follows<sup>1</sup>.

The agent receives input from observations, represented as a set of facts  $S$ . Alternatively, the agent may start with a set of initial goals, represented by a set of decision variables  $S$ . Now the agent tries to match each rule  $A \rightarrow B$  against  $S$ . If  $A$  is contained in  $S$ , and the facts of  $B$  do not contradict a fact in  $S$ , the rule is applicable. However, there may be several rules applicable to  $S$ , from the same and from different mental attitudes, each with a different possible outcome. Using a priority ordering, the agent selects one rule – this is called conflict resolution – and applies it: the result  $B$  is now added to  $S$ . This process continues, until a stable set of facts is reached, to which no further additions can be made. Such a stable set, an extension, represents one of the possible outcomes of the decision making.

The decision making behavior of an agent crucially depends on the way the conflicts among the mental attitudes are resolved. Different priority orders may lead to different extensions, which represent sets of goals and hence lead to different behavior.

*Example 1 (Beer and Smoking.).* An agent has the following inclination. Whenever it finds itself at a party, it wants to drink beer. And whenever it is drinking beer, it needs a

---

<sup>1</sup> Technical details of the reasoning is expressed using input/output logics [18]. Their application to Normative Multiagent Systems is explained in [5, 4].

cigarette. The agent knows that smoking is bad, and therefore the agent has the intention not to smoke. What will happen?

Belief: `at_party`

Goal 1: `at_party` → `drink_beer`

Goal 2: `drink_beer` → `smoke_cigarette`

Goal 3:  $\neg \text{smoke\_cigarette}$

What will happen depends on the relative strength of the agent's urge to smoke and its resolution not to. These relative strengths can be expressed by a priority ordering on sets of rules. Rules in one set are considered equally important. In general, belief rules outrank goal rules; otherwise the agent would suffer from *wishful thinking* [7]. There are two possibilities. If the urge to smoke is too strong, and the agent has already drunk a beer, it will smoke despite its resolution not to. But if the resolution outranks the urge, the agent will refrain from smoking, or refrain both from drinking and smoking.

Priority: `Belief` > {Goal 1, Goal 2} > Goal 3

Outcome: {`at_party`, `drink_beer`, `smoke_cigarette`}

Priority: `Belief` > Goal 3 > {Goal 1, Goal 2}

Outcomes: {`at_party`, `drink_beer`}, {`at_party`}

What do we observe in this example? Before trying to achieve its goals, an agent will consider the previously derived goals in the extension along with their consequences. This process is called *goal generation* [7]. Goal generation precedes the planning process. Roughly, there are two kinds of goals. Achievement goals are satisfied once some state of affairs has been realized. An example is the goal to reach some location. Maintenance goals on the other hand, are only satisfied for as long as some state of affairs continues to hold. Consider for example the goal to maintain some safety standard.

### 3.2 Norms and Obligations

Now what about norms and obligations? Some people have observed that the stabilizing effect of goals for individual agents, is similar to the stabilizing effects of norms on a community of agents. Norms protect long term interests of the group against individual deviators. But how can we model norms in this setting?

The general idea is to use a reduction of obligations to goals of the normative system, where the normative system itself is seen as a separate agent. This may be summarized by the slogan "Your wish is my command": the wishes of the normative system count as commands for the individual agent, provided that the normative system has authority over the individual agents.

The reduction makes use of a so called *violation predicate* [1]. Although violation predicates have been known for a long time, making a reduction to goals rather than to modal or epistemic operators, does make a difference. It expresses that the normative agent makes a conscious decision to detect or sanction a violation. Violation detection is a specific kind of goal of the normative system. It may lead to the addition of a belief in case a violation is indeed detected. Violation detection is separated from sanctioning.

Sanctioning concerns the decision to sanction an agent in case a violation has been detected, and what sanction is most appropriate.

To make sure that obligations are translated correctly, a number of conditions must be observed [4].

**Definition 1 (Conditional Obligations).** Agent  $a$  is obliged to bring about  $x$  under sanction  $s$  in circumstances  $C$ , with respect to a normative system  $n$  in a given model  $M$ , written as  $M \models O_{a-n}(x, s|C)$ , iff:

1. Goal of  $n: C \rightarrow x$   
If agent  $n$  believes that  $C$ , then it has as a goal that  $x$  should be brought about.
2. Goal of  $n: C \wedge \neg x \rightarrow Viol(\neg x)$   
If agent  $n$  believes that  $C$  and  $\neg x$  is the case, then it has the goal  $Viol(\neg x)$ , i.e., to recognize  $\neg x$  as a violation by agent  $a$ .
3. Goal of  $n: \neg Viol(\neg x)$   
By itself, agent  $n$  has no goal for violations. This is to prevent arbitrary detection.
4. Goal of  $n: C \wedge Viol(\neg x) \rightarrow s$   
If agent  $n$  believes that  $C$  and detects  $Viol(\neg x)$ , then it has as a goal to apply sanction  $s$ .
5. Goal of  $n: \neg s$   
By itself, agent  $n$  does not have a goal to apply sanction  $s$ . This is to prevent arbitrary sanctions.
6. Goal of  $a: \neg s$   
Agent  $a$  has a goal not be sanctioned. Without this condition, the sanction would not deter agent  $a$  from violating the obligation.

*Example 2 (Common-Pool Resources).* Consider the following situation. There is a group of agents  $A = \{a_1, \dots, a_n\}$  that share access to a common pool resource  $R$ . Think for example of common fishing grounds, which may suffer from overfishing in spring [22]. Only if all fishermen have a modest spring catch, the fish stocks can be sustained. For this reason, a normative system is set up, with a norm that in spring, no fisherman is allowed to catch more than some predetermined quota<sup>2</sup>.

The model  $M_{pool}$  is set up as follows. We use boolean variables ‘spring’ to represent the fact that it is now spring, ‘overfish’, which stands for fishing more than the quota allow, and ‘fine’ for the penalty that must be paid. Variable  $a$  ranges over agents  $a_1, \dots, a_n$ . Now we have an obligation of agent  $a$  towards the community, not to overfish in spring against the penalty of paying a fine. Using definition 1, this obligation can be modeled as follows.

$M_{pool} \models O_{a-n}(\neg overfish, fine|spring)$ , if and only if:

1. Goal of  $n: \text{spring} \rightarrow \neg \text{overfish}$
2. Goal of  $n: \text{spring} \wedge \text{overfish} \rightarrow Viol(\text{overfish})$
3. Goal of  $n: \neg Viol(\text{overfish})$

<sup>2</sup> Such cases have been studied extensively in economics, for example using evolutionary game theory. See for example [25]. Our example is only meant for illustrative purposes.

4. Goal of  $n$ :  $\text{spring} \wedge \text{Viol}(\text{overfish}) \rightarrow \text{fine}$
5. Goal of  $n$ :  $\neg \text{fine}$
6. Goal of  $a$ :  $\neg \text{fine}$

In addition to these aspects of the norm, we must also model the usual beliefs and goals of an agent. In particular, we must model the fact that without the obligation, overfishing is profitable. This aspect is modeled here by a simple goal. Moreover, for the sake of simulation, let us suppose that it is now spring and that everybody believes this.

7. Goal of  $a$ : overfish
8. Belief of  $a, n$ : spring

### 3.3 Recursive Modeling

How are these rules applied? Suppose it is spring and agent  $a$  has a goal to overfish. Agent  $a$  also has a goal not to be fined. So we get an initial extension of the form  $\{\text{spring}, \text{overfish}, \neg \text{fine}\}$ . Because the consequences are to a large part controlled by other agents, agent  $a$  will try and predict the goals and actions of other agents. It applies all the rules it knows that other agents possess, including those of the normative agent  $n$ . A set of rules of another agent is called a *profile*. Crucially, profiles contain rules that are affected by the actions of the agent itself. That is why we call this process *recursive modeling*. However, for most applications no infinite recursion is needed; recursive models up to three levels of embedding are quite sufficient.

In the first step,  $a$  applies  $n$ 's rule 2 to its own initial extension, which will trigger a goal to detect a violation. On the other hand,  $n$  has goal 3 not to detect violations. Which of these will get priority depends on  $a$ 's profile of  $n$ . Let's assume that  $a$  believes that  $n$ 's conditional goal to report a violation outranks its general goal not to detect violations. That will produce an extension  $\{\text{spring}, \text{overfish}, \text{Viol}(\text{overfish}), \neg \text{fine}\}$ .

In the second step,  $a$  applies  $n$ 's goal 4 to sanction, weighed against  $n$ 's goal 3 not to sanction. If  $a$  believes that  $n$ 's goal not to sanction outweighs its goal to sanction, for example because of prohibitive costs of sanctioning,  $a$  will form the goal to overfish. But if  $a$  believes  $n$  will indeed punish detected overfishing, this would produce an extension  $\{\text{spring}, \text{overfish}, \text{Viol}(\text{overfish}), \text{fine}, \neg \text{fine}\}$ , which contains a contradiction. This conflict will have to be resolved by  $a$ 's own priority order. In case the goal not to be fined outweighs the goal to overfish – in practice: if the penalty is larger than the expected profits – the agent will not form a goal to overfish.

The example shows that an obligation will only work, if two conditions are met. First, the sanction must outweigh the benefits of overfishing. In our model we can express this by a priority constraint.

Goal of  $a$ :  $\neg \text{fine} > \text{Goal of } a: \text{overfish}$

Second, the perceived chances of being detected and fined, must be sufficiently high. In our model we can express this by adding priority constraints to the profiles of  $n$  used by  $a$  during recursive modeling.

Goal of  $n$ :  $(\text{spring} \wedge \text{Viol}(\text{overfish}) \rightarrow \text{fine}) > \text{Goal of } n: \neg \text{fine}$

Thus a lot depends on the enforcement mechanism. Because the normative system, embodied by the village council for example, has no physical power, violation detection and sanctioning must be delegated to specific agents, such as a police force. There are examples of self-organizing communities in which these tasks are performed by ordinary community members, who may even behave altruistically, in the sense that they are not compensated for their detection and punishing efforts [22].

### 3.4 Constitutive Norms

So how do we model the fish quota? Suppose there is an accepted belief among the fishing community that a catch of more than two tonnes a week ‘counts as’ overfishing. The idea is to use *constitutive norms* [24] to model such general beliefs of the community. A constitutive norm applies only under certain circumstances and is intimately linked to an institution. This institution can be a (legal) person or an abstract entity such as a community of users. Whether a rule applies depends on the jurisdiction of the institution. In case we have a normative system, we can re-apply index  $n$  to stand for the institution. Thus constitutive rules are of the form “ $x$  counts as  $y$  under circumstances  $C$  in institution  $n$ ”.

In the example, the quota for overfishing in spring can be expressed as a belief rule of the institution  $n$ , and all agents  $a$  that fall under its jurisdiction.

Belief of  $a, n$ : spring  $\wedge$  (weekly catch  $> 2$  tonnes)  $\rightarrow$  overfish

Constitutive norms can create new institutional facts. The prototypical example is a declarative speech act [2]. For example, the utterance “I name this ship Johanna”, uttered by a lady at an appropriate christening ceremony, will create the institutional fact of the ship bearing this name. So the right kind of utterance uttered by the right person under the right circumstances ‘counts as’ as the creation of an institutional fact.

### 3.5 Value Objects

One way to express that objects represent a value in a community of agents, is as a goal to obtain such objects, attributed to an abstract agent that represents the shared conventions in the community. Just like beliefs can capture ‘counts as’ rules, goals can capture value, up to a point. Note that the attribution of value need not be shared. There are conventions about what is considered valuable in a community, but the value that an individual agent attaches to an object is not modeled. In our example, the fact that overfishing is generally considered profitable, clause 7, is an example of a value expression.

Goal of  $a$ : overfish

Although we do not explicitly model value, as for example [11], the effect of the relative priority of goals in the recursive simulation and violation games that agents play with each other, is similar to that of value.

## 4 Case Analysis

The case hinges on several aspects. There is a penalty mechanism that creates an incentive for energy suppliers to acquire evidence of having distributed a minimal level of renewable energy to British customers. The supplier has a choice whether to buy the obligatory amount of “green” energy, to buy ROCs from other suppliers, or else to pay a buy-out fee that corresponds to the amount of ROCs missing. The penalty aspect can be modeled by violation detection and sanctioning goals, similar to the obligation in section 3.

We consider a set of suppliers  $\{s_1, \dots, s_n\}$  and the normative system, appropriately called OfGEM. The main variables of the model  $M_{ROC}$  are ‘meet\_target’, which represents the fact that a supplier has collected enough ROCs, either by selling renewable energy to customers, or by buying them of other suppliers, ‘buy-out’ which represents that a supplier must pay a buy-out fee, and ‘enforcement\_order’ which represents the sanctioning action of OfGEM in case the Renewables Obligation has been breached.

The normative status of the buy-out fee is interesting. There are two options. On the one hand, paying a buy-out fee is a legitimate way of conforming to the Renewables Obligation. A supplier who pays a buy-out fee is not in breach of the law. This would suggest that there is one obligation, that can be fulfilled in two different ways<sup>3</sup>:

$$O_{s-\text{OfGEM}}(\text{meet\_target} \vee \text{buy-out}, \text{enforcement\_order} | \top)$$

On the other hand, paying a buy-out fee is not the preferred option. The buyer can choose to pay a buy-out fee, but the penalty level is set by OfGEM in such a way, that paying the fee is always more expensive than the appropriate number of ROCs would have been. Hence, we believe that the buy-out fee is better classified as a sanction for not meeting the target. That suggests that there are in fact two obligations. When the supplier is in breach of this second obligation, to buy the buy-out fee, OfGEM can call for a further sanction: an enforcement order. This will force the supplier to pay, against a penalty of further legal sanctions.

$$O_{s-\text{OfGEM}}(\text{meet\_target}, \text{buy-out} | \top)$$

$$O_{s-\text{OfGEM}}(\neg \text{buy-out}, \text{enforcement\_order} | \neg \text{meet\_target})$$

The first of these clauses is worked out in more detail.

$M_{ROC} \models O_{s-\text{OfGEM}}(\text{meet\_target}, \text{buy-out} | \top)$ , if and only if:

1. Goal of OfGEM:  $\text{meet\_target}$
2. Goal of OfGEM:  $\neg \text{meet\_target} \rightarrow \text{Viol}(\neg \text{meet\_target})$
3. Goal of OfGEM:  $\neg \text{Viol}(\neg \text{meet\_target}, a)$
4. Goal of OfGEM:  $\neg \text{meet\_target} \wedge \text{Viol}(\neg \text{meet\_target}) \rightarrow \text{buy\_out}$
5. Goal of OfGEM:  $\neg \text{buy\_out}$
6. Goal of  $s$ :  $\neg \text{buy\_out}$

Again, we have to state that without the RO scheme, not meeting the target for renewable energy would be profitable, because of the initial investments required.

7. Goal of  $s$ :  $\neg \text{meet\_target}$

---

<sup>3</sup> Note that the representation does not account for the fact that a partial shortfall in ROCs may also be supplemented with a corresponding partial buy-out fee.

How can we predict the behavior of an average supplier? What matters is the relative priority of goals 6 and 7. This priority is determined by economic considerations of the supplier agents. Since renewable energy is much more expensive to purchase for the supplier than non-renewable energy, the supplier will only consider purchasing renewable energy, when the level of the buy-out fee is set high enough. So, if the buy-out fee is higher than the extra cost for purchasing renewable energy, then the supplier will prefer 6 to 7. However, if the buy-out fee is less than the extra cost for purchasing renewable energy, then the supplier will prefer 7 to 6. A detailed profitability analysis to determine this preference ordering is beyond the scope of this paper. Detailed tools exist for such purposes. An example is the  $e^3$ -value tool [11]. We will investigate the links with profitability in further research.

*Regulative Aspect.* The Renewable Obligation will only work, in case some further conditions are met. First, the probability of being detected must be perceived to be high by the energy suppliers. Since the ROC scheme makes use of evidential documents and the burden of proof lies with the suppliers, this condition is taken care of. Second, the perceived probability of the OfGEM actually collecting buy-out fees, must be high enough. Currently, buy-out fees flow back into the system. They are used to finance the OfGEM itself and for other renewable energy stimulation. Thus, it is in the interest of OfGEM to actually collect buy-out fees. Moreover, if a supplier fails to pay the buy-out fees, a more severe sanction is invoked. In the ROC case this is called an enforcement order. The obligation  $O_{s-\text{OfGEM}}(\neg\text{buy-out}, \text{enforcement\_order} \mid \neg\text{meet\_target})$  is modeled analogously to clause 1-7 above.

In the two years that the scheme has been up and running, OfGEM has managed to collect most buy-out fees that were due. Recently, two electricity suppliers, Atlantic Electric and Gas Ltd and Maverick Energy Ltd were fined for their likely breach of the Renewables Obligation. Because both companies went into administrative receivership (bankruptcy), OfGEM decided not to issue further enforcement orders [19]. This shows that OfGEM is willing and able to enforce the Renewables Obligation.

*Evidential Aspect.* The ROC documents provide *evidence* of a certain amount of energy having been produced from renewable sources. This evidence is needed for different control purposes, such as to verify that the energy is indeed from an accredited renewable generator, that suppliers meet their Renewables Obligation, and that ROCs keep their value when traded.

This use of ROC evidence documents can be modeled by specific constitutive rules of the normative system, in this case embodied by the OfGEM. So OfGEM guarantees that within the wider community of energy trade, these ROCs ‘count as’ evidence that renewable energy was produced, and can be traded as such. In our formalization, this comes out as a simple belief rule of the normative system, and of all relevant suppliers. Similar reasoning holds for the establishment of the target number of ROCs that a supplier must meet. Note that this target number is relative to the size of the supplier.

Belief of  $s, \text{OfGEM}$ : 1 ROC  $\rightarrow$  1 Mwh renewable electricity

Belief of  $s, \text{OfGEM}$ :  $x \text{ ROC} \wedge (x > \text{target}) \rightarrow \text{meet\_target}$

OfGEM delegates authority to accredited renewable energy producers. Only accredited producers are allowed to issue ROCs. The effect of this delegation relation also shows

as part of the ‘counts as’-rules. Thus, only documents by accredited energy producers will count as true ROCs.

As long as the OfGEM continues to guarantee the validity of the evidence documents, these documents themselves can be traded as valuables. The identity of the holder of a ROC does not matter. For the overall objective of the scheme – to increase the amount of renewable energy produced – it does not matter whether a ROC or real renewable energy is traded, since every ROC stands for a certain quantity of renewable energy having been generated at some point. Because of the existence of a trading market, a supplier may specialize. Some, most notably in Scotland, are better at producing renewable energy; others are better at distributing it.

*Other Aspects.* Regarding some of the other aspects of the case, like the total flow of value between parties, and the fact that no ROCs must disappear from the system, our approach can still benefit from insights of other formalisms. In particular, modeling the decision of a supplier whether to produce ROCs, buy them or pay the buy-out fee, and in what relative proportions, would require the use of more detailed profitability analysis tools. We refer to [14, 15] for an analysis of the case that makes use of the  $e^3$ -value method [11], which does provide such quantitative tools. In an extension to  $e^3$ -value, called  $e^3$ -value<sup>+</sup>, Gordijn and Tan [10] also incorporate aspects of trust and control into the  $e^3$ -value approach. A promising direction for further research, is to investigate further combinations of value-based and normative approaches to virtual organizations.

## 5 Conclusions

In this paper we analyze how normative systems help create a sustainable network or virtual organization. Participants can make a choice to participate in a network or not. When there is no normative structure, participants may suffer from opportunistic behavior of others. The example of the common pool resource illustrated how this can be remedied by a normative system. On the other hand, when too severe a sanction is imposed, agents may not survive, or leave the network. Thus participants are autonomous in their decision to remain in the network. For this reason, it makes sense to apply agent-based modeling techniques. In particular, we apply Normative Multiagent Systems (NMAS) [5, 4].

We model a normative system as an autonomous agent, with explicit beliefs and goals. Norms are reduced to beliefs and goals of the subjects and the normative system, according to the slogan “My wish is your command”. Having an explicit agent to represent the normative system helps in particular to deal with delegation relations, and the way norm enforcement mechanisms such as detection and sanctions are implemented. It also makes it possible to explicitly capture the objectives of a norm, namely as the goals of the normative system.

The normative multiagent systems approach is validated with a case study of the Renewables Obligation in the United Kingdom. This control mechanism was introduced to stimulate the production of energy from renewable sources. What lessons did we learn from this modeling effort?

With respect to the RO case, we learned that the status of the buy-out fee is unclear. It can either be regarded as a legitimate way of complying with the Renewables

Obligation, or as a kind of sanction. The difference comes out clearly in the model. We choose to model the Renewables Obligation as a combination of two obligations: one to meet the target, and a second one that is conditional on not meeting the target, to pay the buy-out fee. This pattern of cascading obligations, in which sanctions are themselves modeled as regulative rules, is quite common in legal texts. Another lesson is that evidential documents like ROCs are difficult to capture. They have an ambivalent meaning because they are both evidential documents and value objects that can be traded themselves.

With respect to the Normative Multiagent Systems approach, we learned that there are in fact normative systems that need the whole range of concepts to be modeled accurately: beliefs, goals, obligations, violation conditions, sanctions and ‘count as’-rules. The case study shows that a complex example can be modeled consistently (internal validity), and that normative multiagent systems can analyze relevant aspects of existing control mechanisms in a virtual organization (external validity).

The model does have important limitations. First, the underlying representation of facts and events is too simple. For example, temporal or organizational relationships can not be conveniently captured. Second, the approach is only qualitative, using relative comparisons to model priorities. Many applications need a quantitative profitability analysis to set the right penalty levels, for the incentive scheme to work. But only after a qualitative analysis has shown the viability of a control mechanism, does it pay to make a detailed quantitative model.

**Acknowledgements.** Thanks to Vera Kartseva for providing details of the RO case.

## References

1. A. Anderson. A reduction of deontic logic to alethic modal logic. *Mind*, 67:100–103, 1958.
2. J.L. Austin. *How to do things with words*. Harvard U.P., Cambridge MA, 1962.
3. G. Boella, J. Hulstijn, Y-H. Tan, and L. van der Torre. Transaction trust in normative multiagent systems. In C. Castelfranchi, S. Barber, J. Sabater, and Munindar Singh, editors, *Proceedings of AAMAS Workshop on Trust in Agent Societies, Utrecht, The Nederlands, 2005 (Trust'05)*, 2005.
4. G. Boella and L. van der Torre. A game theoretic approach to contracts in multiagent systems. *IEEE Transactions on Systems, Man and Cybernetics - Part C*, to appear 2006. Special issue on Game-theoretic Analysis and Stochastic Simulation of Negotiation Agents.
5. G. Boella and L. van der Torre. Security policies for sharing knowledge in virtual communities. *IEEE Transactions on Systems, Man and Cybernetics - Part A*, to appear 2006. Special issue on Secure Knowledge Management.
6. Roger W.H. Bons, Frank Dignum, Ronald M. Lee, and Yao-Hua Tan. A formal analysis of auditing principles for electronic trade procedures. *International Journal of Electronic Commerce*, 5(1):57–82, 2000.
7. J. Broersen, M. Dastani, J. Hulstijn, and L. Van der Torre. Goal generation in the BOID architecture. *Cognitive Science Quarterly*, 2(3-4):431–450, 2002.
8. J. Castro, M. Kolp, and J. Mylopoulos. Towards requirements-driven information systems engineering: the TROPOS project. *Information Systems*, 27:365–389, 2002.
9. Jacques Ferber. *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*. Addison-Wesley, 1999.

10. J. Gordijn and Y.-H. Tan. A design methodology for trust and value exchanges in business models. *International Journal of Electronic Commerce*, 9(3):31, 2005.
11. Jaap Gordijn and J.M. Akkermans. Value-based requirements engineering: Exploring innovative e-commerce ideas. *Requirements Engineering*, 8(2):114–134, 2003.
12. B. Hansson. An analysis of some deontic logics. *Nóus*, 3:373–398, 1969.
13. A. Jones and J. Carmo. Deontic logic and contrary-to-duties. In D. Gabbay, editor, *Handbook of Philosophical Logic*, pages 203–279. Kluwer, 2002.
14. Vera Kartseva, Jaap Gordijn, and Yao-Hua Tan. Value-based business modelling for network organizations: Lessons learned from the electricity sector. In Jrg Becker and Freimut Bodendorf, editors, *Proceedings of the 12th European Conference on Information Systems (ECIS'04), Turku, CD-ROM*, 2004.
15. Vera Kartseva, Jaap Gordijn, and Yao-Hua Tan. Towards a modelling tool for designing control mechanisms in network organisations. *International Journal of Electronic Commerce*, forthcoming.
16. Alex King. 2005/6 review of the renewables obligation statutory consultation. Technical Report september 2005, Department of Trade and Industry, United Kingdom, 2005.
17. R.M. Lee and R. W. H. Bons. Soft-coded trade procedures for open-edi. *International Journal of Electronic Commerce*, 1(1):27–49, 1996.
18. D. Makinson and L. van der Torre. Input/output logics. *Journal of Philosophical Logic*, 29:383–408, 2000.
19. OFGEM. Renewables Obligation - Atlantic Electric and Gas Limited and Maverick Energy Limited. Technical Report R/65, Office of Gas and Electricity Markets, United Kingdom, 2004. [www.ofgas.gov.uk](http://www.ofgas.gov.uk).
20. OFGEM. The renewables obligation: Annual report 2003-2004. Technical Report HC 786, Office of Gas and Electricity Markets, United Kingdom, 2004. [www.ofgas.gov.uk](http://www.ofgas.gov.uk).
21. Alexander Osterwalder and Yves Pigneur. An eBusiness model ontology for modeling eBusiness. In Jozé Gricar, editor, *Proceedings of the 15th Bled Electronic Commerce Conference – eReality: Constructing the eEconomy (Bled'02)*, pages 75–91, 2002.
22. Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, New York, 1990.
23. M. B. Sarkar and C. Butler, B. andSteinfield. Cybermediaries in electronic marketspace: Toward theory building. *Journal of Business Research*, 41:215–221, 1998.
24. J.R. Searle. *The Construction of Social Reality*. The Free Press, New York, 1995.
25. R. Sethi and E. Somanathan. The evolution of social norms in common property resource use. *American Economic Review*, 86(4):766–789, 1996.
26. Don Tapscott, Alex Lowy, and David Ticoll. *Harnessing the Power of Business Webs*. Harvard Business School Press, Boston, MA, 2000.
27. P. Timmers. Business models for electronic markets. *Electronic Markets*, 8(2):3 – 8, 1998.
28. Michael J. Wooldridge, Nicholas R. Jennings, and David Kinny. The Gaia methodology for agent-oriented analysis and design. *Autonomous Agents and Multi-Agent Systems*, 3(3):285–312, 2000.