

# When and How to Violate Norms

Trevor Bench-Capon and Sanjay Modgil

*Department of Computer Science, The University of Liverpool, UK*

*Department of Informatics, King's College, London, UK*

**Abstract.** There is an increasing need for norms to be embedded in technology as the widespread deployment of applications such as autonomous driving and warfare becomes ever closer. Current approaches to norms in multi-agent systems tend either to simply make prohibited actions unavailable, or to provide a set of rules (principles) which the agent is obliged to follow. We argue that both these approaches are inadequate: in order to meet unexpected situations agents must be capable of violating norms, when it is appropriate to do so. This in turn requires that agents be able to reason about what they should do without reference to the norms. One way to achieve this is to conduct value based reasoning using an argumentation scheme designed for practical reasoning. Such reasoning requires that agents have an acceptable set of values and an acceptable ordering on them. We discuss what might count as an acceptable ordering on values, and how such an ordering might be determined. Law breaking is illustrated through a simple road traffic example.

**Keywords.** Norms, Violations, Agents, Values. Argumentation

## 1. Introduction

<sup>1</sup>Today there is an increasing need for norms to be embedded in technology as the widespread deployment of applications such as autonomous driving and warfare and big data analysis for crime fighting and counter-terrorism becomes ever closer. Current approaches to norms in multi-agent systems tend either to simply make prohibited actions unavailable e.g. [31] or to provide a set of rules (principles) which the agent is obliged to follow, in the manner of Asimov's Three Laws of Robotics. Neither of these methods can be seen as satisfactory ways of providing moral agents: agents able to reason with and act in accordance with norms, since not only is it in the nature of norms that they *can* be violated, but that circumstances can arise where they *should* be violated. In fact norms, especially legal norms, are, in real life and also in MAS, typically backed by sanctions [8]. The idea behind sanctions is to change the consequences of actions so as to make compliance more pleasant and/or violation less pleasant<sup>2</sup>. As noted in [8], sanctions can be seen as *compensation* (like library fines) when they can

---

<sup>1</sup>A shorter version of this paper was presented at the AI for Justice workshop, ECAI 2016.

<sup>2</sup>In decision theoretic terms, the ideal is for the action to yield an overall negative utility.

be viewed as a charge for violation, which makes the situation acceptable to the norm issuer, or as *deterrents*, where the sanctions are meant to ensure compliance by relying on the self-interest of the norm subject. When the norm *should* be violated sanctions may be problematic as they dis-incentivise the agent. This problem can be lessened in cases where the violation can be condoned and the sanction not applied, but this requires an agreement between the agent and the agent imposing the sanction, that the violation was justified (often not the case: consider dissidents such as Gandhi and Mandela). Moreover sanctions need to be expected to be enforced, otherwise agents may risk escaping punishment, and violate the norm when there is no acceptable reason to do so.

An important reason for thinking in terms of norms is the recognition that on occasion they need to be violated [22]. While the norm is intended to provide a useful heuristic to guide behaviour, allowing for a quick unthinking response, unreflecting adherence to such moral guidelines is not what we expect from a genuinely moral reasoner. R.M. Hare, a leading moral philosopher of the last century, expressed it thus [20]:

There is a great difference between people in respect of their readiness to qualify their moral principles in new circumstances. One man may be very hidebound: he may feel that he knows what he ought to do in a certain situation as soon as he has acquainted himself with its most general features ... Another man may be more cautious ... he will never make up his mind what he ought to do, even in a quite familiar situation, until he has scrutinized every detail.

Hare regards both these extreme positions as incorrect:

What the wiser among us do is to think deeply about the crucial moral questions, especially those that face us in our own lives, but when we have arrived at an answer to a particular problem, to crystallize it into a not too specific or detailed form, so that its salient features may stand out and serve us again in a like situation without so much thought.

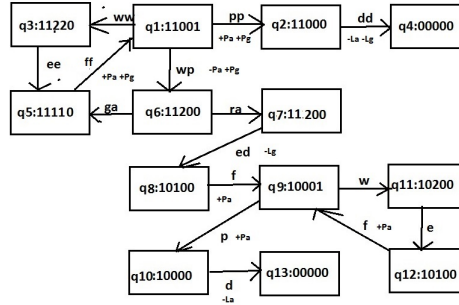
So while principles may serve well enough most of the time, there are situations where we need to think through the situation from scratch. The law, in effect, provides us with a set of ready made principles to guide our behavior, saving us the trouble of formulating our own principles. But no less than with moral principles, circumstances arise when obeying the law will be undesirable, and we need to think for ourselves. Often such circumstances will be captured within the law by exceptions, but it is impossible to envisage every case, and so the exceptions will not be exhaustive. We will therefore here consider any violation of the main norm or law, and not whether particular situations may be covered by exceptions. In this paper we will consider how we can give software agents the capacity to perform quasi-moral reasoning<sup>3</sup> so that they can determine whether they should violate a norm or law, and the form this violation should take.

---

<sup>3</sup>We say “quasi-moral” since software agents do not themselves have ethical status, or can be considered to share our values. In this paper we will see such agents as proxies for human beings in simulations or transactions, and so their values will be those of the human they are representing. Developing a set of values applicable to software agents would be the topic of another paper. To see that human values are not applicable to software agents consider the fact that their life is of little value, since they can be easily reproduced or replaced, they don’t feel

## 2. Problems With Current Treatments

There are two main approaches to enforcing normative behaviour in MAS: either by removing prohibited actions (e.g. [31]), often called *regulation*, or by including explicit rules expressing the norms, often accompanied by sanctions. Neither are entirely satisfactory. We illustrate our discussion with a model of the fable of the *Ant and the Grasshopper* previously used in [13]. The model takes the form of an Alternating Action-Based Transition (AATS) [31], augmented with value labels [4]. The transition system, in which the nodes represent the states the agent may reach and the actions it may use to move between them (in an AATS they are joint actions, one action for each relevant agent), is a typical ingredient of Multi Agent Systems (MAS): the value labelling provides the basis for moral reasoning.



**Figure 1.** AATS+V for the Example: w = work, p = play, a = ask, g = give, r = refuse, e = eat, f = feast d = die

In the fable the ant works throughout the summer, while the grasshopper sings and plays and generally indulges herself. When winter comes and the ant has a store of food and the grasshopper does not, the grasshopper asks the ant for help. The ant refuses and says the grasshopper should have foreseen this, and so the grasshopper starves. The model also can be used to represent the parable of the *Prodigal Son*: in the parable the father makes a different choice from the ant, welcomes the repentant prodigal back, and gives him food.

Using the first approach we would enforce the behaviour recommended by the fable by removing the transition from  $q_6$  to  $q_5$  or the behaviour of the parable by removing the transition from  $q_6$  to  $q_7$ . A real life example in which actions are made unavailable is erecting bollards to prevent vehicles from entering a park (to use the famous example of Hart [21]). What can be wrong with this approach? After all, we can *prove* that the undesirable situation will not be reached, either using model checking [16] or analytic methods. Thus we can prove that universal compliance with the norm will achieve the desired results. This may be so, so long as the situation envisaged in the model is in operation. But suppose some state not modelled arises: perhaps someone has a heart attack in the middle of the park and so it is essential for an ambulance to enter the park in order to save that person's life. Now the bollards will prevent the person from being saved, and

---

pleasure or pain, nor happiness nor sorrow, and have no experience of liberty or fraternity.

the object of the norm, i.e. the value that the norm is designed to serve, the safety of park users, will be demoted rather than promoted. While the norm is effective in an ideal world, we do not live in an ideal world, and in a sub-ideal world it is often the case that adhering to norms applicable to an ideal world will not lead to the most desirable of the results which remain possible<sup>4</sup>.

Similarly, principles may cease to prescribe the best course of action in unforeseen situations. The whole point of Asimov’s three laws as a fictional device is that following them may lead to outcomes that the principles were designed to avoid. While a set of principles may provide good guidance most of the time, it is not difficult to think of situations where following the principles will lead to undesirable results, and so they need to be disregarded. The problem is not improved by the existence of sanctions, and indeed may be made worse since the threat of possible punishment makes violation less attractive to the agent.

Thus while either of the approaches may be effective in closed systems (providing they are simple enough for a model covering every eventuality to be constructed), they cannot be sure to cope with the unexpected events and states that will arise in an open system, where not every possibility can be envisaged or modelled<sup>5</sup>. In such cases we may find that the very reasons which led to the adoption of a norm will require the agent to violate that very same norm. This is even more the case in everyday life than in open agent systems.

Irrespective of which option is chosen, the regulation of behaviours at the level of norms does not allow for agents to appropriately violate norms, even in cases where compliance with the normatively prescribed behaviours results in demotion of the values that these norms are designed “to serve”, or even the demotion of values preferred to those served by the norm. Hence, we argue that agents should be equipped with the capacity to reason about values, the extent to which normatively prescribed actions serve these values, which values are more important than other values (i.e. value orderings qua ‘audiences’), and the ability to derive these orderings from a variety of sources, including experience, the law, and stories prevalent in the culture. These capacities constitute moral reasoning from first principles; the kind of reasoning required to deal with new and unexpected situations in which blind compliance with norms may lead to undesirable outcomes. This paper serves as a call to further develop reasoning of this kind, building on a number of existing developments.

### 3. Value-Based Reasoning

A method for value-based reasoning was formalised using an AATS labelled with values in [4] and further articulated in [3], and which gave nine reasons for action in terms of the promotion and demotion of values. The basic idea is that the transitions which promote values form the basis of arguments for the action which will allow that transition to be followed, and that the transitions which demote values will supply arguments against actions which permit these transitions. Fur-

---

<sup>4</sup>This is known in economics as the *Theory of the Second Best* [23].

<sup>5</sup>As Wilde put it in *An Ideal Husband*: “To expect the unexpected shows a thoroughly modern intellect”.

ther arguments may come from assumptions about the current state and the state that will be reached by following a particular transition. These arguments and the attack relations between them (determined according to the so-called critical questions listed in [4]) define an argumentation framework [19]. Moreover since the arguments will be associated with values, the framework is a value-based argumentation framework (VAF) [7]. In a VAF, arguments are evaluated from the perspective of an *audience* (cf [29]) characterised as an ordering on values, and attacks which are unsuccessful for an audience are distinguished from those which succeed (*defeats*). The result is a set of arguments acceptable to a particular audience. If there are no cycles in a single value, this will be a non-empty unique set [7]. The relation between norms and value based reasoning is considered in [9].

If we consider the ant's choice in  $q_6$  of Figure 1, he may either refuse or give. Which is chosen will, using the labels of Figure 1, depend on whether the ant prefers his own pleasure to the life of the grasshopper. The application of value based reasoning to moral decisions was considered in [5], which suggested that moral acceptability required that one's own lesser values should not be more highly ranked than more important values relating to others. Similarly equality before the law is an important principle, so if the value of life is preferred it should not matter whose life is at issue. This would not (morally or legally) allow the preference of the ant's pleasure over the grasshopper's life, and so require the ant to give food to the grasshopper. But the labelling in Figure 1 is not the only one possible. If we think more abstractly we may see the ant's refusal as promoting *Justice*, since the grasshopper knew full well that food would be required in the winter and not working in the summer would mean later exploitation of the good nature of the ant. Similarly we could label the giving of the food as *compassion* or *mercy*. Preferring justice to mercy becomes more legitimate if we consider the role of the moral code to be producing a sustainable society, which requires that working in the Summer be seen as the norm. As shown in [25] the sustainability of norms requires that transgressions be punished, and that failures to punish are themselves punished. Thus punishing the grasshopper may be seen as the duty of the ant. Note too that in the parable, the prodigal is repentant, and so the father will only be expected to show compassion once. Representing such things as repentance will require an extension to the state descriptions, but will allow a preference for justice over compassion to be dependent on the misbehavior being repeated. Tolerance of limited misbehaviour before enforcing punishments is explored through simulation in [24].

Yet another way of describing the problem would be to recognise that the singing of the grasshopper may be a source of pleasure to the ant as well as to the grasshopper (many societies do support entertainers). Seen this way, we need change the action descriptions, so that the ant does not *give* but *pays* for services rendered. This in turn requires requires recognition that it is the duty of the ant to pay for the services received and so justice is now promoted by following the transition from  $q_6$  to  $q_5$ , not  $q_7$ . Moreover since a single grasshopper may entertain a whole colony of ants, the burden falling on a single ant may be relatively small.

If, however, there is only a single ant, suppose that the harvest fails, and there is no surplus to pay the grasshopper. Should the ant follow the norm, pay the grasshopper and starve or renege on the agreement and watch the grasshopper

starve? Here we will have a genuine moral dilemma, in which the ant must choose between justice and its life. The ant may choose death before dishonour, but may also choose to renege with good authority. Thomas Aquinas writes:

if the need be so manifest and urgent, that it is evident that the present need must be remedied by whatever means be at hand (for instance when a person is in some imminent danger, and there is no other possible remedy), then it is lawful for a man to succor his own need by means of another's property, by taking it either openly or secretly: nor is this properly speaking theft or robbery.<sup>6</sup> [1], Question 66, Article 6. See also [18].

Thus the ant has a choice, and either can be justified. What the ant will do will depend on its value preferences. Arguably the original contract was foolhardy - on the part of both - since the possible failure of the harvest could have been foreseen by both parties, and whichever suffers has only themselves to blame. When designing legal systems the legislators have a similar choice, and the choice made reflects the values of the society which adopts the laws. Very often the law has not recognised necessity as defence: examples include anti-poaching laws, and transportation of people stealing food in Victorian England.

#### 4. What Makes a Moral Audience?

As the last example shows, there may be more than one morally acceptable ordering on values. Some other orderings, such as a refusal to pay an entertainer even when there a surplus available to do so, are not acceptable. What we must do is to provide our agents with an acceptable ordering on which to base their reasoning. In order to do so, we need to look at the value order prevailing in society. As noted above, the decisions made by courts often manifest an ordering on values. Case law decisions often turn on the value preferences the judge wishes to express. This use of social purposes to justify judicial decisions was introduced to AI and Law in [12] and more formally presented in [11]. Thus we may look to the law as one source for our value orderings: the assumption being that the moral order is at least compatible with the order reflected in legal decisions. Note that this legal order need not be static and may reflect changing social views and priorities. Although courts are supposed to be bound by precedents (the doctrine of *stare decisis*) as noted by Mr Justice Marshall in the US Supreme Court case of *Furman v Georgia* (408 U.S. 238 1972) there are occasions when “*stare decisis* would bow to changing values”.

Several methods of deriving an audience, in the sense of a value ordering, from a set of cases have been proposed. In AGATHA [17] the value ordering which best explains a set of cases was discovered by forming a theory to explain a set of cases, and then attempting to provide a better theory, in terms of explaining more cases, until the best available theory was found. In [10], given a VAF and a set of arguments to be accepted, the audiences (if any) to which that set is acceptable is determined by means of a dialogue game. The ordering may not be fully determined (a *specific* audience): it may be possible for the desired set of

---

<sup>6</sup>This would, of course, also justify the grasshopper stealing from the ant.

arguments to be accepted by several audiences, represented as a partial order on the values. In [27], the VAF is rewritten as a meta-level argumentation framework [28], from which value orderings can emerge, or be formed, as a result of dialogue games based on the rewritten frameworks. In [27] explicit arguments for value orderings can be made in the manner of [26].

As well as legal cases, we can identify the approved value orderings from stories, using techniques for deriving character motives from choices with respect to actions, originally targetted at explaining the actions of people involved in legal cases [15]. Stories are often used to persuade people to adopt particular value orders, as with the fable and the parable we have considered in this paper. The notion of using didactic stories as arguments for value orderings was explored in [14] and [13]. Since stories like fables and parables were written specifically to advocate particular value orderings, they are highly suited to our purposes. The values concerned are typically clear, the choices sharp and the correct decisions clearly signposted, leaving little room for doubt as to the recommended preference.

We do not propose data mining or machine learning methods here. Although such methods can discover norms from a set of cases represented as facts and outcomes (e.g [30]), the discovered norms derive their authority from the amount of support in the dataset. They are suited to finding rules, but not exceptions, and it is exceptional cases, where norms need to be violated, that interest us. In law single cases may form important precedents, identifying apparent exceptions to existing norms, closing gaps and resolving conflicts, often revealing or making a choice between value orderings.

As noted above, these methods may produce not a specific audience, but a set of audiences all of which conform to and explain the prevailing decisions. If this is so the question arises as to whether it is desirable or undesirable for all agents to be drawn from the same audience. To unify the audience would be to impose the designer's view as to what is moral, albeit constrained by the social decisions. In practice a degree of diversity may prove useful, leading to different agents occupying different social roles.

## 5. Road Traffic Example

This section considers an area where violation of the general norm is very common: the law that drivers should drive on the left (in the UK, on the right in many other countries). The law is intended to avoid collisions, and so promote the values of *Progress* and *Safety*. But on every journey, it is necessary to violate this law if progress is to be maintained: obstructions such as parked cars and roadworks, the need to overtake slow moving vehicles and cyclists and emergencies such as a pedestrian or animal stepping in from of the car, may all lead drivers to move to the right. But the question remains: when is it permissible to do so?

Space precludes giving the full AATS, but we will give a sufficient fragment to allow the consideration of some relevant situations. Our AATS will represent relevant features of the states that can occur. For our example we consider:

- Whether there is an obstruction and whether it is moving or stationary (0 = clear, 1 = slow moving, 2 = stationary).

**Table 1.** Joint actions for self and on-coming in state 2100

self: on-coming	continue	slow	stop	change lane	mount pavement
continue	J1	J2	J3	J4	J5
slow	J6	J7	J8	J9	J10
stop	J11	J12	J13	J14	J15
change lane	J16	J17	J18	J19	J20
mount pavement	J21	J22	J23	J24	J25

- Whether there is an on-coming vehicle and whether it can stop or not (0 = no on-coming, 1 = can stop, 2 = cannot stop).
- Whether our own vehicle can stop safely (0= can stop, 1 = cannot stop).
- Whether there is a collision (0 = no collision, 1(x,y) = x collides with y).

For actions, both our own vehicle and the on-coming will be able to continue, change lane, stop, slow, or mount the pavement. For values we consider our own progress and safety (P(s) and S(s)), the progress and safety of the on-coming (P(o) and S(o)) and the safety of any pedestrians in the area (S(p)).

Now consider the transitions from the state where there is a stationary obstacle, and both ourselves and the on-coming could stop, and there is no collision (2100). In this case there are a number of joint actions involving self and on-coming as shown in Table 1. Additionally if one or both mount the pavement pedestrians may or may not be hit, depending on whether we get lucky or unlucky (often represented by including *Nature* in the joint action).

In the actions J1-J15 self obeys the norm, whereas J16-J20 and J21-J25 represent different ways of violating the norm. J1-J10 all result in self colliding with the obstacle, which demotes both P(s) and S(s). J11-15 do not demote S(s) but do demote P(s). Thus complying with the law will demote one or both of the values the norm was designed to promote, (although it does allow the on-coming to continue without threat to any of its values). We should therefore consider violating the norm. Suppose we go on to the right hand lane. Now J16 and J17 result in a head-on collision, which demotes all of P(s), P(o), S(s) and S(o). J18 demotes P(o) and J19 demotes both P(o) and S(o). J20 may or not demote S(p) for a variable number of pedestrians. Similarly J21-J25 will jeopardise the safety of an unknown number of pedestrians. We can therefore make a choice. If our desire for progress is insufficient to lead us to risk our safety (and the safety of others) we have to stop. If, however, we are sufficiently reckless that our desire for progress is such that we are willing to risk a collision we should change lane and hope that J18 is achieved, so that while P(o) is demoted, the threat to safety is avoided. This relies on the (normally acceptable) assumption that the on-coming agent will be less reckless than we are. J20-J25 are possible if we don't trust the on-coming to stop, but this poses the risk of an even more serious accident if we hit pedestrians. At this point we could either construct arguments for the other agents involved acting in certain ways (*the on-coming driver can't be as reckless as I am*, or *there will not be any pedestrians at this time of night*) in the manner of [4], or perform an expected value calculation as recommended in [6]. Here most of us will choose to obey the norm. But if there is no on-coming, then we can change lane and violate the norm with no risk to safety. This will be better both than obeying the law or mounting the pavement, however unlikely we consider it



to be that pedestrians are present, telling us both to violate the norm and how to violate it.

## 6. Summary

As the use of agents spreads and as they adopt the autonomous performance of ever more critical tasks, there is a need to provide them with the capacity for moral reasoning. Neither of the approaches popular in current multi-agent systems is entirely suitable. Moral behaviour requires and includes the recognition that on occasion it is right to violate norms, because while norms may be best observed in an ideal world, we need to be able to cope with the sub-ideal, and with the unforeseen. Unforeseen events may occur which mean that following a norm results in undesirable effects, perhaps even subverting the very values the norm was designed to promote. Moreover when another agent transgresses norms, so producing a sub-ideal situation, it may be necessary to deviate oneself, either to punish the transgression or because the case is altered, and in the particular circumstances two wrongs *do* make a right.

But violation of a norm for moral reasons presupposes that the agent be able to recognise when the norm should be violated and what form the violation should take. This in turn requires that the agent be able to reason morally without reference to norms, which requires the agent not only to apply an ordering on values to the current situation, but also to *reason about* the relationship between actions and the different values they promote or demote, the relative importance of values (i.e. value orderings qua ‘audiences’), and to have the ability to derive these orderings from a variety of sources, including experience, the law, and stories prevalent in the culture<sup>7</sup>. If we provide agents with these capacities, we can rely on the agents to make moral choices which might not be the case if they were to blindly follow a fixed set of norms. We have identified value based argumentation [7] as a starting point for formalising such capacities. However, to formalise reasoning *about* values and their orderings will require further investigation into the use of extended argumentation frameworks [26] and (more generally) metalevel argumentation [27,28]. Future work will also need to investigate how value orderings may be induced, how value orderings can be applied to situations that differ (in some tangible way that suggests novelty) from the ones that originally gave rise to them, and whether and how the kinds of reasoning from first principles described in this paper may be sufficiently persuasive so as to instigate modifications to norms (in the form of exceptions).

## References

- [1] Thomas Aquinas. *Summa theologica*. Authentic Media Inc, 2012, written 1265-74.

---

<sup>7</sup>In so doing we provide a morality in the virtue ethics tradition of Aristotle [2], as opposed to the consequentialism and deontology represented by current MAS approaches. We take an agent-centric view in which agents act in accordance with value orderings that encode what they consider virtuous. Note that to replace all reference to value-promotion/demotion by reference to utility-maximisation/reduction, would be to make the impractical assumption that agent’s explicitly apply a ‘view from the angels’ utilitarian calculus in their moral decision making.

- [2] Aristotle. *The Nicomachean Ethics of Aristotle*. Heinemann, 1962, written 350BC.
- [3] K. Atkinson and T. Bench-Capon. Taking the long view: Looking ahead in practical reasoning. In *Proceedings of COMMA 2014*, pages 109–120.
- [4] K. Atkinson and T. Bench-Capon. Practical reasoning as presumptive argumentation using action based alternating transition systems. *Art. Intell.*, 171(10):855–874, 2007.
- [5] K. Atkinson and T. Bench-Capon. Addressing moral problems through practical reasoning. *Journal of Applied Logic*, 6(2):135–151, 2008.
- [6] K. Atkinson and T. Bench-Capon. Value based reasoning and the actions of others. In *22nd European Conference on Artificial Intelligence*, pages 680–688, 2016.
- [7] T. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.
- [8] T. Bench-Capon. Transition systems for designing and reasoning about norms. *AI and Law*, 23(4):345–366, 2015.
- [9] T. Bench-Capon. Value-based reasoning and norms. In *22nd European Conference on Artificial Intelligence*, pages 1664–1665, 2016.
- [10] T. Bench-Capon, S. Doutre, and P. Dunne. Audiences in argumentation frameworks. *Artificial Intelligence*, 171(1):42–71, 2007.
- [11] T. Bench-Capon and G. Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1):97–143, 2003.
- [12] D. Berman and C. Hafner. Representing teleological structure in case-based legal reasoning: the missing link. In *Proceedings of the 4th ICAIL*, pages 50–59. ACM, 1993.
- [13] F. Bex, K. Atkinson, and T. Bench-Capon. Arguments as a new perspective on character motive in stories. *Literary and Linguistic Computing*, 29(4):467–487, 2014.
- [14] F. Bex and T. Bench-Capon. Understanding narratives with argumentation. In *Proceedings of COMMA 2014*, pages 11–18, 2014.
- [15] F. Bex, T. Bench-Capon, and K. Atkinson. Did he jump or was he pushed? *AI and Law*, 17(2):79–99, 2009.
- [16] D. Bošnački and D. Dams. Discrete-time promela and spin. In *Formal Techniques in Real-Time and Fault-Tolerant Systems*, pages 307–310. Springer, 1998.
- [17] A. Chorley and T. Bench-Capon. An empirical investigation of reasoning with legal cases through theory construction and application. *AI and Law*, 13(3-4):323–371, 2005.
- [18] G. Christie. *The notion of an ideal audience in legal argument*. Springer, 2012.
- [19] P.M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Art. Intell.*, 77(2):321–357, 1995.
- [20] R. Hare. *Freedom and reason*. Oxford Paperbacks, 1965.
- [21] H. Hart. *The concept of law*. OUP Oxford, 2012.
- [22] M. Jones, A. and Sergot. Deontic logic in the representation of law: Towards a methodology. *AI and Law*, 1(1):45–64, 1992.
- [23] R. Lipsey and K. Lancaster. The general theory of second best. *The review of economic studies*, 24(1):11–32, 1956.
- [24] M. Lloyd-Kelly, K. Atkinson, and T. Bench-Capon. Emotion as an enabler of co-operation. In *ICAART (2)*, pages 164–169, 2012.
- [25] S. Mahmoud, N. Griffiths, J. Keppens, A. Taweel, and M. Bench-Capon, T. and Luck. Establishing norms with metanorms in distributed computational systems. *AI and Law*, 23(4):367–407, 2015.
- [26] S. Modgil. Reasoning about preferences in argumentation frameworks. *Artificial Intelligence*, 173(9):901–934, 2009.
- [27] S. Modgil and T. Bench-Capon. Integrating object and meta-level value based argumentation. In *Proceedings of COMMA 2008*, pages 240–251, 2008.
- [28] S. Modgil and T. Bench-Capon. Metalevel argumentation. *Journal of Logic and Computation*, pages 959–1003, 2010.
- [29] Ch. Perelman. *The new rhetoric*. Springer, 1971.
- [30] M. Wardeh, T. Bench-Capon, and F. Coenen. Padua: a protocol for argumentation dialogue using association rules. *AI and Law*, 17(3):183–215, 2009.
- [31] M. Wooldridge and W. van der Hoek. On obligations and normative ability: Towards a logical analysis of the social contract. *J. Applied Logic*, 3(3-4):396–420, 2005.