Value-Based Reasoning and Norms

Trevor Bench-Capon [1]

Abstract. In this paper we explore how value-based practical reasoning relates to norms and their evolution. Starting from a basic model of a society and the norms that can arise from it, we consider how additional values, extended state descriptions and finer grained descriptions of actions, can lead to more complex norms, and a correspondingly more complex social order.

1 Introduction

Norms are a topic of considerable interest in agents systems [60], [51], [62], [58], [50], [49], [3], [54], [39]. In particular, in open agent systems, it is not possible to assume that all agents will behave according to the same ethical code, and the open nature of the system means that the designer cannot simply impose norms that can be assumed to be followed by all. Of course, it is possible to construct so-called *regulated* systems, where the agent can only perform permissible actions (e.g.[28], [58], [2]). However, since, unlike norms found in legal and moral systems, such norms cannot be violated, it can be argued that (e.g. [35], [30]) they should not be seen as norms all, because the agents have no choice beyond compliance or non-participation. Such rules are thus like the rules of a game, not moral and legal norms.

An excellent starting point for considering the emergence of norms is [57], which does, of course, considerably pre-date multi agent systems, but none the less contains many relevant considerations. In that work, Ullmann-Margalit uses simple two player games, such as the prisoner's dilemma (PD) [46], to explore the topic. In such games there are two players and each can cooperate or defect, and the choices determine the payoffs. In PD as used in [57] mutual cooperation gives a payoff of 3 to each player and mutual defection 1 to each player, while if the actions differ the defector receives 5 and the cooperator receives zero. Some key results concerning PD are that the Nash Equilibrium [48] is where both defect (since defection is the *dominant* action, and will receive the better payoff whatever the other player does) and that a successful strategy in iterated PD (where the players play one another repeatedly) is *Tit-Fot-Tat* [11] (but see [18]). Using *Tit-Fot-Tat* an agent will cooperate in the first round, and then copy its opponent's previous move in every subsequent round. Importantly PD is non-zero sum game: the aggregate utility of mutual cooperation is greater than any other payoff, and the equilibrium in fact yields the lowest collective utility. Thus, if would in fact be mutually beneficial if one offered a payment to the other if they cooperated: this could secure payoff of 3 and 2, so that both would gain over mutual defection. Such agreements are, however, not possible in the game, which does not allow for prior negotiations.

Public goods game have formed the basis of several studies of the emergence of norms in multi-agent systems such as [51], [50], [53], [17], [54] and [39]. An alternative approach is to model a situation as a State Transition Diagram (STD), and to investigate how norms can emerge from agent interaction in such situations [62], [3]. In these latter models, agents are typically represented using the Belief-Desire-Intention (BDI) model [45], [61], inspired by [20]. The BDI model supposes agents to have a set of *beliefs* and a set of dispositional goals (*desires*). Actions are chosen by identifying the desires than can be realised in the current situation (candidate *intentions*), and then committing to one or more of these intentions, and choosing a course of action intended to realise the associated goals. This, however, leaves open the question of where the desires come from in the first place.

Empirical studies suggest, however, that public goods games do not provide a very realistic model of actual human behaviour. Experiments using the public goods games are very common and have formed the subject of metastudies. For example [27] examined 131 examples of the Dictator Game and [42] was based on 37 papers reporting Ultimatum Game experiments. In none of these many studies was the canonical model followed. Although the metastudy of [33] was smaller, looking at only 15 studies, it is particularly interesting in that the studies considered highly homogeneous societies. In BDI systems, there is no explanation of where goals come from. Often they are completely fixed, and even systems where they can be derived from the current state [44], there is a fixed set of potential desires some of which are active in a given situation.

An alternative approach to action selection (often called practical reasoning [47]) is provided by Value-Based Reasoning, in which agents are associated with a set of social *values*, the aspirations or the purposes an agent might pursue, such as liberty, equality, fraternity, wealth, health and happiness, and these values provide reasons why certain situations are considered goals by the agent. The basic idea is that agents have a set of such values and their aspirations and preferences are characterised by their ordering of these social values. Acceptance of an argument as to what to do depends not only on the argument itself - for it must, of course, be a sound argument - but also on the audience to which it is addressed [43]. This notion of audience as an ordering on values was computationally modelled in [31] and made more formal in Value-Based Argumentation Frameworks (VAFs) [12]. VAFs are an extension of the abstract Argumentation Frameworks (AFs) introduced in [24], but whereas in an AF an argument is defeated by any attacking argument, in a VAF an argument is *defeated for an audience* by an attacker only if the value associated with the attacking argument is ranked at least as highly as the attacked argument by that audience. In this way different audiences will accept different sets of arguments (preferred semantics [24] is used to determine acceptance), and, as is shown in [12], provided the VAF contains no cycles in the same value, there will be a unique non-empty preferred extension.

Use of VAFs provides a way of explaining (and computing) the

[1] Department of Computer Science, University of Liverpool, UK. email: tbc@csc.liv.ac.uk

different arguments accepted by different audiences. Value-Based Reasoning been used as the basis of practical reasoning in, amongst others, [29], [6], and [59], and applied in particular areas including law [13], e-democracy [22], policy analysis [55], medicine, [9], experimental economics [14], and rule compliance [21]. Complexity results for VAFs were established in [25] and [41]. Here we will discuss norms and their evolution in terms of the Value-Based approach to practical reasoning.

2 Background

In this section we provide some essential background: the structure with which we use to model our "world", Alternate Action Based Transition Systems (AATS), and the valued-based arguments that agents can use to justify their actions in this environment; the running example we will use to instantiate our model; and the three types of ethical theory we will consider.

2.1 Alternate Action Based Transition Systems

Based on Alternating Time Temporal Logic [4], AATS were originally presented in [62] as semantical structures for modelling game-like, dynamic, multi-agent systems in which the agents can perform actions in order to modify and attempt to control the system in some way. As such they provide an excellent basis for modelling situations in which a set of agents are required to make decisions.

The definition in [62] is:

Definition 1: AATS. An *Action-based Alternating Transition System* (AATS) is an $(n + 7)$-tuple $S = \langle Q, q_0, Ag, Ac_1, \ldots, Ac_n, \rho, \tau, \Phi, \pi \rangle$, where:

- Q is a finite, non-empty set of *states*;
- $q_0 \in Q$ is the *initial state*;
- $Ag = \{1,...,n\}$ is a finite, non-empty set of *agents*;
- Ac_i is a finite, non-empty set of actions, for each $ag_i \in Ag$ where $Ac_i \cap Ac_j = \emptyset$ for all $ag_i \neq ag_j \in Ag$;
- $\rho : Ac_{ag} \rightarrow 2^Q$ is an *action pre-condition function*, which for each action $\alpha \in Ac_{ag}$ defines the set of states $\rho(\alpha)$ from which α may be executed;
- $\tau : Q \times J_{Ag} \rightarrow Q$ is a partial *system transition function*, which defines the state $\tau(q, j)$ that would result by the performance of j from state q. This function is partial as not all joint actions are possible in all states;
- Φ is a finite, non-empty set of *atomic propositions*; and
- $\pi : Q \rightarrow 2^\Phi$ is an interpretation function, which gives the set of primitive propositions satisfied in each state: if $p \in \pi(q)$, then this means that the propositional variable p is satisfied (equivalently, true) in state q.

AATSs are particularly concerned with the joint actions of the set of agents Ag, J_{Ag}: j_{Ag} is the joint action of the set of n agents that make up Ag, and is a tuple $\langle \alpha_1,...,\alpha_n \rangle$, where for each α_j (where $j \leq n$) there is some $ag_i \in Ag$ such that $\alpha_j \in Ac_i$. Moreover, there are no two different actions α_j and $\alpha_{j'}$ in j_{Ag} that belong to the same Ac_i. The set of all joint actions for the set of agents Ag is denoted by J_{Ag}, so $J_{Ag} = \prod_{i \in Ag} Ac_i$. Given an element j of J_{Ag} and an agent $ag_i \in Ag$, ag_i's action in j is denoted by j^{ag_i}. This definition was extended in [6] to allow the transitions to be labelled with the values they promote.

Definition 2: AATS+V. Given an AATS, an AATS+V is defined by adding two additional elements as follows:

- V is a finite, non-empty set of values.
- $\delta : Q \times Q \times V \rightarrow \{+, -, =\}$ is a *valuation function* which defines the status (promoted (+), demoted (–) or neutral (=)) of a value $v_u \in V$ ascribed to the transition between two states: $\delta(q_x, q_y, v_u)$ labels the transition between q_x and q_y with one of $\{+, -, =\}$ with respect to the value $v_u \in V$.

An *Action-based Alternating Transition System with Values* (AATS+V) is thus defined as a $(n + 9)$ tuple $S = \langle Q, q_0, Ag, Ac_1, ..., Ac_n, \rho, \tau, \Phi, \pi, V, \delta \rangle$. The value may be ascribed on the basis of the source and target states, or in virtue of an action in the joint action, where performing that action itself promotes or demotes a value.

2.2 Reasons for Action

The values give agents reasons to perform or not to perform the various actions, based on the argumentation scheme proposed in [6]. A number of such reasons are given in [8] (the "N" suffix denotes reasons not to perform the action: ϕ is a *goal*, which holds or fails to hold in a given state, and which agents may attempt to realise, maintain, avoid or remove).

R1 We should participate in j in q in which ϕ holds to maintain ϕ and so promote v.
R2N We should not participate in j in q in which ϕ holds since it would remove ϕ and so demote v.
R3 We should participate in j in q in which $\neg\phi$ holds to achieve ϕ and so promote v.
R4N We should not participate in j in q in which $\neg\phi$ holds since it would avoid ϕ and so fail to promote v.
R5 We should participate in j in q to ensure ϕ and so promote v. Note that ϕ may be contingently realised or unrealised in q and that, in some variants, the promotion of v might not be immediate, or permanent. This also applies to R5N and R6.
R5N We should not participate in j in q which would ensure $\neg\phi$ and so demote v.
R6 We should participate in j in q to prevent $\neg\phi$ and so promote v. Note that $\neg\phi$ may be contingently realised or unrealised in q.
R6N We should not participate in j in q which would prevent ϕ and so fail to promote v. We suggest that to make the reason worth consideration we should only use variants which prevent ϕ immediately and permanently.
R7 We should participate in j in q in which $\neg\phi$ to enable ϕ to be achieved and v to be promoted on the next move.
R8N We should not participate in j in q in which ϕ which will risk ϕ being removed on the next move which would demote v.
R9 We should participate in j in q because performing j^{ag} promotes v.
R9N We should not participate in j in q because performing j^{ag} demotes v.

Objections to these arguments can be formed by questioning whether the state is as claimed, the consequences of the action will be as specified, whether the goal is realised and whether the value is indeed promoted. The arguments and attacks are then organised in a Value-Based Argumentation framework (VAF) [12] and evaluated according to an ordering on the values. These value orderings will depend on the subjective preferences of the particular audience, and so different agents may choose different actions.

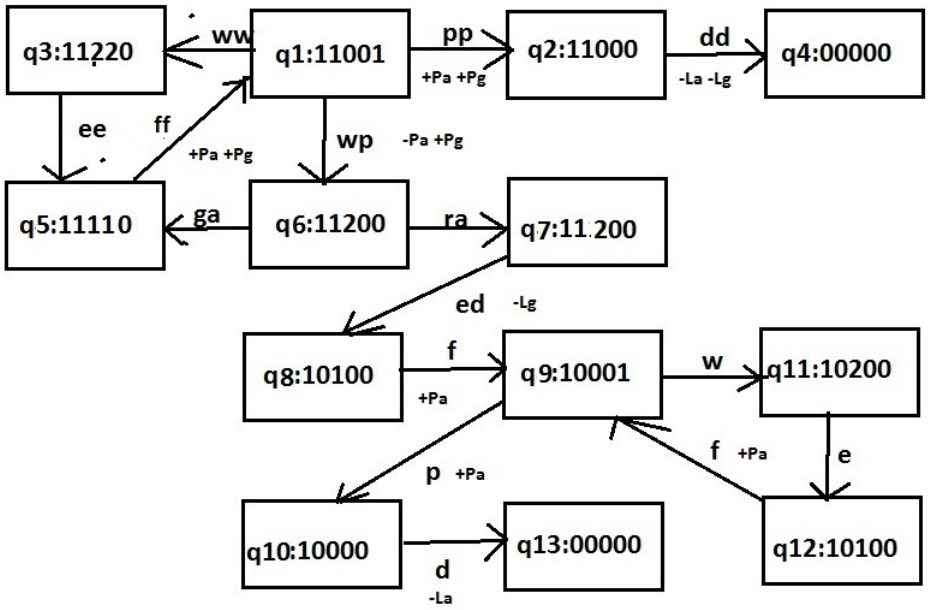

Figure 1. AATS+V for the Example: w = work, p = play, a = ask, g = give, r =refuse, e = eat, f = feast d =die. The same AATS+V is used for both the fable and the parable. Joint actions are ant/father, grasshopper/son. States are: ant/father alive, grasshopper/son alive, ant/father has food. grasshopper/son has food, summer/winter

2.3 Example

An AATS+V was used in [16] to model the states and actions found in both the fable of *the Ant and the Grasshopper* [1] and the parable of *the Prodigal Son* (Luke 15:11-32). Fables and parables are suitable examples for us because they are stories with a moral point. In *the Ant and the Grasshopper*, the story is that during the summer the grasshopper sings and plays while the ant works hard storing up food for the winter. When the winter comes, the grasshopper has no food: nor will the ant give away any of its store, and so the grasshopper dies. In *the Prodigal Son* the prodigal wastes his inheritance on idle play but when destitute asks his father for forgiveness: the father does forgive and takes him back into his household.

An AATS based on the model of [16] is shown in Figure 1. In our example, food is sufficiently abundant in Summer that one can gather food and eat without effort. Growing food for the winter is, however, a full time effort (digging, planting, weeding, reaping, storing) and produces a surplus, but the nature of the activity is that it is either done or not: the amount produced in not proportional to the effort. The food does not last into the summer: therefore the winter ends with a period of carnival (q5, q8 and q12) when the surplus is consumed with feasting. The state has five propositions. The first two indicate whether the ant (father) and the grasshopper (son) are alive. The third and the fourth whether the ant (father) and the grasshopper (son) have no, enough or abundant food, and the fifth whether it is summer or winter. The key decisions are in the initial state (q1) where both grasshopper and prodigal choose to play rather than work and in q6 where the ant refuses the grasshopper (action *r*) while the father gives to the prodigal (action *g*). In the other states there are no choices to be made.

We have labelled the diagram in Figure 1 with just four values. Life for the ant (father) and grasshopper (son) (L_a and L_g) and Pleasure for the ant (father) and the grasshopper (son) (P_a and P_g).

2.4 Ethical Theories

Broadly, as a considerable simplification, ethical theories can be divided into three types:

Consequentialism: An action is right if it promotes the best consequences. For example, Mill's *Utilitarianism* [40].
Deontology An action is right if it is in accordance with a moral rule or principle. For example, Kant [36]
Virtue Ethics: An action is right if it is what a virtuous agent would do in the circumstances. For example, Aristotle [5]

3 Developing a Moral Code

In this section we consider how a moral code might develop from a consideration of value-based practical reasoning in the example scenario.

3.1 Arguments in q_1

We now consider the arguments available to an agent in q_1, based on the values of pleasure and life. The agent's own pleasure and life will be denoted P_s and L_s, the pleasure and life of the other as P_o and L_o. Our arguments are derived from the reasons of section 2.2, expressed in terms of only the agent's own action and the value, e.g. *you should perform* α *since it will promote* v, where α is the agent's action in the justified joint action, and v is the value promoted.

A You should not play since it will risk L_s being demoted (R4N)
B You should work since it will enable P_s to be promoted (R7)
C You should play to promote P_s (R9)
D You should not work since it will demote P_s(R9N)

Thus we have reasons pro and con working: the pro reason is the future pleasure it enables, and the con reason is the immediate loss

of pleasure which accompanies it. Play in contrast affords immediate pleasure, but risks the loss of life. The risk associated with argument A is substantial: avoiding death requires both that the other works, and that the other will forgo its own pleasure in order to save one's life. Therefore (assuming life is preferred to pleasure) only the most risk taking individuals will choose to play in q_1.

Viewed from the perspective of the three moral theories:

- Consequentialism will make work obligatory (or forbid play), to avoid both the undesired consequence of being in q_2 with its unavoidable successor q_4, and the normative collapse that will result from the encouragement given to free loading if the idler is fed [39].
- Deontology will also make work obligatory (or forbid play), since it is not possible to will that both agents play.
- Virtue Ethics will require that life is preferred to pleasure, so that the virtuous agent will choose to work..

All three of these theories will support the continuing existence of the community. We believe a moral code should be *sustainable*, in the sense of ensuring the continuance of the community and avoiding the collapse of its norms [39].

3.2 A first set of moral norms

So let us suppose that the society has the moral norm:

MN1: It is forbidden to play

If all agents act morally the situation can continue indefinitely round the loop q_1, q_3, q_5, q_1. But there will always be the temptation to violate MN1: the value L_s is only threatened, and if q_6 is reached, there is the possibility that that the other agent will give the required food. Work on norms such as that reported in [10] and [39] suggests that unless there is some reinforcement mechanism, norms are liable to break down. The reinforcement mechanism is for violations to be punished by other members of the group, which in q_6 would mean that food is withheld. Moreover, to avoid the norm collapse [39], it is necessary to punish those who do not punish, and so punishment needs to be seen as obligatory. This in turn means that we need a moral norm applicable in q_6:

MN2: It is forbidden to give food

Now refusal to give food would also be justified by an argument based on R6N, *you should not give since that will fail to promote* P_s. Given the counterargument based on R5N, you should not refuse since this will demote L_o, this requires a preference for P_s over L_o. But this does seems selfish rather than moral, and acts against sustainability. It does not seem morally good to prefer a minor value in respect of oneself to an important value in respect of the other: in [7], for example, acting morally was characterised as not regarding lesser values enjoyed by oneself as preferable to more important values in respect of others, which would oblige the agent to give. We can, however, introduce another value, Justice (J), which will justify refusal. This has some intuitive appeal, since the foodless agent has chosen to be in that position, and is attempting to take advantage of the efforts of the other. Thus recognising justice as a third value (labelling the transition q_6-q_7), preferred to L_o, will justify the punishment of violations of MN1. This would be difficult to justify under a consequentialist perspective (since it means the grasshopper dies), but is capable of universal adoption, and it is not difficult to see a preference for justice as virtuous, since it can be justified in terms of equity and sustainability, by preventing the collapse of MN1.

The result will be a pair of norms which are capable of resisting collapse, according to the empirical findings of [39]. The result is a rather puritan society (relieved only by a brief period of hedonism), based on subsistence farming, with a strong work ethic, and an aversion to what in the UK is currently termed a "something for nothing" society. An alternative would be to introduce a fourth value, Mercy, preferred to Justice, and labelling the transition $q6$-q_5. Ranking Mercy above Justice is very possibly the recommendation of the parable of *The Prodigal Son*, and would also allow society to continue, at the expense of a sacrifice of pleasure by the ant. But it is a feature of the parable that the son repents, and there is a tacit understanding that the son will not repeat the pattern, but will choose work in future. We might therefore wish to modify MN2 to something like

- **MN2a** It is allowed to give food only once. We might wish to go further and to accompany this with
- **MN2b** It is obligatory to meet the first request for food

This would represent a preference for Mercy, but enforce *a two strikes and you are out* policy, so that justice is still respected. It also opens the possibility for the ant to play at the grasshopper's expense on some future cycle (cf. children supporting elderly parents). Whereas simply removing MN2 would lead to the possibility of exploitation, and so devalue Justice, the pair of norms MN2a and MN2b retain some respect for Justice, while allowing scope for Mercy until the object of the mercy proves incorrigible. This will require the state vector to have an extra term to record violations.

Our developments beyond the basic scenario of Figure 1 will necessarily be less detailed, both because of space limitations and because of the large number of possible variations. Of course it would, in future, be desirable to extend the scenario at the same level of detail and provide an expanded AATS+V, but we hope that it is clear that the discussion given below in the following sections is making use of the same techniques.

3.3 Critique

Although the norms NM1 and NM2 will give rise to an equitable and sustainable society, we might expect to see thinkers in such a society as questioning the worth of the society. There might be a number of grounds for critiques. For example:

- There is no net pleasure in the society: the displeasure of working is offset by the pleasures of feasting at carnival, but there is no net gain. Such a society lacks progress and reward and any point beyond it own continuance.
- There is no choice or diversity in the society: the path taken is determined at all times by the moral code.
- The pleasure enjoyed in this society is of a rather basic kind, whereas the pleasure it denies itself might be seen as an higher pleasure. The hard line utilitarian might adopt Bentham's view [15] that "Prejudice apart, the game of push-pin is of equal value with the arts and sciences of music and poetry", but others, like Mill would disagree: "it is better to be a human being dissatisfied than a pig satisfied; better to be Socrates dissatisfied than a fool satisfied" [40]. Such higher pleasures can only be provided by (certain forms of) play, not by feasting.

Therefore critics might seek a way of changing the moral code so as to improve the society in one or more of these respects. Certainly, it is considered essential to a civilised society that it is able to generate a surplus of food and so allow scope for the development

of arts and sciences, or even the simple enjoyment of leisure time. There is therefore some push behind finding a justification for allowing some relaxation of the rigid morality represented by MN1 and MN2. In order to further the discussion we will distinguish between three types of pleasure, by introducing different values for different pleasurable activities. We will retain P for the bodily pleasures associated with carnival (feasting and the like): toil will also continue to demote this value. We will also distinguish between approved activities made possible by not working (e.g. arts and sciences) which we will term *culture* (C), and deprecated activities (e.g. gaming or mere idleness) which we will term *frivolity* (F). We thus need to distinguish between approved play ($play_a$) (i.e engagement in culture producing activities) and deprecated play ($play_d$) (i.e. engaging in frivolity). We might even modify MN2b to give food only to someone in need because of $play_a$, and to withhold food from those in need as a result of $play_d$.

4 Allowing For Play

There are a number of ways in which we can accommodate players. Some require disparity between agents, while other require a redescription of the world, additional values and a finer grained description of activities and states.

4.1 Power

We first consider a disparity of power. In this situation some agents are sufficiently more powerful than the others to be able to compel them to surrender their food. We assume that the powerful agents comprise less than half the population. This is modelled by allowing the powerful agents to *demand*, rather than *request*, food in q_6, and to render it impossible to refuse a demand, so that there is no *rd* transition between q_6 and q_7. This removes argument A for the powerful, since there is no longer any risk in playing because the demands must be acceded to. Might the powerful play and demand all the food from the ant so that they can also feast? This would result in the ant starving and so would be a short term expedient, since the ant would die and the powerful be forced to work in subsequent years. So we should perhaps have a norm to prevent this:

MN3 It is forbidden to demand non-surplus food.

This can be based on a value preference for the Life of the other over Pleasure.

Remember now that we have distinguished between three types of pleasure so that argument C needs to be split into two arguments:

- **C1**: You should $play_a$ to promote culture (C).
- **C2** : You should $play_d$ to promote frivolity (F).

Now the powerful will not choose to work unless they prefer P to both C and F. They also have a choice of leisure activity, depending on whether they prefer culture of frivolity. Of course, this moral preference is built into the names of the values, and the moral norm, *applicable only to the powerful*, will be

MN4 It is forbidden to $play_d$.

This norm allows the choice to work to be morally acceptable. MN4 is, like Bentham, comfortable with a preference for pleasure over culture. Alternatively we can represent Mill's position with

MN4a It is obligatory to $play_a$

(also directed at the powerful). The problem here is that this means that there is one norm for the powerful and one norm for the powerless. To justify this distinction, there needs to be some kind of social order, recognised by all, so that the class difference between those able to demand in q_6 and those not so able is seen as acceptable. This is not at all unusual in actual societies: for example Mrs Alexander's well known hymn *All Things Bright and Beautiful*, often seen as particularly directed towards children, contains the verse (seldom sung these days):

> "The rich man in his castle, The poor man at his gate, God made them high and lowly And ordered their estate."

This is unproblematic for Consequentialist Theories: indeed given Mill's view that not all pleasures are of equal worth, the consequences are an improvement: since only the powerful *can* act to promote culture, it is good that they do so, even if it is at the expense of the powerless, since culture is preferred to pleasure. Nor is it a problem for Virtue Ethics, since it can enjoin a preference ordering: $L \succ C \succ P \succ F$.

One example of such a society is Feudalism. A good model for such a society is where some agents own the land and allow tenant farmers to work the land in exchange for the payment of rent. The nature of such a society is coloured by the ratio of powerful to powerless. If there are relatively few powerful, they can demand low rents and so leave some surplus to the tenants and allowing some degree of feasting to them (so shortening rather than removing the carnival period). This will also mean that there will be some surplus food available after the needs of the powerful have been met, some of which can be demanded to give the powerful pleasure as well as culture.

What is important, for the sustainability of such a society, is that the powerless respect the social order and do not rise up and overthrow the elite. Revolutions must be avoided. The social order can be reinforced by including a value *deference* (D), promoting by working if one has no power to demand, and by giving when food is demanded, and so promoted by the transitions q_1-q_3 and q_6-q_5. This gives the powerless arguments to respect the social order, to "know their place". Deference can reinforce the preference for C over F by being seen as promoted by the transition using $play_a$ and $work$, but not the transition $play_d$ and $work$ (the idle masters do not command respect). This value recognises two different roles: the powerful are required to promote culture (MN4a) and the rest are required to enable them to do so. Acceptance can be reinforced in several ways including patriotism, in which the powerless are encouraged to take pride in the cultural achievements of their masters; or religion, as in the hymn quoted above. As a further reinforcement, prudence suggests that the rents should not be too high.

A further possibility is that some workers may be taken out of food production and used for other purposes of benefit to all, which might be additional cultural activities (e.g. minstrels), building works (e.g. the pyramids), or whatever, and then fed from the tribute. Thus, once the despot's own needs have been considered, the surplus can be apportioned by them between allowing some retention by its producers and some public works ("bread and circuses"). Oddly the fewer in the powerful class, the greater the scope for ameliorating the lot of the powerless, and hence the society is more likely to be stable. In feudal societies it seems that the powerless suffer more when there is a weak king and squabbling barons rather than when there is a powerful king who keeps the barons in check[2]. The proportion that

[2] For example, in the Robin Hood legends the people favour Richard over his weak brother John.

is taken has been investigated in behavioural economics [38][3]. At the limit, where the classes are equally divided, there is no leeway: there the powerful requires all the surplus.

In addition to Feudalism, there are other models: slavery is one, and the kind of brigandry depicted in the film the *Magnificent Seven* is another. But these afford far less opportunity for keeping the powerless content, and so are liable to breakdown. In the film the banditry is stopped by force, and slavery was abolished, whereas Feudalism evolved into a different social order, rather than being abolished or overthrown (at least in the UK: in France things were ordered differently). The key distinction is restraint on the powerful so that revolution is not seen as worthwhile[4]. To reinforce this, we often find notions of "*noblesse oblige*" or philanthropy. We will term the associated value as *generosity* (G), and it is the *quid pro quo* of deference. This might form the basis of the moral norm:

> **MN5** It is obligatory to be generous in your treatment of the less fortunate

and the virtue ethic ordering: $L \succ C \succ G \succ D \succ J \succ P \succ F$. We still need $C \succ G$ because the point of this social order is to permit $play_a$. G is there to encourage stability, not as an end in itself. Note that, part of this accommodation is to play down which persons actually enjoy the various pleasures. Culture is now seen as a public good and $play_a$ a duty. People are expected to promote the values they can, given their social position. We have accordingly omitted the suffices indicating beneficiaries. Note, however, that generosity could lead the powerless to give away food to the needy: it could replace mercy as a motivation for MN2a and MN2b.

4.2 Wealth

In post-feudal societies we find that class and disparity remain, but that this disparity is manifested as wealth rather than physical coercion. In a sense this transition began in the feudal age, when power began to take the form of (enforceable) land ownership rather than force of arms.

When wealth is the source of power, the forcibly coercive demands of the powerful are replaced by the ability to buy the surplus. So here the transition between q_6 and q_5 becomes buy and $sell$ rather than *ask* (or *demand*) and *give*. In this model, selling is not compulsory and so the possibility of reaching q_7 is there. However not selling restricts the hoarder to promoting P and jeopardises L_o, whereas selling not only avoids demoting L_o, but also opens up the possibility of enjoying some $play_a$ or even $play_d$. For example, by selling half the surplus for two cycles, a worker would be able to save so as to accumulate sufficient wealth to spend the third in play of one or the other kinds and then buy food for the winter. This is the underlying idea of holidays, pensions, and more recently of "gap years". The balance between how the surplus is distributed between $work$, $play_a$ and $play_d$ can be left to the individuals and so made to depend on the preferences of individuals, or there may be norms imposing limits. At his point it is useful to distinguish been values that are maximisers, for which more is always better, and values which are satisficers[5] for which enough can be enough and more is of no benefit and possibly of harm: for example, one will become sated with too much feasting.

In its purest form, this model should lead to a fair degree of equality, since eventually the initially wealthy will have spent all their money, and so be forced to work, since there is no other source of income. There are, however, mechanisms which tend to allow the wealthy to maintain their position:

- The wealthy may own the land (or the means of production) and be in a position to take some proportion of the labour of others in the form of rent or profit. The situation is little different from the feudal, except that payment is now in money, not in kind. The flexibility afforded by money is more suitable to an Industrial society where production requires more than land and labour, and where produce is not bread alone, but a whole range of manufactured goods.
- The wealthy may choose to lend money at interest. Since many will regard a "bird in the hand as worth two in the bush", there is likely be takers for such loans, allowing for people with initial wealth to pay for their needs from the interest and maintain their wealth, and perhaps even, given sufficient borrowers or high enough interest rates, to increase it. Note, however, this requires some way of ensuring that the lenders can be confident that the interest will be paid, and the debt repaid. This in turn requires some kind of norm, e.g.

  > **MN6a** It is obligatory to repay debts.

  This would be associated with a new value of *trustworthiness* or *honesty* (H), promoted by observance of debts (and contracts and agreements generally) and demoted by reneging on such agreements. In order to make this more general we might prefer to use the formalation:

  > **MN6** It is obligatory to honour agreements.

- Some people may have access to wealth from outside. For example, in the sixteenth century, the Spanish rulers had a seemingly inexhaustible supply of gold and silver from the Americas.
- Deference or Generosity may mean that some agents are not required to work or pay but are simply given some kind of tribute. For example monks or priests may be supported by tithes or donations, or the infirm by alms. The latter, where the motivating value is generosity, are perhaps covered by MN4, but this could be modified to be more specific, perhaps as a version of MN5, applicable to all. But the rephrasing as MN5a means that we broaden the notion of *unable to support themselves* from incapacity to include those engaged in some other, worthwhile but unremunerative, activity. This allows us to subsume mercy under generosity, while the qualification still acknowledges justice as a value.

  > **MN5a** It is obligatory to give alms to those unable to support themselves.

  The introduction of honesty may give a value ordering.

  $$L \succ H \succ C \succ G \succ D \succ J \succ P \succ F$$

There is some scope for variation: e.g. P may be ranked higher than J without causing real problems to our moral vision. It is vital

[3] The powerful find themselves in the position of the Dictator in the Dictator Game, or Proposer in the Ultimatum Game. Both of these have been much studied in behavioral economics ([27] and [42]). These studies have suggested that it is rare for people to keep as much as they can for themselves, and that Respondents in the Ultimatum game will take nothing if offered less than what they consider to be a fair amount. Explanations for behaviour in the two games in terms of value-based argumentation can be found in [14].

[4] In the words of the blues song *Custard Pie Blues* by Sonny Terry and Brownie McGhee "You have to give me some of it, or I'll take it all away".

[5] The distinction introduced by Simon [52], although he uses it to describe the attitudes of different people with respect to a single value, namely 'utility'. See also [56] and [14]

that honesty be given such a high ranking as there will normally be reasons based on some other value to break an agreement. Indeed it could be argued that H should even be preferred to L_s since it is always possible (and perhaps desirable) to avoid entering agreements which would risk demoting L_s

We might see a conflict between MN5a and MN2 and its relaxations MN2a and MN2b. In fact what we are doing is recognising a difference between those who cannot work, and whose requests should be granted, and those who could work but choose not to do so[6]. The distinction is intended to enforce MN1, but to allow for some excusable violations (e.g. on the grounds of illness).

4.3 Turn Taking

In the previous subsection we considered situations with an initial imbalance of wealth. But it is possible, given a norm such as MN6, to enable the beneficial trade of surplus production for opportunities for $play_a$, through the mechanism of turn-taking. This arrangement, expressed here as one agent plays this year supported by another agent in return for supporting the play of that agent the following year, is in fact very common as an informal arrangement at the personal level. Many couples or groups living together will come to such an arrangement regarding chores, and the idea of "turn taking" is very common amongst children.

Turn taking also emerged in the empirical work of [37] in which a society of agents played a number of iterated prisoner's dilemma games. The agents had different degrees of *tolerance* (readiness to punish) and *responsiveness* (readiness to cooperate). What emerged was a number of stable situations: mutual cooperation and mutual defection, of course, but also some stable turn taking cycles. These turn taking cycles sometimes benefited the two agents equally, but even where one gained more from the arrangement than the other, it could still be beneficial to both, and to their combined score, when compared with mutual defection. Therefore we might well see such an arrangement emerge, even in an initially equal society, given that C is preferred to P and there is a reinforcing norm such as MN6. As has been noted above, such arrangements are likely to be especially common in domestic situations, where trust is likely to be high. This in turn suggests that it might be possible to differentiate H according to whom it is directed. It is not uncommon to regard it as wrong to cheat family and friends (H_f), dubious to cheat other individuals (H_i), but acceptable (where possible) to take advantage of large ("faceless") organisations (H_o). Such discrimination is rarely enjoined by any ethical theory (although it is possible that, in some circumstances, it would be endorsed by some forms of consequentialism), but is a commonly argued for (and practiced) behaviour. Over-claiming on insurance is not uncommon and is seen by some as a "victimless" crime, suggesting that some might give H_o a very low rank, perhaps even below F.

4.4 Service Provision as Work

In several of of the scenarios discussed previously it came about that because of the preference for C over some other values, certain agents may be enabled to $play_a$ because the consequent promotion of C was such that other agents were inclined to support this activity out of their surplus in preference to P. This is likely to be particularly so in the case pf powerful agents who will choose to act as patrons to certain agents to allow and encourage certain kinds of $play_a$. But similar kinds of patronage may be attractive to other individuals as well, who may be prepared to part with a (typically) small part of their surplus. It is possible that this may emerge with just two agents. The ant may find the singing of the grasshopper so entertaining that he is willing to sacrifice his entire surplus for the privilege of listening to her. But, since the singing of a single grasshopper may entertain a whole colony of ants, it is even more attractive if the cost of supporting the grasshopper can be shared across a large number of individuals. Where this is so, a variety of entertainers can be supported, and other services performed. Money greatly assists this arrangement, and places it on a formal, contractual footing, so that it falls under MN6. As such we might expect the emergence of a service and entertainments sector, where some agents were able to adopt the role of providers of C promoting activities willingly supported by groups of other agents.

This is likely to be increasingly the case when productivity rises, so that workers generate larger surpluses. Now we can adjust our notions of the difference between $play_a$ and $play_d$. We can see $play_a$ as being non-work activities for which people are prepared to pay, and $play_d$ as non-work activities for which people are not prepared to pay. This will require consideration of the agent as well as the activity: people will pay to watch Lionel Messi play football, but no one will pay to watch me play football. We therefore combine norms MN1 and MN4a into single norm:

MN1a It is obligatory to $play_a$ or to $work$.

This differs from MN4 because that norm was directed at only a subset of agents, whereas MN1a can be seen as universal. Interestingly a norm like MN1a may be better supported by a system of reward for $play_a$ rather than punishment for $play_d$. Indeed the payment for the services provided for $play_a$ may well be seen in terms of reward for norm compliance. For a discussion of enforcing norms with rewards rather than punishments see [19].

4.5 Emergence of a State

As well as choosing to spend their surplus on providing themselves with culture, through paying others to $play_a$, agents may choose to pay others to do their duties. In [39] it was shown empirically that to avoid norms collapsing it is necessary that they not only be backed by the punishment of violators, but that those who fail to punish must themselves be punished. Since punishment has a cost, however, there are reasons not to punish, and in societies where violations are comparatively rare, the cost of punishment falls unevenly and unpredictably. We saw how punishment for violating MN1 can naturally be expressed as MN2 (which actually is cost free for the punisher), but when we move to more sophisticated norms such as MN6, punishment may not have a simple manifestation as a norm. Recognising the need to punish is an important aspect of social cohesion: as expressed in [39]:

> This move from enforcement by vigilantes (those taking the law into their own hands) to seeing law enforcement as the social duty of responsible citizens is an important milestone in the development of a society that respects its laws.

Once punishment is seen as a social duty it is a small step to organise and pay for a third party to punish violators. Assuming relatively

[6] The distinction between the deserving and undeserving poor was a central concern of the UK 1834 Poor Law Amendment Act, and is enjoying a revival in popular attitudes expressed in the UK today. It contrasts with the underlying philosophy of the UK Supplementary Benefits Act 1976, which saw a certain minimal level of support as the right of every citizen.

few law breakers a small levy will enable a dedicated agent to be paid to enforce the norms. Of course, non-payment of the levy will also be subject to punishment. From this it is a small step to taxation, and the provision of services such as law enforcement by the state. And if law enforcement, why not other duties? Thus MN5a may be better observed by contribution to a central fund responsible for identifying those who should be supported and providing that support.

In this way States may emerge, first as a Hobbesian *Leviathan* [34], but, once established, available to take on the performance of other duties. Further the State may take on the role of intervention to resolve conflicts of interest between its citizens [23], or to educate its citizens [37]. An emergent State may also lead to new values such as *self-reliance*, *freedom* and *community*, and the relative preferences of these new values, and how they are promoted and demoted by different models of the State may provide insight into the form in which the State emerges. In some circumstances the State may take on an even broader role, and become itself the arbiter of what constitutes $play_a$, by itself supporting certain activities. Thus we often find subsidies for opera, but never for football. Of course, allowing the state to determine what counts as culture in this way will be controversial, and so we may find that we need to to distinguish between two types of $play_a$: high culture as approved by the state and subsidised ($play_{sa}$) and popular culture approved by citizens and paid for out of their own retained surplus ($play_{pa}$). This provides another example of how increasing the level of sophistication of the model necessitates the finer grained discrimination of values and actions.

5 Discussion

As observed by Hare [32], for most people, most of the time, following moral norms involves little more that applying a set of learned principles. Hare, however, also says that there will be occasions when we need to think out a moral problem from fist principles, and that the recognised norms are a useful summary of such reasoning.

> 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.

When thinking about the emergence of norms, it is is this deep thinking that gives rise to the norms that we need to model. In this paper we have argued that value-based practical reasoning applied to a model of society expressed as an AATS+V provides the machinery to model this kind of reasoning. Much current work on norm emergence is done using either simulations of public goods games or by proving properties of such games as in [51], or by performing model checking on state transition diagrams as in [62]. The first approach has given some insights, but the simplification necessary, and assumptions about the homogeneity of agents, suggest that there are limitations to the approach. These doubts are strengthened by the fact that the behaviour of people observed empirically in experiments using such games does not support the model used [27] and [42]. The second approach also has a view of agents as highly goal directed, and tends to simplify its representation of norms by removing transitions representing forbidden actions. This means that it is highly effective at proving properties of the system, when the norms are complied with and for verifying the design of norms, but less good in explaining where the norms come from in the first place, and why the agents wish to pursue them, If we are looking for emergence rather than imposition by a designer this is a problem. We believe that the use of value-based argumentation provides a finer grainer account of the reasoning involved, and is therefore better placed to account for the norms that emerge from different social set-ups.

In section 3 we described how two norms might emerge in a simple society. One is a primary norm, the other provides a natural way of punishing transgressions of the primary norm (and a way of removing transgressors). We believe that although the model is simple, it is a not implausible representation of a primitive agricultural society. Subsequently we described how making the model more sophisticated would lead to other norms, and more importantly to the need to introduce additional values (some of which may be *metavalues* promoted and demoted by value orderings rather than actions) and to make finer grained discriminations both in values and in actions. Thus $play$ becomes seen as the socially beneficial $play_a$ and the indulgent $play_d$ and a need to discriminate the value of honesty according to the relationship between the agents involved in the transaction may become apparent. Unfortunately the provision of detailed models, and the particular arguments that they support, is beyond the scope of this workshop paper: all that is possible here is to sketch how additions to the model would result in different norms, and so give a flavour of the process.

We believe that such detailed models would indeed provide a fruitful way of analysing and explaining social developments. Our account here for example, coheres well with the account of social development found in Durkheim [26]. Durkheim suggests that in a "primitive" society people act and think alike with a collective or common conscience, which is what allows social order to be maintained. In such a society laws tend to be highly repressive. Both of these are true of the model presented in section 3, where there is a norm (MN1) to be followed by all and transgressions are effectively punished by death through MN2. Durkheim further argues that in an advanced, industrial, capitalist society, the complex division of labor means that people are allocated in society according to merit and rewarded accordingly, and that diversity is embraced rather than opposed. This accords with our discussion of the norms that develop as surplus production increases, and the development of exchanges enabled by MN6, leading to the increasing prevalence and diversity of service work, rather than food production. Within this framework we could, for example, explore the different norms that emerge when the surplus comes from a general rise in productivity from where it comes as the result of an external boost to wealth, as in sixteenth century Spain. Note also that the sophisticated societies require increased cooperation (supported by norms such as MN6 and values such as trust and honesty) and tended to increase the degree of commercial exchanges between agents. It was these two factors that were found to lead to the greatest deviation from the classical model in the Ultimatum Games studied in [33], supporting the view that the more sophisticated the society the less adequate the model provided by simple public goods game simulations. Thus even if these simulations provide a good account of how initial norms *emerge*, investigating their *development* may require a finer grained approach.

As a final remark we may return to the types of ethical theory mention in 2.4. The consequentialist approach to ethics is reflected in both public goods game simulations which picture agents as homogeneous utility maximisers, and the STD based reasoning of [3] which designates states as desirable and undesirable. In contrast the value-based approach, which allows for agents to have different desires and aspirations represented by their different ordering on values, is more in the virtue ethics tradition. Norms encourage a value order such that agents will want to choose the "right" actions.

REFERENCES

[1] Aesop, *Fables, retold by Joseph Jacobs*, volume Vol. XVII, Part 1, The Harvard Classics. New York: P.F. Collier and Son, 1909-14.

[2] 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*, pp. 145–152, (2009).

[3] T. Ågotnes and M. Wooldridge, 'Optimal social laws', in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, pp. 667–674, (2010).

[4] R. Alur, T. Henzinger, and O. Kupferman, 'Alternating-time temporal logic', *Journal of the ACM (JACM)*, **49**(5), 672–713, (2002).

[5] Aristotle, *The Nicomachean Ethics of Aristotle, translated by W.D. Ross*, Heinemann, 1962.

[6] K. Atkinson and T. Bench-Capon, 'Practical reasoning as presumptive argumentation using action based alternating transition systems', *Artificial Intelligence*, **171**(10-15), 855–874, (2007).

[7] K. Atkinson and T. Bench-Capon, 'Addressing moral problems through practical reasoning', *Journal of Applied Logic*, **6**(2), 135–151, (2008).

[8] K. Atkinson and T. Bench-Capon, 'Taking the long view: Looking ahead in practical reasoning', in *Computational Models of Argument - Proceedings of COMMA 2014*, pp. 109–120, (2014).

[9] K. Atkinson, T. Bench-Capon, and S. Modgil, 'Argumentation for decision support', in *Database and expert systems applications*, pp. 822–831. Springer, (2006).

[10] R. Axelrod, 'An evolutionary approach to norms', *American political science review*, **80**(04), 1095–1111, (1986).

[11] R. Axelrod, *The evolution of cooperation*, Basic Books, 1987.

[12] T. Bench-Capon, 'Persuasion in practical argument using value-based argumentation frameworks', *J. of Logic and Computation*, **13**(3), 429–448, (2003).

[13] T. Bench-Capon, K. Atkinson, and A. Chorley, 'Persuasion and value in legal argument', *J. of Logic and Computation*, **15**(6), 1075–1097, (2005).

[14] T. Bench-Capon, K. Atkinson, and P. McBurney, 'Using argumentation to model agent decision making in economic experiments', *Autonomous Agents and Multi-Agent Systems*, **25**(1), 183–208, (2012).

[15] J. Bentham, *The rationale of reward*, John and HL Hunt, 1825.

[16] 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).

[17] C Bicchieri, *The grammar of society: The nature and dynamics of social norms*, Cambridge University Press, 2005.

[18] K. Binmore, 'Review of robert axelrod complexity and cooperation', *Journal of Artificial Societies and Social Simulation*, **1**(1), (1998).

[19] A. Boer, 'Punishments, rewards, and the production of evidence', in *Proceedings of JURIX 2014*, pp. 97–102, (2014).

[20] M.E. Bratman, *Intention, Plans, and Practical Reason*, The David Hume Series, Cambridge University Press, 1999.

[21] B. Burgemeestre, J. Hulstijn, and Y-H. Tan, 'Value-based argumentation for justifying compliance', *AI and Law*, **19**(2-3), 149–186, (2011).

[22] D. Cartwright and K. Atkinson, 'Using computational argumentation to support e-participation', *Intelligent Systems*, **24**(5), 42–52, (2009).

[23] A. Chorley, T. Bench-Capon, and P McBurney, 'Automating argumentation for deliberation in cases of conflict of interest', in *Proceedings of COMMA 2006*, pp. 279–290. IOS Press, (2006).

[24] P. M. Dung, 'On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games', *Artificial intelligence*, **77**(2), 321–357, (1995).

[25] P. Dunne, 'Tractability in value-based argumentation', in *Proceedings of COMMA 2010*, pp. 195–206, (2010).

[26] E. Durkheim, *The division of labor in society*, Simon and Schuster, 2014. First published 1893.

[27] C. Engel, 'Dictator games: a meta study', *Experimental Economics*, **14**(4), 583–610, (2011).

[28] M. Esteva, D. De La Cruz, and C. Sierra, 'Islander: an electronic institutions editor', in *Proceedings of AAMAS 02*, pp. 1045–1052, (2002).

[29] A. Garcez, D. Gabbay, and L. Lamb, 'Value-based argumentation frameworks as neural-symbolic learning systems', *J. of Logic and Computation*, **15**(6), 1041–1058, (2005).

[30] G. Governatori, 'Thou shalt is not you will', in *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pp. 63–68. ACM, (2015).

[31] F. Grasso, A. Cawsey, and R. Jones, 'Dialectical argumentation to solve conflicts in advice giving: a case study in the promotion of healthy nutrition', *International Journal of Human-Computer Studies*, **53**(6), 1077–1115, (2000).

[32] R. M. Hare, *Freedom and reason*, Oxford Paperbacks, 1965.

[33] J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, and R. McElreath, 'In search of homo economicus small-scale societies', *The American Economic Review*, **91**(2), 73–78, (2001).

[34] Thomas Hobbes, *Leviathan, 1651*, Scolar Press, 1969.

[35] A. Jones and M. Sergot, 'Deontic logic in the representation of law: Towards a methodology', *AI and Law*, **1**(1), 45–64, (1992).

[36] I. Kant, *Groundwork of the Metaphysics of Morals*, CUP, 1998. First Published 1785.

[37] M. Lloyd-Kelly, K. Atkinson, and T. Bench-Capon, 'Fostering cooperative behaviour through social intervention', in *Proceedings of (SIMULTECH), 2014*, pp. 578–585. IEEE, (2014).

[38] G. Loewenstein, 'Experimental economics from the vantage-point of behavioural economics', *The Economic Journal*, **109**(453), 25–34, (1999).

[39] S. Mahmoud, N. Griffiths, J. Keppens, A. Taweel, T. Bench-Capon, and M. Luck, 'Establishing norms with metanorms in distributed computational systems', *AI and Law*, **23**(4), 367–407, (2015).

[40] J.S. Mill, *Utilitarianism*, Longmans, Green, Reader, and Dyer, 1871.

[41] Samer Nofal, Katie Atkinson, and Paul E Dunne, 'Algorithms for decision problems in argument systems under preferred semantics', *Artificial Intelligence*, **207**, 23–51, (2014).

[42] H. Oosterbeek, R. Sloof, and G. Van De Kuilen, 'Cultural differences in ultimatum game experiments: Evidence from a meta-analysis', *Experimental Economics*, **7**(2), 171–188, (2004).

[43] C. Perelman, *The new rhetoric*, Springer, 1971.

[44] I. Rahwan and L. Amgoud, 'An argumentation based approach for practical reasoning', in *Proceedings of AAMAS 06*, pp. 347–354, (2006).

[45] A. S. Rao and M.P. Georgeff, 'Modeling rational agents within a bdi-architecture', in *KR 91*, pp. 473–484, (1991).

[46] A. Rapoport and A. Chammah, *Prisoner's dilemma: A study in conflict and cooperation*, volume 165, University of Michigan press, 1965.

[47] J. Raz, *Practical Reasoning*, Oxford University Press, Oxford, 1979.

[48] A.E. Roth and J. K. Murnighan, 'Equilibrium behavior and repeated play of the prisoner's dilemma', *Journal of Mathematical psychology*, **17**(2), 189–198, (1978).

[49] B. Savarimuthu, Maryam Purvis, Martin Purvis, and S. Cranefield, 'Social norm emergence in virtual agent societies', in *Declarative Agent Languages and Technologies VI*, 18–28, Springer, (2008).

[50] S. Sen and S. Airiau, 'Emergence of norms through social learning.', in *IJCAI*, volume 1507, p. 1512, (2007).

[51] Y. Shoham and M. Tennenholtz, 'On the emergence of social conventions: modeling, analysis, and simulations', *Artificial Intelligence*, **94**(1), 139–166, (1997).

[52] H. A. Simon, 'Rationality as process and as product of thought', *The American economic review*, **68**(2), 1–16, (1978).

[53] Brian Skyrms, *Evolution of the social contract*, CUP, 2014.

[54] T. Sugawara, 'Emergence and stability of social conventions in conflict situations', in *IJCAI*, pp. 371–378, (2011).

[55] I. Tremblay, J.and Abi-Zeid, 'Value-based argumentation for policy decision analysis project in québec', *Annals of Operations Research*, **236**(1), 233–253, (2016).

[56] A. Tversky and D. Kahneman, 'Judgment under uncertainty: Heuristics and biases', *science*, **185**(4157), 1124–1131, (1974).

[57] E. Ullmann-Margalit, *The emergence of norms*, Clarendon Press Oxford, 1977.

[58] W. van Der Hoek, M. Roberts, and M. Wooldridge, 'Social laws in alternating time: Effectiveness, feasibility, and synthesis', *Synthese*, **156**(1), 1–19, (2007).

[59] T. van der Weide, F. Dignum, J-J Ch Meyer, H. Prakken, and GAW Vreeswijk, 'Multi-criteria argument selection in persuasion dialogues', in *Argumentation in Multi-Agent Systems*, 136–153, Springer, (2011).

[60] A. Walker and M. Wooldridge, 'Understanding the emergence of conventions in multi-agent systems.', in *Proceedings of ICMAS 95*, pp. 384–389, (1995).

[61] M. Wooldridge, *An introduction to multiagent systems*, John Wiley & Sons, 2009.

[62] M. Wooldridge and W. van der Hoek, 'On obligations and normative ability: Towards a logical analysis of the social contract', *Journal of Applied Logic*, **3**, 396–420, (2005).