Noname manuscript No.
(will be inserted by the editor)

Transition Systems for Designing and Reasoning about Norms

T.J.M. Bench-Capon

the date of receipt and acceptance should be inserted later

Abstract The design and analysis of norms is a somewhat neglected topic in AI and Law, but this is not so in other areas of Computer Science. In recent years powerful techniques to model and analyse norms have been developed in the Multi-Agent Systems community, driven both by the practical need to regulate electronic institutions and open agent systems, and by a theoretical interest in mechanism design and normative systems. Agent based techniques often rely heavily on enforcing norms using the software to prevent violation, but I will also discuss the use of sanctions and rewards, and the conditions under which compliance by autonomous agents (including humans) can be expected or encouraged without sanctions or rewards. In the course of the paper a suggested framework for the exploration of these issues is developed[1].

1 Introduction

We see norms expressed in various ways every day. *Smoking is forbidden* tells us what is prohibited, whereas *Smoking is permitted in designated areas* tell us what is allowed: in different situations different deontic operators are favoured. *You are requested not to smoke* is directed at individuals, whereas *Liverpool Airport is a smoke free zone* is couched in objective terms. *No smoking* is very abstract, whereas *Smoking of pipes, cigars and cigarettes, and the use of e-cigarettes is not allowed* is rather specific. Finally there is the assumed context: when smoking was allowed on public transport in some carriages, sometimes we would see *smoking* signs on the smoking carriages and sometimes we would see *no smoking* signs on the non-smoking carriages. The former assumes no smoking to be the default, the latter assumes permission to smoke to be the default. Note that it was never obligatory to smoke

T.J.M. Bench-Capon
Department of Computer Science
University of Liverpool
Liverpool, UK
E-mail: tbc@csc.liv.ac.uk

[1] This is a revised and extended version of a paper presented at Jurix 2014 [11].

in a smoking carriage[2]. These four dimensions, the default status (desirable or undesirable), the viewpoint (legislator, norm subjects), the level of abstraction (high or low) and the favoured deontic operators (permission, prohibition or obligation) were identified in a paper by Winkels and den Haan published in ICAIL 1995 [26] but since (undeservedly) rather neglected[3]. In that paper Winkels and den Haan explore the roles that "deep structures" (formal descriptions of the environment in which the norms will operate) might play in drafting legislation. Sometimes an interesting idea appears but remains largely unexplored until the technology to handle it is developed and it reappears in a new context: I feel that this is what has happened to [26]. That paper concerned the automated drafting of legislation, and discussed how the same objectives could be achieved by different norms, and that these four perspectives could be used to frame laws in ways appropriate to different circumstances. But the use of "deep structures" can play a wider role in the design and analysis of norms, as work in other areas of Computer Science has since shown: for example process compliance (e.g. [18]), another area where norms are central and which has arguably received too little attentention.

Since 1995 a major feature of Computer Science research has been the development of Multi-Agent Systems (MAS), and within MAS the emergence of Electronic Institutions (e.g. [13], [16], [5]) and the general idea that open agent systems (e.g.[14]) can be controlled by norms (which can be traced back to [23]). This has led to a continuing growth of interest in the topics of designing and representing norms in the MAS community and norms for multi-agent systems is now a much studied topic. As well as the practical uses in regulating open systems, there has grown up a significant strand of more theoretical work: for one important and representative strand see e.g. [25], [2], [1][4]). The role of the "deep model" of [26] in the MAS systems is played by a transition diagram, a structure long used in the specification and design of software systems. The transition system indicates the effects of an agent's action in the context of the system, and norms are specified as constraints on such behaviour. Specification in [25] and subsequent papers such as [2] and [1] is in terms of Kripke structures and Computation Tree Logic (CTL) [15], but we will not use any particular formal structures in this paper, remaining at the level of the transition diagrams.

The term "norm" will be used in this paper to refer to any statement intended to influence behaviour, including e.g. laws, social conventions, and rules for electronic institutions. Typically there will be an authority which issues the norm (e.g. the legislature or systems designers[5]) and a group of subjects intended to comply with it (e.g. citizens or software agents). The MAS work is directed towards software systems and its concerns are mainly software engineering concerns, such as effectiveness, efficiency and liveness. Our discussion will return to the theme of [26] and relate the MAS work to human-directed norms: we will be thinking

[2] I once encountered a cafe in Paris with three areas: smoking, non-smoking and mixed, but I think this was a joke on the part of the *patron*.

[3] At the time of writing it had received only nine citations on Google Scholar.

[4] Awarded best paper prize at AAMAS-2011.

[5] Social conventions are an exception, although one's peers will soon make one aware of transgressions. Books of etiquette are published which are initially intended to describe existing social norms, but may, if sufficiently accepted, become authoritative sources, such as Nancy Mitford's essay about word usage *The English Aristocracy* published in *Encounter* magazine 1954.

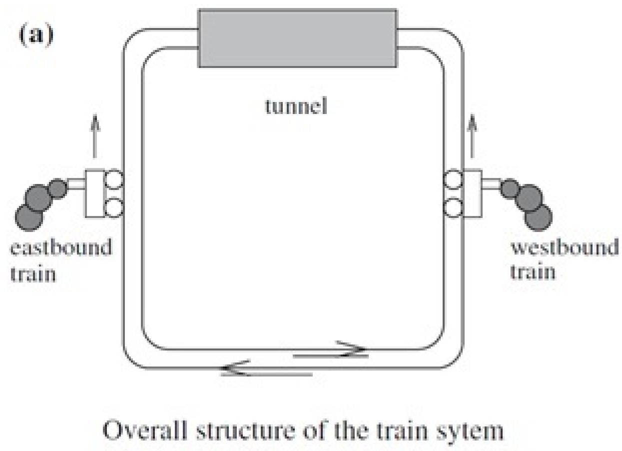

Fig. 1 Train system in [25]

mainly in terms of human agents and their legal systems. We will also relate the MAS work to the four dimensions identified in [26].

2 Models

In this paper we will focus our discussion on a model taken from [25], and use this simple model to explain how norms are represented in MAS. We will also make use of an important motivating example from [24], Ullmann-Margalit's seminal work on norms, in which she uses a variety of public goods games to discuss social norms and their origins.

2.1 A Simple Train System

Our main example is of two trains travelling in different directions around a circular track. Mostly the tracks are separate, but at one point both tracks pass through a narrow tunnel and a crash will occur if both trains are in the tunnel at the same time. The trains can either move forwards or remain where they are. With respect to the tunnel, a train may be in the tunnel, waiting to enter, or away from the tunnel. A picture of the system, taken from [25] is shown in Figure 1. Note that this description is very abstract: it may be that there are several points away from the tunnel any of which can be considered *away*, and this further level of detail is something we might need to consider in some circumstances.

Two viewpoints are relevant: that of the trains (the norm subjects) and that of the system as a whole, which will be the viewpoint of the train company or the norm issuing authority. We can represent these viewpoints using transition diagrams. Figure 2 shows the individual perspective, while Figure 3 shows the system perspective. A train may be *in*, *waiting* or *away*, and may *move* or *idle*. The

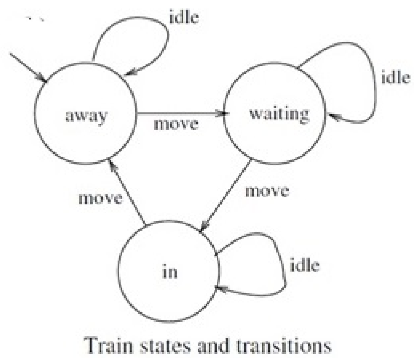

Fig. 2 Transition Diagram from Individual perspective taken from [25]

system perspective is more complicated because each train can be in any of the three states, giving rise to nine states instead of the original three, and because the transitions must now consider the *joint actions* of *both* agents[6] so that we have four actions available in each state (note that in Figure 2 we have omitted the joint actions where both trains are *idle*: these simply involve a transition back to the same state, and would clutter the diagram for little benefit) instead of two. The system perspective is also more informative: we can see that the state where both trains are *in* has no exit: so that idleness is enforced in that state. That moving in the waiting state may or may not result in a crash is not captured by Figure 1.

The possibility of a crash does, however, need to be represented from the individual perspective. Since we cannot represent the other agent from this perspective, we need to see the state reached by moving from the waiting position as indeterminate: the train does not know whether or not a crash will result. We indicate this in Figure 4 by labelling actions with a variable outcome with a "?". Also we add a state to indicate that a crash occurs. Note that the crash state does not have an exit transition.

The technique of [25] is to identify an objective (for example in [25] the objective is taken as avoiding crashes, i.e the state where both trains are in the tunnel), and express a norm as a behavioural constraint which will ensure that the objective is achieved. The behavioural constraint used in [25] is to prohibit both trains from waiting in the tunnel and to prohibit the *eastbound* train from entering the tunnel, except where it is waiting and the westbound train is away. This norm gives priority to the westbound train, which is never obliged to idle. The norm can be represented by removing all the transitions which include waiting in the tunnel

[6] One type of transition system, introduced in [25], which has joint actions as its transitions is Action-Based Alternating Transition Systems (AATS), adapted from the Alternating Transition Systems used to underpin Alternating-time temporal logic [3].

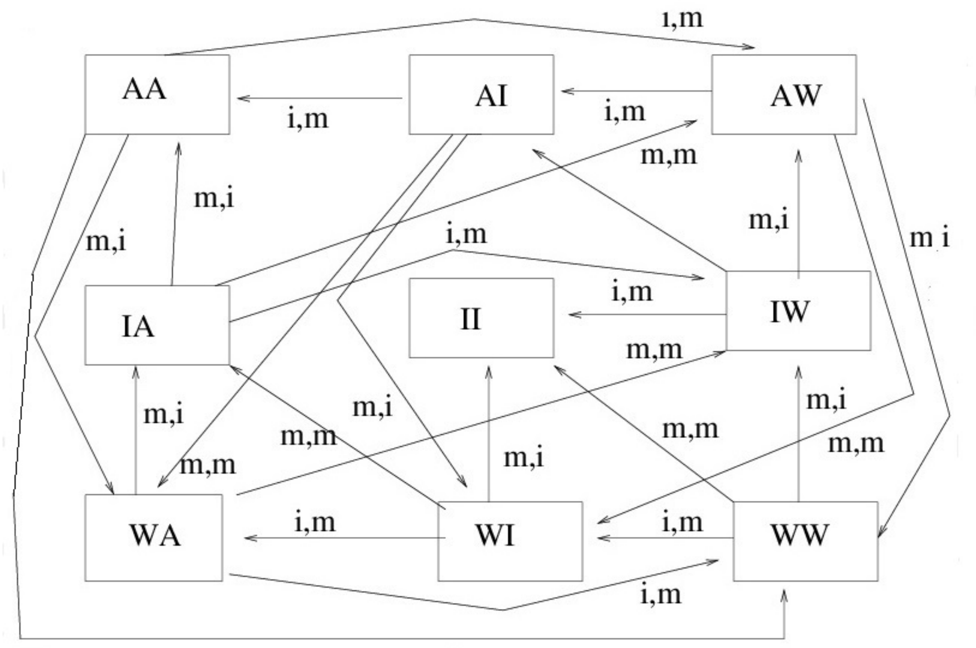

Fig. 3 System perspective on trains

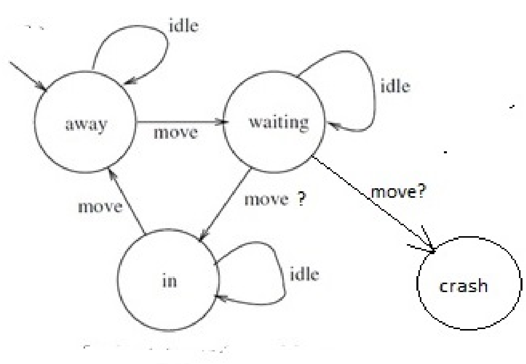

Fig. 4 Individual perspective of train with crash

and those which include the eastbound train moving in states WI and WW as part of the joint action (as shown in Figure 5). Notice how the number of transitions have been reduced, reflecting the prohibition of several actions, and, importantly, that the undesirable state can no longer be reached. For the individual perspective we now revert to Figure 1, although the eastbound train will have an additional precondition affecting its movement when waiting, since it has to look to see that the westbound train is *away*. Model checking tools, such as the MOCHA system of [4] can now be used on the amended transition diagram in Figure 5 to verify the effectiveness of the norm.

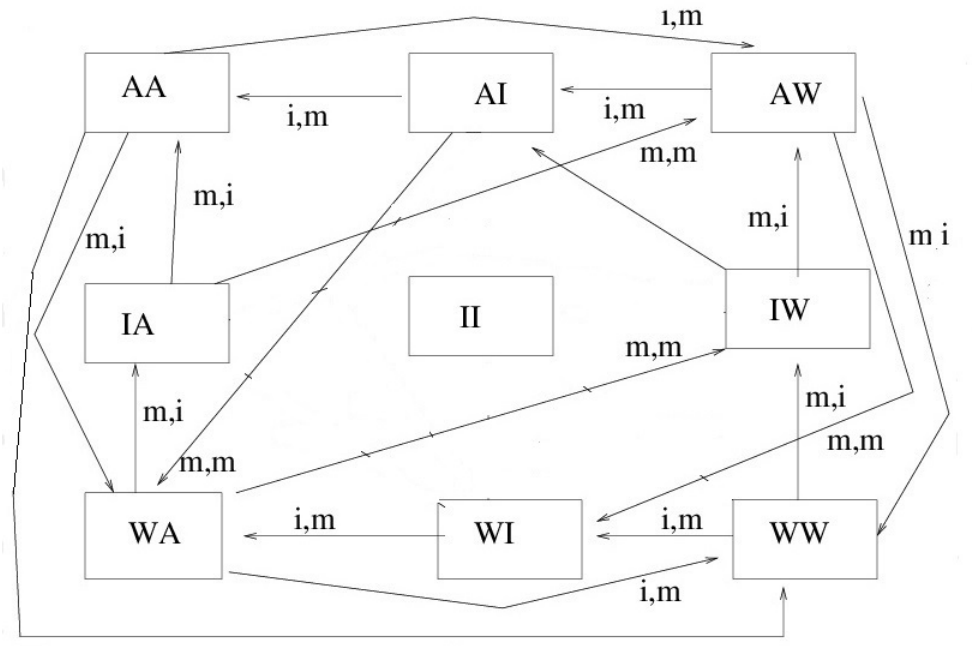

Fig. 5 System perspective on trains with norm

2.2 The Machine Gunners Dilemma

The example from [24] is the *Machine Gunners Dilemma*, a variant on the well known *Prisoners Dilemma*. I will extend the example slightly here. The idea is that there are two machine gunners who need to delay the enemy if reinforcements are to arrive, and victory is to be won. If both run away, the enemy will win, and both will be taken as prisoners of war. If one stays and one runs, the deserter will escape unharmed while the other will delay the enemy sufficiently even though in the end he will die (remember the Alamo). If both stay there is a possibility that both will survive (as at Rorke's Drift) or that both will die, but only after having succeeded in their mission (as with the three hundred of Thermopoylae), according to the skill and determination of the enemy. A payoff matrix for the gunners is shown in Table 1 and the payoff for the army is shown in Table 2.

	B stays	B runs
A stays	Hero and Possibly Survive, Hero and Possibly Survive	Hero but Die, Survive
A runs	Survive, Hero but Die	Captured, Captured

Table 1 Pay off matrix for the Gunners for Machine Gunners Dilemma

	B stays	B runs
A stays	Success, Possible Triumph	Success
A runs	Success	Failure

Table 2 Pay off matrix for the Army for Machine Gunners Dilemma

The objective of the army will be that at least one, or preferably both, the machine gunners remain at their post. The machine gunners most want to avoid dying, but would also like to avoid capture. The dominant game-theoretic strategy is to run, but both running is not best for the gunners themselves and is the one option that the army is anxious to avoid. Ullmann-Margalit discusses three possibilities in which the behaviour of the soldiers can be made to better fulfill the objectives the army as a whole. One is simply to make desertion impossible, by chaining the gunners to their guns: this is very similar to the usual MAS approach - the undesired actions are made unavailable. The second is to impose sanctions: if deserters are shot, then they will both choose to stay as this is the only possibility of survival. This works not by removing actions but changing the states they reach: capture and escape are no longer possible. The third possibility is the use of norms to change the preferences of the agents: two military sayings are *Death or Glory*[7] and *Death Before Dishonour*[8]. The idea is to instill sufficient trust in or loyalty to their comrades or sufficient confidence in the fighting qualities of their regiment that they will choose to remain at their posts. Military training typically attempts to induce all three in their soldiers through teaching regimental history, team building etc. The sanctionless military norms encountered in this example are unwritten, but the principles that underlie them have implications for legal norm formulation.

The following three sections will discuss enforcement, sanctions and norms which rely neither on enforcement nor sanctions.

3 Enforcement

As described above and shown in Figure 5, enforcement is the method normally adopted in MAS specification[9]. It is quite natural, and relatively easy to implement, in a software system. For example in an Electronic Institution users may be offered a menu of actions, and the forbidden actions may be greyed out, or otherwise made non-executable in certain contexts. But in the real world it is harder to enforce norms in this way: it is difficult to make it impossible for the eastbound train to move and the idea of chaining soldiers to their guns is surely contrary to the Geneva Convention. Moreover, it is usually not even desirable to make the prohibited action *impossible*: sometimes there is an overriding reason to violate the norm. Perhaps the westbound train is stuck in the tunnel and the only way to remove its passengers is for the eastbound train to (slowly and cautiously) enter the tunnel. The use of deontic modalities (as opposed to alethic modalities which result

[7] The Regimental Motto of the 17th/21st Lancers, now part of the Queen's Royal Lancers.

[8] A saying apparently widely used in the US Marines Corps.

[9] It has been argued that removing the possibility of violation means that we are no longer really dealing with norms at all (e.g. [20], [19]). None the less it is very common in MAS and needs to be discussed here.

from making prohibited actions actually impossible) is important when violations need to be considered [20]. In general the option of rendering actions impossible is not available to legislatures, although it remains a reasonable way to model a fully complied with norm, and so remains a convenient way of testing the effects of a norm (assuming universal compliance) using model checkers or even by proving properties of the system.

We can consider the solution of [25] and other MAS systems which use enforcement on the dimensions of [26]. Note that this applies to the transition system: the formulation of norms corrsponding to these modifications remains problematic [19]. First the viewpoint is that of the system: we need to consider both trains and their interaction. Second the default is that all actions are permitted. Third the level of abstraction is high: it is useful to be able to describe a whole set of possible positions as the single state "away". Lastly the preferred modality is prohibition: it is easy and natural to remove the transitions corresponding to prohibited actions. An alternative (for a software system) would be to start with no actions available to the agents, and then add permitted actions and their transitions. This is unrealistic in an open system where the agents participating are designed by people other than those running the institution, and so their repertoire of actions is beyond the control of the institution. This is even more so when the norm subjects are real people. In order to represent an obligation we must remove every transition except the obligated one from the relevant state. Prohibition (removing an action) will be at least as simple as, and typically simpler than, obligation, which will often involve removing a large number of actions. For example, if there are five actions available to an agent, we would need to remove four to represent the obligation: when each action forms part of several joint actions, the number of transitions to remove can increase quickly.

We can also see from this example that the rule and exception structure widely found in law is a very natural and concise way of specifying norms. In [26] the norm is specified on a state by state basis. But this is only feasible for very limited problems where the number of states is small. Much more concise is to say that *a train must move except if it is eastbound and waiting when it must idle except when the westbound train is away when it must move*. This now has the familiar *norm, exception, exception to exception* structure so often found in legislation, and long recognised and exploited in AI and Law e.g. [17]. Note, however, that the preferred modality has now become obligation.

The systems perspective is more natural than the individual: in an open MAS we are designing a system to control agents, not designing agents specifically to participate in a particular system. When norms are enforced in this way the agents themselves do not choose to conform to the norms, the regulation removes the possibility of transgression.

One should also note the role of abstraction. In the train example all locations of the train other than *in* the tunnel and *waiting* to enter are represented as the single state *away*. If *away* is a single location the eastbound train will not have very long to wait for the westbound train to reach the away state, and can see that the westbound train is away, and so will know when it is permitted to enter the tunnel. But at a more detailed level, we might find that the circuit is much longer so that there are many locations contained within the *away* abstraction. In this case the westbound train might be invisible to the eastbound train for long periods, making it uncertain whether it can enter the tunnel, and making the wait until it can be certain that it is safe to enter rather long. Another situation might be that the eastbound train has only a short distance between *in* and *waiting*, while the westbound train has a much longer circuit. In this case the eastbound train would have the longer wait for the other train to reach *away*, and so the norm should be reversed to give priority to the eastbound rather than the westbound train, if we wish to maximise the number of circuits the trains can perform. Two points are important: one is to frame the norms so that the subjects can know whether they are complying with them or not, and the other is that an asymmetry may appear to be acceptable at one level of abstraction but not at an another. We must therefore, when framing norms, be careful to use the appropriate level of abstraction, both to encourage compliance (since too long a wait may lead to disregard for the norm, especially since the longer circuit reduces the possibility of a crash) and to frame acceptable norms. If the track is long and so compliance becomes doubtful, we might need to introduce some kind of signaling system, whereby a signal would be lowered when the westbound train reached the waiting position and raised when it was in the tunnel. With the signaling system a norm such as *trains are prohibited from moving when the signal is down* will avoid crashes and unnecessary waiting on the part of the eastbound train.

Before moving on to sanctions, observe that in MAS treatment of norms the objective is to prevent a *state*, but the norm forbids an *action*. This is also very common in general (consider the ten commandments). Why is this so? I believe the main reason is the uncertainty of the effect of actions, because agents do not know in which joint action they will participate. Because we do not know for certain whether an action will lead to the unwanted state it is necessary to prohibit any action which *could* result in the unwanted state. In the train example, moving into the tunnel will often not result in a crash: waiting is only necessary if the other train enters the tunnel at the same time. But to *ensure* that this does not occur, that action must be prohibited unless the other train is *away*. *Keep off the Grass* signs exist not because and particular act of walking on the grass is undesirable: it is the wear on the grass caused by many individual acts of walking that is undesired. But to ensure that the grass is not worn the action is prohibited for all: calculating the consequences in terms of wear cannot reliably be left to individuals. None the less, there are many approaches in MAS, especially those closely related to deontic logics, which do forbid and obligate states. I shall, however, not discuss these in this paper.

4 Sanctions

The approach of enforcement of norms by making undesirable actions actually impossible left the transition diagrams unchanged, except for removing certain transitions so as to make certain states unreachable. Sanctions operate very differently since they add *additional* states, actors and transitions to the diagram to represent the application of sanctions and the agents responsible for applying them. Suppose we are unable to prevent the eastbound train from moving and so it would be unrealistic to remove the actions prohibited by the norm. We might instead attempt to ensure compliance with the norm proposed in [25] by imposing sanctions if the norm is violated: for example, fining the eastbound train if it violates the norm. We now need an additional agent to impose the sanctions, and

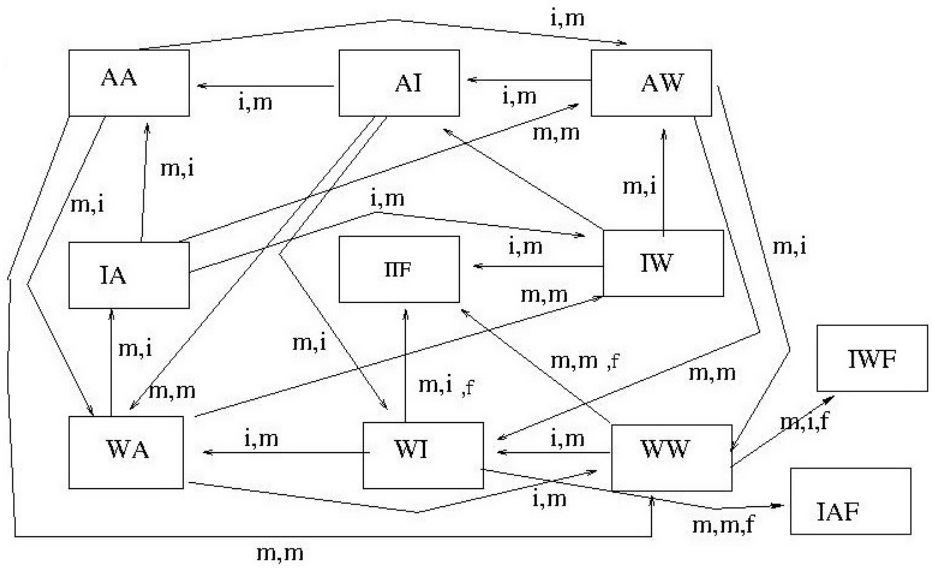

Fig. 6 Trains with Sanctions

we need to represent the action of imposing the sanction performed by that agent (and the corresponding joint actions) and we need to supply additional states to represent the situations in which the norm is applied. Thus the system perspective transition diagram ceases to be that of Figure 3 and becomes the diagram shown in Figure 6 (transitions following fines are omitted from the diagram, since we will not look beyond the sanction state). If we can assume that the sanctions are enforced whenever a violation occurs we get the situation shown in Figure 6. We can ignore any states reachable only by unpunished violations (since violations always meet with sanctions). Given a well designed norm this will include all the undesirable states, so that the undesirable states are not simply unreachable, but impossible.

We can now classify the states which result from sanctions to reflect the different ways in which sanctions can operate. It may be that the receipt of the fine makes the situation actually desirable to all the agents concerned. An example of this may be library fines: the library may well prefer to receive the income for the book than have it unborrowed on their shelves, and the library user may well be willing to pay to retain the book, or to return it at a more convenient later date. In this case the sanction makes everyone happy: it is less a fine than a charge for an extra service. The sanction works by making all states desirable: a win-win situation. If there is no level of fine which can make everyone happy, the fine should be sufficient that at least the norm-issuing authority is made happy. Certain motoring offences, such as parking (and, some motorists claim, speeding) may fit this. The fines provide revenue, which the authority welcomes, but which the users resent. But this does not matter: users can choose to obey the norm, and the states so reached will also be acceptable to the authority. Effectively the grievance of the norm subjects is as much that the amount charged for violation is too high as that the norm exists: motorists are relatively willing to pay a standard parking charge, when parking is metered rather than prohibited. In some cases norm subjects may disagree as to whether the sanction serves as a fee or as a discouragement. Some people see parking fines as worth paying and will not worry about accumulating parking tickets, whereas others will avoid parking illegally because of the fines. Essentially there is now a market in parking spaces, it which the more convenient ones attract a significantly greater fee.

The third possibility is that the consequences of the undesired action are so bad (from the perspective of the agent issuing the norms) that no level of sanction is able to compensate. Here the purpose of the sanction is not to *compensate* for the violation and so make the action acceptable to the norm-issuer, but to *deter* the norm subject from performing the action, since the sanction will make the situation unacceptable to the norm subject. The crash situation in the train example is one such. The objective is to ensure that there are no crashes: any revenue that might result from violations that do not lead to crashes is neither here nor there, since the important thing is to eliminate violations and hence the possibility of a disaster. Here the sanctions have to be such as to make the situation definitely undesirable to all norm subjects, whatever their resources and inclinations, so that they will choose to comply with the norm, unless there is some abnormal situation which makes violation absolutely necessary. Whereas fines are natural sanctions for some offences, they do not work so well as a deterrent: it will always be possible that there are some people who consider the price worth paying. Thus such violations tend to be associated with custodial (or even corporal or capital) sanctions, which are a very different type of sanction from fines, since imposing them represents a cost to the enforcer rather than compensation. For this reason they should be reserved for avoiding situations which are so undesirable that they cannot be compensated for.

The above assumes that the sanctions will be always enforced. This is often the case, where detection is not a problem: library fines can always be exacted when the book is returned. In the train example a single CCTV camera would ensure that violations were detected. But speeding offences often go undetected, and speed limits are frequently exceeded without punishment as a result. Thus the norm subject may be uncertain whether violating the norm will reach the state where the penalty is imposed or a state where the violation goes unpenalised. If the subject is to be deterred by the penalty, the subject will need to assess the likelihood of detection. Given sanctions at the appropriate level, if it is certain that the sanction will be imposed, the norm will be obeyed and if it is certain not to be imposed the norm will be violated, but if the sanction is a more or less likely consequence, the norm subject must weigh the benefits of violation against the risk of the sanction. In this case the authority must make the sanction larger than it would be if detection was thought certain, so that the *expected value* will be undesirable to the norm subject, even with the smaller risk of detection. For example if norm subjects consider that there is a 50% chance of detection and the sanction applied, the fine would perhaps need to be twice as large to keep the same expected values.

Briefly returning to the Machine Gunners Dilemma, suppose that a sniper is stationed so that any gunner attempting to flee will be shot, but that the sniper will only be able to shoot one target. The payoff matrix becomes that shown in Table 3.

	B stays	B runs
A stays	n% Risk of Dying, n% Risk of Dying	Die, Die
A runs	Die,Die	50% Risk of Dying, 50% Risk of Dying

Table 3 Pay off matrix for Machine Gunners Dilemma with Sniper

From this the choice of action will turn on the way the gunner assesses the fighting abilities of his colleague and the enemy. If there is a greater than 50% chance that both will survive (a Rourke's Drift rather than a Thermopolyae situation) then staying is the rational course. If not, perhaps taking the chance with the sniper will seem the better option, unless there is also a belief that his colleague will stay, leaving him as the only target. This illustrates a second way in which regimental loyalty can operate: it is not necessary for the gunner himself to be loyal, it may be enough that he believes that his colleague will remain loyal.

4.1 Rewards

Before leaving the topic of sanctions we should consider the possibility of rewards, as raised in [12]. Just as norms may be couched in terms of prohibitions or obligations, it is possible to encourage obedience to norms by rewarding compliance rather than punishing violation. Essentially the idea is that instead of making the situations undesired by the norm issuer less attractive to the norm subject, the situations desired by the norm issuer are made more attractive to the norm subject. In the analogue of the sanctions as fees situation, the norm issuer is offering to pay for compliance, rather than charging for violation. Sometimes the difference may be no more than a matter of framing. Suppose the norm is that three months notice should be given when renewing a passport. A discount for early application and a surcharge for late application, may not imply any difference in fee. If the fee is \$100 with a 10% discount if three months notice is given, it is equivalent to a \$90 fee with a \$10 surcharge for late application. In such cases the choice can be simply pragmatic: which framing is found to be more effective in securing timely applications?

In practice sanctions seem to be more common that rewards. This may be because rewards have a perceived cost for the norm issuer, whereas sanctions may even provide revenue. But this neglects enforcement costs. Boer's idea in [12] is that it is a matter of evidence. In a reward regime the norm subject supplies evidence of compliance, whereas in a punishment regime the norm authority needs to detect (and produce evidence of) violation. Thus in situations where evidence of violation is difficult (or costly) to collect, a reward regime may well be more effective, and perhaps even less costly.

Note also that a reward regime need not fall foul of the objection to formulating law in terms of obligation rather than prohibition in section 3. There it was observed that whereas representing a prohibition required the removal of only one transition, obligation might require the removal of several transitions. With sanctions and rewards, there is no such asymmetry: both require modification of one state.

Returning briefly to the Machine Gunners Dilemma, rewards have been widely used in military contexts. Medals are one such reward (eleven Victoria Crosses, the highest British Army award for gallantry, were issued at Rorke's Drift), and War Widows pensions are another. Here, especially since the risks required to be taken are so high, and the enforcement of the norms so difficult in the fog of war, it becomes particularly important that the norm subjects themselves choose to comply. This theme is continued in the next section.

5 Norms Without Sanctions

In this section we will consider norms which are intended to be obeyed without enforcement and without sanctions. These are norms which secure compliance of the third type considered by Ullmann-Margalit when discussing the Machine Gunner's Dilemma in [24], where the norm subjects choose to obey. The first norms here are norms of *co-ordination*. Norms are needed when the norm subjects are confronted by a choice, and one choice will, or may, lead to an unacceptable situation. Often the norm subject will prefer the choice which may lead to the unacceptable situation: in the train example, the train wants to move, but there is a need to avoid crashes. In such cases norms serve to remove the choice (enforcement) or to alter the outcomes by imposing sanctions so that norm subject no longer wishes to make the undesirable choice, provided the norm subject believes that the sanction will be imposed.

But in some cases the norm subject is indifferent as to the action performed, and the undesirable situation arises from two or more agents making choices which together have unfortunate consequences. The classic example is norms relating to the side of the road that cars should drive on. Drivers do not particularly mind which side of the road they are required to drive on: some countries say left and others say right and no one argues for change. What matters is that everyone makes the same choice. Thus once the norm has been established, so that the behaviour of others is known, compliance becomes the only sensible choice: agents obey not because they cannot do otherwise, nor because they wish to avoid sanctions, but because it is in their interest to do so.

The train example can be seen as a coordination norm if we assume that, although both trains want to move, they are more concerned with safety. If this is the case, without the norm neither will dare to enter the tunnel and the rail system will grind to a halt. But if the norm is promulgated, and the trains assume that it will be obeyed, the westbound train can enter with confidence and the eastbound will be happy to suffer a short wait to ensure its safety. Similar solutions are used for narrow bridges on roads: by giving priority to one direction, accidents can be avoided and the certainty the norm affords more than compensates for any delay for the unfavoured direction.

This leads to the notion of reasoning about the behaviour of others, since what the other chooses often determines whether the outcome will be good or bad. The train problem was discussed in [8], which used the machinery of [6] to describe the reasoning of agents considering the choices. In brief [6] augments an AATS by labelling transitions with the (social) values[10] promoted or demoted by following the transition. These values supply reasons for and against choosing the action required for the transition, and the action is then chosen according to the agent's (subjective) preferences between the values, using a Value Based Argumentation framework [10] to resolve conflicting reasons. In the train examples the relevant values are *Progress* and *Safety*. Progress will be promoted by a transition in which a train moves, and Safety will be demoted by a transition into the state where both are in the tunnel. If the trains prefer Safety, there will be standoff as they wait for one another, unless there is a norm to constrain their behaviour.

In this case the norm works by altering the labels on the transitions. Effectively we can see the existence of the norm as introducing one or more additional values, such as *Compliance*, promoted by complying with the norm and demoted by violating it. Such a value can itself be a reason for action, and a strong one: some people will obey the law simply and purely because it is the law. (Kant argued in [21] that this was the only moral reason to obey the law, and it is very often at least a contributory reason.) This additional value makes obeying the law more attractive (unless breaking the law is seen as a value in itself, as with the adolescent playing chicken). An alternative is to make the undesired choice less attractive, which can also be effected by a norm: if society attaches a stigma to breaking the law, then the norm violator's reputation will suffer. There is a similar contrast in the slogans mentioned in relation to the Machine Gunner's Dilemma: *death or glory* gives a value-based reason to comply, while *death before dishonour* gives a value-based reason not to violate. Different types of norm may relate to different values: obeying the law is a different value from conforming to a social convention: either may be rated more highly by an agent and they may even conflict (e.g in gang culture). The same norm can work differently for different agents: the norm will still work provided that one or other of the values, or a combination of the two, can serve to direct the agent into the desired choice. Sanctions go further: the effect of the sanction is to change the state reached by the undesired act and in so doing cause the transition to that state to demote some value prized by the norm subject. This will in turn provide additional reasons against violation when the agent chooses an action.

5.1 The Role of Values

In MAS, [25] for example, the objective of avoiding of some particular state or states is simply a given. While it is intuitively plausible that the railroad would wish to avoid crashes, the precise reason is not given in [25] and so the reason why the state is undesirable is not explained. It could be because a crash means that movement is no longer possible so that the railroad will cease to function; or it could be concern for the safety of passengers and staff, or it could even be because of the cost of repairing or replacing the trains. If, however, we think in terms of the values of the norm authority and the norm subjects, it is possible to derive the objectives from the values, using the techniques of [8] and [7]. Suppose that an important value for the railroad and both trains is Progress (**P**), so that any

[10] Like [6], this paper does not consider accrual when multiple values are promoted or demoted by an action, not differing degrees of promotion and demotion. These issues are important, but require further investigation.

transition involving movement of either train will promote **P** for the railroad, and any transition involving movement of a particular train will promote **P** for that train. Now we can see that this will give rise to an avoidance goal of not entering the state where the trains crash, since this will preclude any further movement by either train, so that all will accept this goal. But there is also an achievement goal in moving to the next state: this applies to the company and to both trains, although different transitions will promote **P** for the different trains. Of the four transitions from a given state, one will promote **P** for both trains and the company, two for one of the trains and the company and one for no one.

But we also need to consider Safety (**S**): when waiting the trains will have a reason to idle, unless they can see that the other train is *away*, since the transition to a crash will demote **S**. Thus when *waiting* a train will have three competing arguments: move to promote P in the short term (P_S); idle to avoid demoting P in the long term (P_L); and idle to avoid the risk of demoting S. Each train has six possible value orderings (although strictly the power of the argument based on S depends on the assessment of, and aversion to, the risk, disregarding the risk has the same effect as valuing S least, and so need not be considered separately). The action that will be chosen on each ordering is shown in Table 4.

Preference ID	Value Order	Action
1	$P_L > P_S >$S	idle
2	$P_L >$S $> P_S$	idle
3	$P_S > P_L >$S	move
4	$P_S >$S $> P_L$	move
5	S $> P_S > P_L$	idle
6	S $> P_L > P_S$	idle

Table 4 Train preferences and consequent actions

Note that the trains will select these actions both when the other train is *waiting* and when it is *in*, since they cannot discriminate the two states. A crash will occur only if *both* drivers have Preference ID (PID) 3 or 4 (because otherwise at least one of them will choose to idle). However, from the train company perspective, neither driver having PID 3 or 4 will also be undesirable since the railroad will come to a standstill with both trains waiting, neither daring to enter the tunnel. Thus it could be that diversity amongst the train drivers would be sufficient give rise to satisfactory operation of the system: this, however, cannot be relied upon and so the norm is required. Note, however, that the norm is no hardship on a driver with PID 1,2,5 or 6, since the driver would choose to behave in this way without the norm. Thus ideally a train company relying on driver preferences to avoid crashes should attempt to arrange that if only one of its drivers has PID 3 or 4 that driver is given the westbound role, since in that way compliance is rendered the more certain. Thus it may be that, despite its apparent symmetry, effectiveness of the norm may depend on the preferences of the driver of train which is required to give way.

Finally we may consider trains which include reasoning about the other train. Suppose there is no norm: now if drivers believe that their colleague has PID 1, 2, 5 or 6, they may decide that it will be safe to move anyway since their colleague will wait. The lower the ranking of P_S attributed to the other driver, the less will be the perceived risk of entering the tunnel, and the lower the weight accorded to S. The danger is, of course, that both drivers may think like this.

Thus, while in many situations values may (through different preferences) enable the system to function, there will always be a danger of calamity. Therefore it probably prudent to have the norm in place, although consideration of the value orderings of the norm subjects can aid in formulating the norm so as the better secure compliance.

6 Perspectives and State Transition Diagrams

Let us now return to the four perspectives of [26] and see how they relate to the transition diagrams we have been discussing.

6.1 Viewpoint

Two viewpoints were discussed in [26], that of the norm-issuer and the norm subject. Differences in viewpoint manifest themselves in several ways in the transition diagram. First the viewpoint can have epistemological implications. Compare the transition diagram in Figures 2 and 4 with that in Figure 3. Figures 2 and 4 confine themselves to what can be known by a particular train: the states do not include the whereabouts of the other train, nor do the transitions include the actions of the other train. In contrast Figure 3 offers a “birds eye” view of the system, in which each of the states of Figure 2 become three different states according to the location of the other train, and each action becomes two joint actions depending on what the other train does. When designing norms, the more general viewpoint is undoubtedly more useful, and the agent itself will find it more profitable to reason with the bigger picture, assuming it is aware of the other train and its capabilities, although this may require it to make assumptions about where the other train is and what it will do. The individual viewpoint may be useful for understanding the reasoning of norm subjects, but when thinking about the system, the wider viewpoint captures more of what needs to be considered.

Viewpoint might also relate to the labelling of the transitions. From the point of view of the train company, movement is promoted when either train moves (and perhaps promoted more strongly when both trains move). But from the viewpoint of an individual train, movement is only promoted when the train itself moves. Thus in Figure 3 when in state $< A, A >$ the transition to $< W, W >$ promotes movement from all three viewpoints, but the transition to $< W, A >$ promotes movement only for the train company and the eastbound train. Since the train company is designing the norm, its viewpoint takes precedence in norm design, but it may wish to consider the viewpoint of norm subjects when thinking about whether the norm will be complied with. The suggested norm for example, prevents the promotion of movement from the viewpoint of the eastbound train when it is *waiting* and the westbound train is not *away*. Thus there may be a temptation for the eastbound train to violate in that state, whereas the westbound train never finds its progress impeded and so will not be tempted to violate the norm for that reason. This may well be useful when considering the practical effectiveness of a norm, or when choosing between two theoretically effective norms.

6.2 Degree of Abstraction

Getting the degree of abstraction right is a large part of the skill in constructing state transition diagrams. It is desirable to keep the number of propositions in a state to a minimum, because n propositions give rise to 2^n states. Thus if we can group together cigarettes, cigars, pipes and e-cigarettes as "smoking materials", we will save ourselves a good deal of complexity. Similarly in the train example, having a single state for *away* even if the circuit is quite lengthy, allows us to keep the number of states within reasonable bounds. Similarly we want to restrict the number of agents and their actions. Given m agents each with n actions, there will potentially be $n * m$ joint actions in every state, each requiring a destination to be assigned and one or more value labellings to be determined. In the train example it is not of relevance how many west and east bound trains there are: the situation of interest is when there is some eastbound and some westbound train in the *waiting* state. Thus abstracting all trains bound in a particular direction into a single train keeps the joint actions to a minimum, without losing anything of importance. The same is true of the actions: the trains are capable of moving at a variety of speeds, and so we could distinguish between *move fast* and *move slow*. Now it could be that if the trains enter the tunnel slowly they will not crash, but become immobilised by blocking one another. Since the state $< I, I >$ remains undesirable (from everyone's perspective) as further progress is now impossible, we need not be represent the different speeds. Moving slowly into the tunnel does not even avoid the risk of a crash: the other train may continue to move fast. Thus decreasing the level of abstraction by including the different movement speeds would have no gain in this case.

In contrast, suppose that the situation was not, as shown in Figure 1, symmetrical, but that the westbound train had a much longer piece of track in the *away* state than the eastbound train, and that much of this was not visible to the eastbound train. This could either be represented by a number of additional *away* states (at a minimum *away and visible to the eastbound* and *away and invisible to the eastbound*. Now while at the level of abstraction in Figure 4 there is no difference whether we require the eastbound to wait for the westbound or the westbound to wait for the eastbound, the new level of abstraction will reveal reasons to make the westbound wait. First, this will maximise movement: the eastbound will have a longer wait that the the westbound. Second, the westbound is more likely to comply than the eastbound: quite apart from any sense of injustice the eastbound might feel, the probability of a crash is less for the eastbound since there are several safe states that the westbound might be in, and the longer wait demotes the progress of the eastbound train to a greater degree. Again therefore, more detail may be required when assessing the practical as opposed to the theoretical effectiveness of the norm.

6.3 Default Status and Favoured Operators

We will consider these together. Normally we create the state transition diagram so as to show all possible actions (given a repertoire at the desired level of abstraction), suggesting that the default that everything is permitted. For the enforcement approach we now remove actions. Preferring to couch norms in terms of prohibitions rather than removals will normally be preferable, since this will require the removal of fewer transitions. Only for closed systems, where we design the agents to fulfill the system specification rather than design norms to control the behaviour of autonomous agents, does the approach of starting with no transitions and then adding transitions to enable the agents to carry out the required behaviour (i.e. starting from a context in which everything is forbidden) seem feasible.

As we have suggested, however, the enforcement approach is less feasible for human societies that societies of software agents. Here we need to distinguish between what is possible from what the norms allow since compliance is not assured, and so all the transitions will remain part of the diagram, since they represent options which the agents can choose. Here we may use sanctions to encourage compliance, and the natural modality will typically be prohibition, so that the sanctions are associated with states reached by violations of the norm, rather than with states reached by non-performance of obligations[11]. Rewards, on the other hand will, by the same reasoning, typically be associated with obligations.

If the norm authority is using neither sanctions nor rewards, it is forced to rely on a desire to comply (or an aversion from non-compliance) on the basis that this value is demoted. Of course, the psychological response of norm subjects may differ, and this may lead the norm to be phrased in different ways. Consider, for example, the difference between a request (*please refrain from smoking*), a command (*no smoking*) and a statement of fact (*Liverpool airport is a smoke free zone*), and the different reactions these might produce. Again pragmatic consideration of which formulation will encourage compliance is a factor in choosing how to express the norm.

Another factor which may influence the ways in which norms are expressed is the type of goal involved. In [7] four different kinds of goals were identified:

- *Achievement goals* which promote a value by making something currently false true.
- *Remedy goals* which promote a value by making something currently true false.
- *Maintenance goal* which promote a value by keeping something true true.
- *Avoidance goals* which promote a value by keeping something false false.

It might be thought that obligation is more natural for the first two of these and prohibition more natural for the latter two. But this is not so: changes occur independently of the norm subjects, and so sometimes action must be taken if the *status quo* is to be preserved. The important point here is these exogenous changes need to be represented in the diagram: we cannot rely on the idea that if an agent does nothing the state will remain the same. For this reason the state transition

[11] Moreover, non-performance of obligations is harder to specify and to detect, see, e.g., [27]. Sometimes, however, this will depend on the formulation of the action: forbidding tax evasion and insisting on payment of taxes require the same action. Here, however, the obligation will be preferred because there is only one way of paying and many ways of evading. In general the formulation requiring the fewest transitions to be sanctioned should be chosen.

diagram needs to reflect these exogenous changes. One common way to do this is to introduce a particular agent (often called *Nature*) to be seen as the actor bringing about these changes. Now the transitions will always be joint actions, at least joint between the norm subject and Nature, so that the idea of the active prevention of change can be modelled. Nature is also useful for the representation of non-deterministic actions, such as tossing a coin. Sometimes a well intentioned action (call it A), can have bad consequences. Now A will have two transitions, one where Nature "co-operates" and one where it does not. This use of Nature can provide a useful way of modelling frame axioms (e.g. [22]), which must, of course, be taken into consideration when designing norms. See e.g. [9] for an example of the use of Nature in modelling an AI and Law problem.

7 Summary

Norms can be modelled using transition systems. From the above we can see that we normally need to adopt a system perspective, so that we can represent the effects of the actions of a group of agents acting simultaneously, since the outcome of an action by one agent will very often depend on what the other agents choose to do. We therefore typically perform our analysis on a transition diagram in which the transitions relate to the joint actions of a group of agents, such as the Alternating Action-Based Transition Systems (AATS) of [25] and [6].

Norms are promulgated by an authority (legislature, social group, software designer, etc) when agents may act so as to realise a state undesirable to the authority (and perhaps also the agents themselves). This can be because of uncertainty, which may relate to uncertainty as to the current state (the train does not know whether the other train is waiting or not), or uncertainty as to what the other will do, or because the effect of the actions is indeterminate and the agents consider the risk acceptable. The norm subjects may, however, wish to enter a state the authority finds undesirable because it is personally desirable to them: there is no reason why the value preferences of the authority and the norms subjects should coincide. In summary a norm is required if an authority wishes to avoid a state and there is one or more norm subjects who can act so as to enter the state, and:

1. The norm subject desires to enter the state
2. The norm subject does not know whether its action will cause the state to be entered because
   (a) It is unsure of the current state;
   (b) It is unsure of the outcome of its action because of uncertainty as to what other agent will do;
   (c) It is unsure of the outcome of its action because the result of its action is indeterminate.

Using the general notion of exploring norms through models in the form of state transition diagrams, we have seen that there are several ways of representing a norm in a transition system.

1. By removing transitions representing prohibited actions;
2. By removing all the transitions from a given state except those representing an obligatory action;

3. By imposing sanctions for violation. This involves modifying the transition diagram to include an additional agent and appropriate transitions to impose the sanction and additional states to represent that the sanction has been imposed. Sanctions may be designed to produce
   (a) A situation acceptable to both authority and norm subject. Such sanctions are akin to *fees*.
   (b) A situation acceptable to the authority. Such sanctions are intended to be *compensation*.
   (c) A situation unacceptable to the norm subject. Such sanctions are intended to be *deterrent*.
4. By offering rewards for compliance. Only the first of the situations applicable to sanctions seems applicable to rewards: here the reward can be seen as a *payment* by the norm authority.
5. By labelling the transitions to represent promoted and demoted values (as in the AATS+V of [6]). The norm will then extend the set of possible labels (since the values will now include at least *Compliance* in addition to the existing values) and these new labels must be applied to the diagram. These values will then motivate the agents to avoid the undesirable situation. They can work
   (a) By removing uncertainty as to what the others will do (*co-ordination* norms).
   (b) By making the undesired state less attractive to the norm subject (compare sanctions).
   (c) By making a choice other than the undesired state more attractive to the norm subject (compare rewards[12]).
   (d) A combination of (5b) and (5c).

Of these, (1) and (2) represent *enforcement*. This is a simple and effective way of modelling a system in which the norm subjects can be forced to comply (such as an Electronic Institution or other software systems) or of determining the properties of the system under the assumption that all the norm subjects do, in fact, comply. It is, however, unsuitable if we want violation to be possible (either because there may be certain situations in which violation is acceptable, or even desirable), or where forced compliance is impossible (as in the normally the case in human societies). (3) represents sanctions and we can distinguish three sanction types (3a), (3b) and (3c). (4) represents rewards and (5) represents norms which require neither enforcement nor actions by the norm authority: it is intended that the agents will adopt the norms for their own reasons. This requires either the right sort of situation (so that what is needed is co-ordination: in particular the norm subjects themselves must wish to avoid the situation the authority finds undesirable) for (5a), or that the norm subjects have the appropriate value preferences (for (5b), (5c) and (5d)). The latter option may require some programme of education or training, to attempt to bring it about that agents have the appropriate values and preferences between them.

[12] In fact, whereas (3) is much more common that (4), (5c) is perhaps more common that (5b): this probably reflects that this method of seeking compliance is used when sanctions are difficult or impossible to enforce, and so compliance relies on the free choice of the agents concerned. This is closely related to the point about evidence in [12].

8 Conclusion

Although the design and representation of norms and their analysis to determine efficiency and efficacy is a potentially important topic for AI and Law, it has received very little attention in that field over the last two decades. In contrast recent years have seen a surge of activity on this topic in MAS, as norms are seen as a practical way of regulating Electronic Institutions and open agent systems in general, and normative systems have become the subject of theoretical study. In this paper I have considered the techniques used for MAS from an AI and Law standpoint, and discussed a variety of factors which can influence compliance with the norms. My intention is to highlight a number of issues, to provide a context in which these matters can be reasoned about, and to provide a framework which I hope will encourage further developments of this topic in AI and Law.

References

1. T. Ågotnes, W. van der Hoek, M. Tennenholtz, and M. Wooldridge. Power in normative systems. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 145–152. International Foundation for Autonomous Agents and Multiagent Systems, 2009.
2. T. Ågotnes, W. van der Hoek, and M. Wooldridge. Robust normative systems. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*, pages 747–754. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
3. Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.
4. Rajeev Alur, Thomas A. Henzinger, Freddy Y. C. Mang, Shaz Qadeer, Sriram K. Rajamani, and Serdar Tasiran. MOCHA: modularity in model checking. In *Proceedings of Computer Aided Verification, 10th International Conference*, pages 521–525, 1998.
5. J. Arcos, M. Esteva, P. Noriega, J. Rodríguez-Aguilar, and C. Sierra. An integrated development environment for electronic institutions. In *Software agent-based applications, platforms and development kits*, pages 121–142. Birkhäuser Basel, 2005.
6. K. Atkinson and T. Bench-Capon. Practical reasoning as presumptive argumentation using action based alternating transition systems. *Artif. Intell.*, 171(10-15):855–874, 2007.
7. Katie Atkinson and Trevor Bench-Capon. States, goals and values: Revisiting practical reasoning. In *Proceedings of 11th Intl. Workshop on Argumentation in Multi-Agent Systems*, 2014.
8. Katie Atkinson and Trevor J. M. Bench-Capon. Taking the long view: Looking ahead in practical reasoning. In *Proceedings of the Fifth International Conference on Computational Models of Argument (COMMA 2014)*, pages 109–120, 2014.
9. Katie Atkinson, Trevor J. M. Bench-Capon, Dan Cartwright, and Adam Zachary Wyner. Semantic models for policy deliberation. In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, pages 81–90, 2011.
10. Trevor J. M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *J. Log. Comput.*, 13(3):429–448, 2003.
11. Trevor J. M. Bench-Capon. Analysing norms with transition systems. In *Legal Knowledge and Information Systems - JURIX 2014: The Twenty-Seventh Annual Conference*, pages 29–38, 2014.
12. Alexander Boer. Punishments, rewards, and the production of evidence. In *Legal Knowledge and Information Systems - JURIX 2014: The Twenty-Seventh Annual Conference*, pages 97–102, 2014.
13. C Di Napoli, C Sierra, M Giordano, P Noriega, and MM Furnari. A pvm implementation of the fishmarket multiagent system. In *ISAI/IFIS 1996. Mexico-USA Collaboration in Intelligent Systems Technologies. Proceedings*, pages 68–76. IEEE, 1996.
14. Mark d'Inverno, Michael Luck, Pablo Noriega, Juan A. Rodríguez-Aguilar, and Carles Sierra. Communicating open systems. *Artif. Intell.*, 186:38–94, 2012.

15. E.A. Emerson. Temporal and modal logic. *Handbook of Theoretical Computer Science, Volume B: Formal Models and Sematics (B)*, 995:1072, 1990.
16. Marc Esteva, David De La Cruz, and Carles Sierra. Islander: an electronic institutions editor. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 3*, pages 1045–1052. ACM, 2002.
17. Thomas F. Gordon. Oblog-2: A hybrid knowledge representation system for defeasible reasoning. In *Proceedings of the First International Conference on Artificial Intelligence and Law,*, pages 231–239, 1987.
18. Guido Governatori. Business process compliance: An abstract normative framework. *Information Technology*, 55(6):231–238, 2013.
19. Guido Governatori. Thou shalt is not you will. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pages 63–68, 2015.
20. A. Jones and M. Sergot. Deontic logic in the representation of law: Towards a methodology. *Artificial Intelligence and Law*, 1(1):45–64, 1992.
21. I. Kant. *Groundwork of the Metaphysic of Morals*. Harper Perennial modern thought. Harper Collins, 2009.
22. Raymond Reiter. Formalizing database evolution in the situation calculus. In *Fifth Generation Computer Systems '92: Proceedings of the International Conference on Fifth Generation Computer Systems.* , pages 600–609, 1992.
23. Y. Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial agent societies (preliminary report). In *AAAI*, pages 276–281, 1992.
24. E. Ullmann-Margalit. *The emergence of norms*. Clarendon Press Oxford, 1977.
25. W. van Der Hoek, M. Roberts, and M. Wooldridge. Social laws in alternating time: Effectiveness, feasibility, and synthesis. *Synthese*, 156(1):1–19, 2007.
26. R. Winkels and N. Den Haan. Automated legislative drafting: Generating paraphrases of legislation. In *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, pages 112–118. ACM, 1995.
27. Adam Wyner. *Violations and Fulfillments in the Formal Representation of Contracts*. PhD thesis, King's College, London, 2006.