Altruism and Agents: An Argumentation Based Approach to Designing Agent Decision Mechanisms

Trevor Bench-Capon
Dept. of Computer Science
University of Liverpool
Liverpool, UK
tbc@csc.liv.ac.uk

Katie Atkinson
Dept. of Computer Science
University of Liverpool
Liverpool, UK
katie@csc.liv.ac.uk

Peter McBurney
Dept. of Computer Science
University of Liverpool
Liverpool, UK
mcburney@liverpool.ac.uk

ABSTRACT

We present an argument-based qualitative decision-making framework in which the social values promoted or demoted by alternative action-options are explicitly represented. We show how this framework may be used to explain the results of experimental economic studies in which human subjects play the Ultimatum Game, an interaction between two participants in which one player divides a sum of money between them, and the other player may accept or reject the offer. The results of these experiments are not explained by a decision-model assuming the participants are purely self-interested utility-maximizers. Some studies further suggest that differences in choices made in different cultures may reflect their day to day behaviour, which can in turn be related to the values of the subjects, and how they order their values. The decision-framework we propose will aid software engineers designing decision-making mechanisms for autonomous agents, particularly for situations requiring agent adaptability, for example, where agents may prefer different outcome states in transactions involving different types of counter-parties.

1. INTRODUCTION

Autonomous agents are expected to make their own decisions. But what is the basis for their choice? A natural inspiration would be the basis for choice used by human agents. One suggestion to explain human behaviour, a foundational assumption of much economic theory, is that humans act so as to maximise their satisfaction, well-being or utility. The idea was stated succinctly by John Stuart Mill [22]:

> "*[Political economy] does not treat the whole of man's nature as modified by the social state, nor of the whole conduct of man in society. It is concerned with him solely as a being who desires to possess wealth, and who is capable of judging the comparative efficacy of means for obtaining that end.*"

This assumption has been explored and questioned in experimental economics. Two experiments that have been widely used are the *Dictator Game* (e.g. [11]) and the *Ultimatum Game* (e.g. [23]). In the Dictator Game the first player is given a sum of money and told that he may give as much or little as he likes to his partner. Once he has decided on an allocation the players receive the amounts proposed, and the game ends. The Ultimatum Game builds on the Dictator Game by allowing the second player an option: the second player may choose to accept the proposed allocation, or reject it. If the proposal is rejected both players receive nothing. If players were really motivated only by self interest, the expectation would be that Dictators would keep all the money, and Proposers in the Ultimatum Game would offer their partner the minimum amount, which would be accepted on the grounds that something, however small, is better than nothing. In practice these expectations are not met. Experimental studies using the Dictator Game suggest that typically 70% of dictators give non-zero sums and transfer around a quarter of the initial sum. None of the many studies offers support for the canonical model. For example, in one typical study [11], given \$10 to distribute, 79% of participants gave away a positive amount, with 20% giving away half. The mode sum given away in that study was \$3. Similar deviations are found in the Ultimatum Game: for example, Nowak and colleagues report that the majority of proposers offer 40–50% and about half of responders reject offers below 30% [23]. These results are robust, and, with some variations, are replicated in all the many studies. Oosterbeek *et al.* [24] report a meta-analysis of 37 papers with 75 results from Ultimatum Game experiments, which have an average of 40% offered to the responder. The experiments of Henrich *et al.* [13], carried out over fifteen small-scale societies in twelve countries of five continents, report mean offers between 26% and 58%, and note than in some societies there is considerable variation in which offers are rejected: however, again none suggests that the canonical model is followed by those making and responding to offers.

To explain why dictators and ultimatum game proposers and responders do not act as predicted, a number of suggestions have been made as to what other considerations are being taken into account. These suggestions include: that the benefit of the other player has some positive utility; that the action of giving, in itself, confers utility; that there is a sense of fairness which suggests to participants that the money should be shared, perhaps even equally, between them; and that people do not wish to appear selfish. A number of experiments have attempted to isolate or control for these factors in an effort to confirm or disconfirm their influence.

Two other points need to be made: first that there is a great deal of heterogeneity between subjects. While studies do identify variations across cultures, there is always a significant amount of variations within cultures. Second, the way the problem is presented can have a significant effect, known as the *framing* effect; this effect is noted by Bolton and colleagues [7]. An experiment by Bardsley designed to explore these framing effects [5] shows that dictators are significantly less generous when it is represented to them that the experimenter has given their partner the money and they can

Cite as: Altruism and Agents: An Argumentation Based Approach to Designing Agent Decision Mechanisms, Trevor Bench-Capon, Katie Atkinson and Peter McBurney, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. XXX-XXX.

take as much as they want for themselves than they are in the usual framing where the dictators are given the money and asked how much they wish to give away.

The contribution of this paper is an argumentation framework for agent decision-making which explicitly represents the social *values*[1] promoted or demoted by alternative action options. This framework generalizes utility-based approaches and provides a more coherent and more complete explanation of the evidence from these empirical studies. Agent software engineers seek to develop autonomous software agents capable of independent decision-making, including the ability to respond appropriately to situations not necessarily envisaged by the designer. If such software agents are to be selfish utility-maximisers (in the manner of *homo economicus*), then mainstream economic theory provides a basis for the engineering of their autonomous computational decision-making mechanisms. Thus, we have already seen a lot of work in this vein in the agents community, e.g., [25]. However, human beings are not always utility-maximisers, as revealed by evidence in experimental and behavioural economics [18]. The invariable deviance from the canonical model suggests that there must be beneficial effects in doing so. While there are considerable variations, suggesting that there is room, or need, for diversity, it is striking that the canonical model is never followed. Accordingly, designers of software agents may need to create autonomous computational decision-mechanisms which are closer to actual human decision-making processes – either because such software agents are acting on behalf of human principals who wish their agents to mimic their human decision-making processes, or because such agents are interacting with other humans or with agents acting on behalf of other humans, or simply because this aids the functioning of an agent society, and produces superior mechanisms for making decisions. It is important, therefore, for agent designers to have access to models of decision-making which provide a good fit to actual human decision-making processes: we present such a model in this paper.

2. BEYOND SIMPLE UTILITY MAXIMISATION

One way of accommodating these results from Dictator and Ultimatum Game experiments is to retain the idea that agents maximise their utility but to make the utility function more complicated, by including these other factors, suitably weighted in accordance with cultural and individual preferences. An alternative, qualitative, approach was, however, proposed in [3] where the authors advocated an approach based on argumentation. The various possible influences are seen as reasons to motivate and justify a choice between the various options, and as a basis for critiquing the reasons offered. This results in a number of conflicting arguments which can be resolved using a technique developed in the argumentation in AI community [6] based on a ranking of the various motivating factors, which can of course vary across individuals and cultures. The approach was applied to the Dictator game in [3] and a number of advantages for the approach were claimed. First, the transparency of the approach: giving explanations and justifications of the choices in terms of arguments is more informative and more open to discussion and criticism than referring to a formula for the utility function which can only be obtained by fitting the function to the choices made. Second, importantly, the argumentation framework can explain the framing effect. Whereas if the choice depended only on estimating the utility of the state reached, we should expect the same individuals to choose the same outcome whatever their initial position, in the argumentation approach the arguments available depend on the initial state as well as the target state. Given that there are therefore different arguments available depending on the way in which the problem is framed, we would even expect to see these framing effects. In the remainder of this section we will describe the argumentation approach to these games. In section 3 we will apply this approach to the Ultimatum Game. Section 4 will discuss how these findings might be used in the design of multi-agent systems. Section 5 will provide some further discussion and offer some concluding remarks and directions for future work.

[1]Values in our sense should not be confused with any kind of quantitative measures. We use "‘values"’ in a sense common in current English usage, in which, for example, the values of the French Republic are liberty, equality and fraternity. This sense is in daily use by politicians and journalists, who appeal to Christian values, socialist values, British values, etc.

2.1 The Argumentation Approach

The argumentation approach to the economic experiments is based on the general argumentation approach to practical reasoning developed in [4] The idea is that an option is presumptively justified by an argument which instantiates an argument scheme based on the practical syllogism. Instantiations of this scheme can then be critiqued by posing critical questions characteristic of the scheme. In turn attempts to rebut these critiques can be made.

The descriptive version of the scheme is as follows:

AS1 In the current circumstances R
We should perform action A
Which will result in new circumstances S
Which will realise goal G
Which will promote value V.

AS1 is an extension of Walton's sufficient condition scheme for practical reasoning [27] in which Walton's notion of a 'goal' is articulated in more detail by separating it into three elements: the state of affairs brought about by the action; the goal (the desired features in that state of affairs); and the value (the reason why those features are desirable). The justification is only presumptive: a set of critical questions can be posed challenging the various components of the scheme: for example one may deny that the current situation is as described, that the action will realise the goal, or that the goal will promote the value. In [4] sixteen of these critical questions are given.

In order to apply this approach to a particular problem, it is first necessary to formulate the problem in such a way that instantiations of the argument scheme and critical questions can be identified. For this we follow [2] and describe the problem as an Action Based Alternating Transition System (AATS) [28]. An AATS is a state transition diagram in which the transitions represent *joint actions*, that is actions composed from the individual actions available to the agents in that state. Additionally we label the transitions with the values promoted and demoted by moving from the source to the target state. A summary of the AATS representation that we use is given below:

- Each state transition diagram comprises a set *Q* of *states* of which one state, *q0*, is designated the *initial state*. A state is a consistent conjunction of literals.
- *Ag* is a set of agents.
- A_i is the set of *actions* available to a particular agent, ag_i.
- *J* is the set of joint actions, where a joint action is a tuple comprising one element of A_i for each agent ag_i in *Ag*.

- The *state transition function* defines the state that results from the execution of each joint action in a given state.
- A goal is a consistent conjunction of literals. A goal *g* is realised in a state *q* if *g* is true in *q*.
- *V* is the set of *values* relevant to the scenario.
- The *valuation function* defines the status (promoted +, demoted –, or neutral =) that labels the transition between two states.

Given this model, arguments can then be generated that propose and attack particular actions based on the values promoted through execution of the actions. For example, consider the simple AATS in Figure 1. States are represented by two propositions (P and Q) which can be true or false. We have two agents, each of which can perform two actions (a or b) in q1. There are thus four possible joint actions and acting so as to move to q2 promotes V1 and acting so as to move to q3 promotes V2. Where a transition promotes a value AS1 can be instantiated to justify an agent in its performance of its component of the corresponding joint action. Thus here the first agent can justify doing action a by the argument *in q1 I should perform a so as to reach q2 which will make P true and promote V1*. This can be critiqued by objections, such as *the second agent may not perform b and so q2 will not be reached* or *performing a means that q3 will not be reached and so V2 will not be promoted*. This second objection may be met if, for example, the agent expresses a preference for V1 over V2.

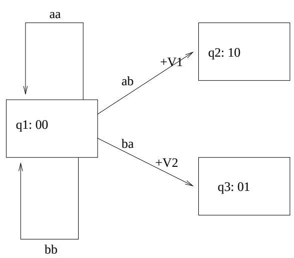

Figure 1: Simple AATS

Instantiating the argument scheme and the critical questions gives rise to a set of conflicting arguments. Once this set has been produced we need some mechanism by which we can evaluate the arguments to determine their acceptability. We do this by organising them into a Value Based Argumentation Framework (VAF) [6]. A VAF is an extension of the standard Argumentation Frameworks (AF) introduced by Dung [8], which provide a formal means of evaluating arguments based on consideration of the attacks between a set of conflicting arguments. An AF can be pictured as a directed graph with the nodes representing arguments, and the edges an attack of one argument by another. The purpose is to find a subset of the arguments which is at once conflict free (i.e. no two arguments in the subset attack one another), and collectively able to defend itself (i.e. any attacker of an argument in the subset is itself attacked by an argument in the subset). The maximal such subset is called a *preferred extension*, and represents a maximal consistent position given the arguments presented. VAFS, and some associated notions are formally defined in Definitions 1 and 2.

Definition 1: Value-Based Argumentation Framework
A Value-Based Argumentation Framework (VAF) is defined by a triple $\langle H(X,A), \nu, \eta \rangle$, where $H(X,A)$ is an argumentation framework, $\nu = v_1, v_2, ..., v_k$ a set of k values, *and $\eta : X \rightarrow \nu$ a mapping that associates a value $\eta(x) \in \nu$ with each argument $x \in X$.*
A specific audience, *α, for a vaf $\langle H, \nu, \eta \rangle$, is a total ordering of ν. We say that v_i is preferred to v_j in the audience α, denoted $v_i \succ_\alpha v_j$, if v_i is ranked higher than v_j in the total ordering defined by α.*

Definition 2: Concepts Relating to VAFs
Let $< H(X,A), V, \eta >$ be a VAF and α an audience.
a. For arguments x, y in X, x is a successful attack on y (or x defeats y) with respect to the audience α if: $< x, y > \in A$ and it is not the case that $\eta(y) \succ_\alpha \eta(x)$.
b. An argument x is acceptable to the subset S with respect to an audience α if: for every $y \in X$ that successfully attacks x with respect to α, there is some $z \in S$ that successfully attacks y with respect to α.
c. A subset R of X is conflict-free with respect to the audience α if: for each $< x, y > \in R \times R$, either$< x, y > \notin A$ or $\eta(y) \succ_\alpha \eta(x)$.
d. A subset R of X is admissible with respect to the audience α if: R is conflict free with respect to α and every $x \in R$ is acceptable to R with respect to α.
e. A subset R is a preferred extension for the audience α if it is a maximal admissible set with respect to α.

VAFs extend AFs in that each argument in the graph is associated with the value promoted by that argument. The purpose of this extension is to distinguish attack from defeat, relative to the audience's preference ordering on the values. Whereas in an AF attacks always succeed, in a VAF they succeed only if the value associated with the attacker is ranked by the audience evaluating the VAF equal to, or higher than, the argument it attacks. Unsuccessful attacks are removed, and then the VAF can be evaluated as a standard AF. The VAF thus accounts for elements of subjectivity in that the arguments that are acceptable are dependant upon the audience's ranking of the values involved in the scenario.

A fully worked example of applying this approach to generating and evaluating actions to the Dictator Game is given in [3]. In the next section we will apply this approach to the Ultimatum Game.

3. MODELLING THE ULTIMATUM GAME

To apply the argumentation approach to the Ultimatum Game we must first construct the appropriate AATS. This will involve first identifying the propositions we wish to include in our states, next the actions the agents can perform, then the values we wish the agents to consider, and finally associate transitions with values.

Obviously the states must include the money held by the two agents. We also wish to represent the reactions of the two players. When the offer is made, it is important whether the second player perceives it as fair, or as insulting. We therefore use a proposition which is true when the second player is annoyed by the offer made. At the end of the game we can consider the reaction of the first player. In particular if the offer is rejected, a first player who made an ungenerous offer is likely to feel regret that he did not offer more. We therefore use a fourth proposition to record whether the first player feels regret.

Next we turn to actions. Obviously we need that the first player can offer n% of the available sum to the second player and that the second player can accept or reject it. The reception the offer receives will, however, depend critically on the size of n. We will therefore distinguish four cases: where n > 50, where n = 50, where n > 0 but < 50 and where n = 0. We should also recognise that the the two actions are not chosen simultaneously, and that the choice to accept or reject will depend on how the second player reacts to the offer of the proposer. We therefore introduce a third action, in which the second player chooses a threshold, t, above which he will regard the offer as just, and below which he will feel insulted. We will assume that t > 0 and t < 50, discounting players who will not be satisfied with even an equal share. While the second player accepts and rejects the first player can do nothing. This gives the set of joint actions shown in Table 1.

Table 1: Joint Actions

Joint Action	Player 1	Player 2
j1	A1:Offer > 50	B1:Set t < 50
j2	A2:Offer 50	B1:Set t < 50
j3	A3:Offer n < 50 and > 0	B1:Set t < n
j4	A3:Offer n < 50 and > 0	B1:Set t > n
j5	A5:Offer n = 0	B1:Set t > 0
j6	A4:Do nothing	B2:accept
j7	A4:Do nothing	B3:reject

Now consider the transitions. An offer will have the effect of moving from the initial state where both players have 0 to one where the first player has 100-n and the second player has n. Moreover where n does not exceed t, the second player will be insulted. Accepting the offer leaves the amounts unchanged, while rejecting the offer returns both amounts to 0. Where the second player is insulted, rejecting the offer expiates the insult. Finally if the first player has offered less than half and has been rejected he will experience regret. The transitions are shown in the AATS in Figure 2.

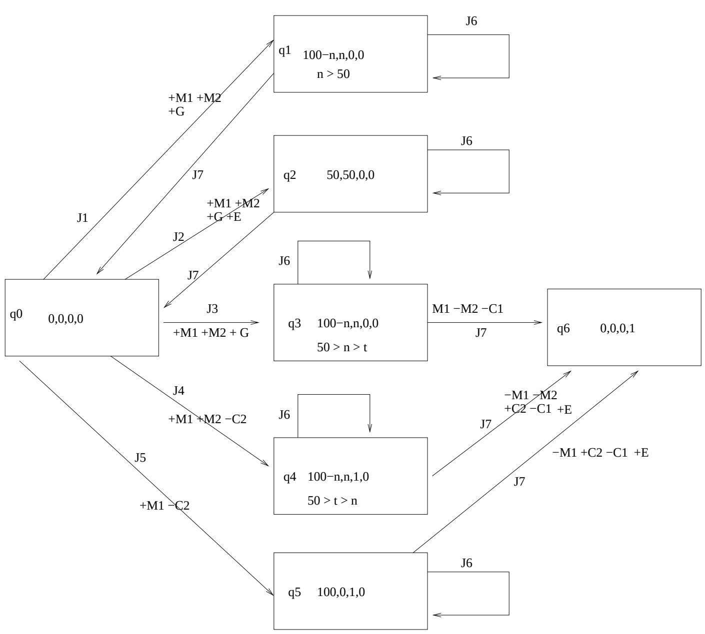

Figure 2: AATS for Ultimatum Game

Now we must identify some values and the transitions which promote and demote them. First there the economic value, the money, which we shall call M. This can be promoted in respect both of player 1 (M1) and in respect of player 2 (M2). These values are promoted to different degrees according to the size of the player's share. Next we take from the literature that some people seem to value fairness, which we shall call E for equality. This is either promoted or not. Third we have the value of generosity, (G) which again has been identified as a motivation by various experimenters. Whereas M will be promoted to varying degrees according to the amount of money, E is either promoted or not. What of G? Experimental evidence suggests that the impact of G does not increase as the amount given increases: we will therefore consider that G, like E, is either satisfied or not, and that any effect of the size of the gift is reflected in M(2). Finally either player may be content with the outcome, and we represent this as C(1) and C(2). Again we will not model degrees of contentment. Labels indicating the promotion and demotion of these values are shown on the AATS in Figure 2.

We can now generate arguments. Each promotion of a value will provide an instantiation of AS1, justifying the agent in its choice of its own component of the corresponding joint action, and each demotion of a value will constitute an objection to that action. Moreover if the value M could be promoted to a greater degree that would be an objection to performing the less lucrative action.

Consider first the reasoning of the second player responding to the offer, who will be in one of q1 to q5. In each of these states the second player needs to consider whether rejection is justified. Accepting stays in the same state, and does not promote any values, and so will be chosen only if the objections to rejecting are preferred to the justification. The arguments justifying rejection and the objections to them are shown in Table 2. An argument for rejecting in q2 is included for completeness, although it provides no justification since no value is promoted, and will therefore be defeated by any value based objection.

Table 2: Arguments for the Second Player

ID	Argument	Objections
S1a	Reject in Q1 to promote E	Demotes M1, M2
S2a	Reject in Q2 (no reason)	Demotes M1, M2
S3a	Reject in Q3 to promote E	Demotes M1, M2
S4a	Reject in Q4 to promote C2 and E	Demotes M1, M2, C1
S5a	Reject in Q5 to promote C2 and E	Demotes M1, C1

This gives rise to the VAF shown in Figure 3.

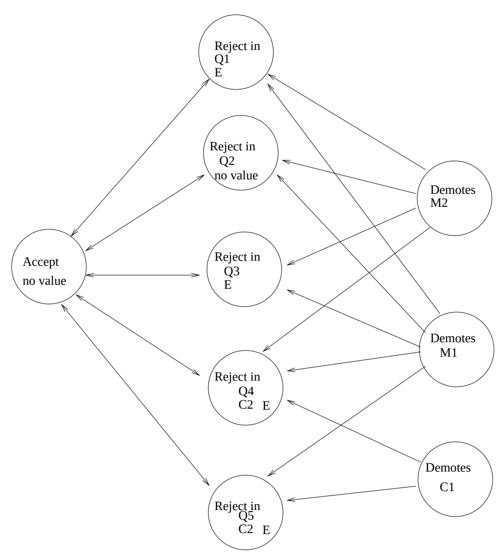

Figure 2: VAF for acceptance or rejection

What the second player will do will depend on how it orders its values. Thus an offer above 50, or below 50 but above the second player's threshold of acceptability (states Q1 and Q3), will only be rejected if the player prefers equality to both its own and the other player's, money: $E > \{M(1),M(2)\}$. Given the set of values we have used, we would expect any player to accept an offer of half the sum, since rejecting in Q2 promotes nothing and demotes money for both players. If the second player is insulted by a non zero offer and so is in Q4, however, he has a choice of whether to punish the first player and so restore its own equanimity, or to take the money. Normally we would expect that the player will prefer its own money and its own contentment to the money and contentment of the other agent, and so require $M(2) > C(2) > \{M(1),C(1)\}$ for acceptance, or $C(2) > M(2) > \{M(1),C(1)\}$ for rejection. If E is preferred to both M(2) and C(2) the second player will also reject the offer, but here motivated by a desire for equality, rather than the insult.

Finally if a zero offer is made we would expect rejection, either because of the insult, or because equality is desired. Indeed a zero offer will only be accepted if the second player prefers the others player's money or contentment to its own contentment: $\{C(1), M(1)\} > C(2)$. This would be an extreme example of altruism, and we would expect it to be rare. These orderings would also lead to acceptance in Q4.

Now consider the first player. The arguments it will consider are shown in Table 3. No argument is proposed to reach q4: if A3 is chosen the reason is that the first player wishes to reach q3. The situation of the first player is considerably more complicated than that of the second player, since there is a much wider range of choice available, and a wider variety of values to promote.

If the first agent is highly altruistic, so that M(2) is its most preferred value, then it should choose A1, since this promotes the other agent's wealth to the greatest extent, and there are no objections resting on M(2). Similarly if the agent prizes equality above all else it should choose A2. One of A1 or A2 should also be chosen if the most important thing is to avoid upsetting the second player, since C(2) may be demoted if another action is selected. If, however, it prefers the feeling of being generous, or its own wealth, or its own contentment, then things become more difficult, because of the uncertainty as to where the other agent will set its threshold. Because A5 will only succeed in promoting M(1) if the second agent prefers M(1) to its own contentment this action will probably be rejected, even by an agent who values only M(1), since the agent cannot take the risk that the other agent will be that altruistic. The agent will therefore be most likely to choose A3, since this is as good as any other with respect to G, and - provided n is chosen correctly - will promote M(1) while not demoting C(2), and risking the demotion of C(1). The exact amount to offer will depend on the agent's view of what will be taken to be a fair offer, and the relative importance it gives to M(1) and C(1). If it prizes C(1) more - seeing the important thing to be the avoidance of regret at not offering enough, n will tend to be higher than the sum the agent would itself accept: if M(1) is preferred then the agent may choose an amount very close to what it would itself regard as acceptable. If G is very important this will also intend to increase the size of the offer, since the higher the offer the more confidence there can be that the action will succeed in promoting G. If an agent had a very strong preference for G or C(1) or both, then the offer might even rise to 50%, since this will ensure that G is promoted and C(1) is not demoted. Thus a cautious agent who prized these values might choose A2, even though equality was not so very important to it. This caution is especially merited if the agent can make no assumptions about the other agent: if both agents come from the same, relatively homogenous, culture they may be able to predict the size of offer that will be expected with more accuracy, and so the reaction to various offers can be more reliably predicted. Cultural effects, particularly

Table 3: Arguments for First Player

ID	Arguments	Objections
F1a	a1 to promote M1	a2, a3, a5 promote M1 more; a2, a3 promotes G as much; a2 avoids demoting C1; a3 may avoid demoting C1 a2 avoids demoting C2; a3 may avoid demoting C2
F1b	a1 to promote M2	
F1c	a1 to promote G	
F1d	a1 to avoid demoting C1	
F1e	a1 to avoid demoting C2	
F2a	a2 to promote M1	a3, a5 promote M1 more; a1 promotes M2 more; a1, a3 promotes G as much; a1 avoids demoting C1; a3 may avoid demoting C1; a1 avoids demoting C2; a3 may avoid demoting C2
F2b	a2 to promote M2	
F2c	a2 to promote G	
F2d	a2 to avoid demoting C1	
F2e	a2 to avoid demoting C2	
F2f	a2 to promote E	
F3a	a3 to promote M1	a5 promotes M1 more; a1, a2 promotes M2 more; a1, a2 promotes G as much; a3 may not promote G; a1, a2, avoids demoting C2; a3 may demote C2
F3b	a3 to promote M2	
F3c	a3 to promote G	
F3d	a3 to avoid demoting C2	
F5a	a5 to promote M1	a1 demotes C2; a1 may demote C1

those from the small scale cultures used in [13] will form the basis of the explanation of diverse behaviour in the next section.

3.1 Explaining the Differences

Using the particular value orderings of individual agents, we can therefore account for the range of behaviour exhibited by subjects in the various Ultimatum game experiments. In all experiments the whole range of behaviours is found, but there are differences in the proportions of people exhibiting the behaviours. The natural explanation of this is that there are cultural differences, which have an effect of the "normal" value ordering in the culture concerned. These cultural differences are explored in [24], which divides the experiments across continents, (with US and Europe further divided into East and West) and [13], which looks at fifteen small scale societies taken from twelve countries on five continents.

In [24] no significant difference in the size of offer was found between their continental groupings. While, when grouped by country, the mean offer varied from 51% in Paraguay down to 26% in Peru, both of these are in the South America group and so tend to cancel each other out. Where the grouping by continent did show regional differences was in the responder's behaviour. Asians had a lower rejection rate than US responders, and Western US rejected less than Eastern US. They were, however, unable to come up with any explanation of this. Four hypotheses were explored based on notions taken from Hofstede [14] on the degree of individualism and the power distance (the expectation and acceptance that power is distributed unequally), and respect for authority, taken from Inglehart [16]. The only hypothesis endorsed was that offers tended to be smaller in a deferential society, although the rejection rate remained the same. There are then, few pointers from this study: perhaps the grouping by continent was not appropriate, and variations are within continent rather than across them.

The results from Heinrich *et al* [13], which looked at smaller, more homogenous, groups are perhaps more interesting. They find that:

- The canonical model is not supported in any society studied;
- There is considerable behavioural variety across societies and the canonical model fails in a variety of ways;
- Group level differences in economic organisation and the extent of market integration explains a substantial portion of the variation - the greater the degree of market integration, the greater the cooperation in the experiments;
- Individual economic and demographic variables do not explain variation;
- Behaviour in the experiments is generally consistent with the economic patterns of everyday life in these societies.

These are interesting conclusions. Two dimensions were considered: how important cooperation was to the economic production of the society, and the degree of market exchange experienced in the daily lives of the society. The lowest offers were made by the Machiguenga people of Peru, whose daily lives involve little or no cooperation in production, exchange or sharing beyond the family unit. This can be contrasted with the three societies making the highest offers, with modes at 50%. The Lamelara of Indonesia, whose mean offer was 56%, are a whaling community, who hunt in large canoes of a dozen or more, and the size of their prey makes sharing obvious. The Ache of Paraguay, it is reported, leave their kill at the edge of camp, pretending to have no success. Their kill is then discovered and meticulously shared by the whole camp. 94% of the Ache made offers above 40% with a mean of 51%. Such behaviour suggests a high degree of trust that the other villagers will behave as expected on behalf of the hunter. The Orma of Kenya related the game to a local institution of village level contributions to public good projects such as schools or roads.

The rejection rates also exhibit interesting variations. The Machiguenga, although making the lowest offers, also have a low rejection rate, rejecting only one in ten offers below 20%. The Lamelara reject no offers, even experimenter generated offers less than 20% (no actual offers this low were made). The highest rejection rates of offers (including offers above 20%) appear in societies ranked around the middle of the dimensions used by Henrich *et al*. Two such groups, the Au and the Gnau of Papua New Guinea rejected not only low offers, but also offers of greater than 50%. From the rejection behaviour, one might conclude that in some societies, like the Lamelara cooperation is simply a way of life, and generous offers are made routinely and accepted routinely. In others, like the Machiguenga, the society is independent, offering little, and hence not resenting being offered little. The Papua New Guineans are given a different explanation in [13]: apparently in these two cultures accepting a gift commits one to reciprocation in the future: members of these societies may well thus reject even good offers to avoid indebtedness. Also in the middle ranking societies, where cooperation is neither essential and natural nor unneeded and unlooked for, the need to maintain the required level of cooperation by punishing low offers becomes greater. In such societies therefore people are likely to be more sensitive to selfish behaviour, and readier to reject a low offer. The highest rejection rate outside of Papua New Guinea came from the Sangu farmers of Tanzania, who rejected 25% of offers, even though only one fifth of these was below 20%. As these example demonstrate, people may have cultural reasons for engaging in or rejecting transactions with particular others, such as members of clan groups and moieties.

The overall conclusion of [13] is that *"the degree of cooperation, sharing and punishment exhibited by experimental subjects closely corresponds to templates for these behaviours in the subjects' daily lives"*, and that *"preferences over economic choices ... are shaped by the economic and social interactions of everyday life"*. In the next section we will discuss the implications for the design of agent systems.

4. IMPLICATIONS FOR MULTI-AGENT SYSTEMS

In designing a system of autonomous agents it is necessary to include some mechanism to enable the agent to motivate a choice between the various candidate actions available in a given situation. This has often been done using a quantitative approach with a utility function (possibly multi-dimensional) to determine expected utility, which the agent can then attempt to maximise. We have proposed a qualitative alternative, in which the agent determines which of its values will be promoted and demoted by the available actions, and then chooses by resolving the competing justifications by reference to an ordering of these values.

The role of the systems designer is thus to consider how best these values should be chosen. The Ultimatum game suggests that one rationale can be provided by the degree of cooperation and economic interaction that is involved in the agent system.

First consider a relatively closed multi agent system in which agents interact with other agents from outside relatively little. The agents may have specialist roles within the system, but they are in fixed and stable relationships, rather akin to a subsistence family group. When such agents need, for example, to compete for a shared resource, the above discussion would suggest that they can be effective on a simple model of maximising their own economic utility. In such a situation it may well be that pure market forces will be able to determine an efficient allocation, and that the canonical economic model is the one to use. In our terms this would emphasise the value M for the proposer and minimise the value C for the responder (a *selfish order*). Another reason why we might expect market forces to be more appropriate to this kind of system is that certain other fundamental assumptions, such as perfect, freely available, information, are more easily satisfied. Closed, single-company, multi-agent systems, such as the telecommunications services maintenance system described in [21], may be viewed as examples of such systems.

At the other end of the spectrum, in some very open multi agent systems it may be that cooperation with unfamiliar agents encountered in a flexible way not determinable in advance, is essential. For example it may be necessary that the operation of such a system requires the agent to gather information from other agents and to share information with them in turn. In such a system the generous and tolerant attitudes of the Lamelara and Ache may prove beneficial. This would be achieved by emphasising the value of E and being relatively indifferent towards which agent M was promoted in respect of (an *altruistic* order). An example of such a multi-agent system could be the multi-agent vehicle traffic insurance claims system of [1], were it to extend beyond the small group of initial companies involved in the project.

Between these extremes there are many multi-agent systems where agents require interaction with unfamiliar agents on an irregular basis, and it is necessary for agents to be able to do business with one another, but where it is proper that some price be demanded. In such applications, the notion of punishment becomes important: it is necessary that agents are kept aware of their responsibilities towards one another, while allowing them to pursue their own interests to a reasonable extent. This would suggest an emphasis on C for the responder and tempering the estimation of M by the use of C by the proposer, to ensure that the agent was sufficiently sensitive to the possibility of punishment for this to be effective (a *responsive* order). Examples of such multi-agent systems may be air-traffic resource allocation systems where prices are used to allocate airport gate access, as in the multi-agent system of [17].

Variation in agents is desirable since the tasks they perform and the contexts in which they operate vary. Different mechanisms will give different behaviours and these will be appropriate to different situations. The use of decision making using value based argumentation offers a relatively simple way in which these differences can be incorporated, and one which can be related to the empirical observation of various human societies. But in order to make use of these variations we have to abandon the idea that there is a single right answer to what agents should want, and instead be prepared to draw inspiration from the diversity of cultures that have emerged in human societies.

The transparency of the explanation also provides distinct advantages when agents are able to discuss their behaviour with each other. The ability to provide explanations and arguments justifying the acceptance or rejection of offers and bids has been found very beneficial in negotiation [26]. Suppose an agent were to reject an offer of a third of the available amount: it could be of future use to know whether this was done because the agent was insulted by the size of the offer, or because of a desire to ensure that both players received the same amount. In the one case the proposer would have no reason to increase the offer (unless he was prepared to offer half), whereas in the other he should perhaps recognise that a larger offer would be more acceptable in the community of agents in which he finds himself. By receiving this kind of information about the value ordering of the other agents, the proposer is better able to predict what the other agents will do. It may also be desirable for the agent to modify its own value order so that its decisions are not continually frustrated by the choices of other agents. In this way, one might see a kind of cultural assimilation. A similar idea, although essentially based on an underlying utility function, can be found in [19]. We believe that our approach will be able to achieve similar benefits, but in a more straightforward and transparent manner.

5. CONCLUDING REMARKS

In this paper we have presented an argumentation-based qualitative decision mechanism for agents facing a choice of actions, and used this mechanism to model observed empirical behaviour by people playing the Ultimatum Game in experimental situations. Our decision-mechanism explicitly recognises alternative values which different actions may promote or demote, and, by assuming an ordering over such values, enables a decision-making agent to rank the possible actions. Because the behaviour of people playing the Ultimatum Game observed in experimental situations is not explained by the standard model of an economic decision-maker as a selfish utility-maximiser, our approach provides a novel and coherent interpretation of this empirical data. As such, our approach also provides a new decision-mechanism for agent designers wishing to implement actual human-decision-making processes.

There are several advantages of our new approach. Firstly, our approach enables a coherent account to be given of the empirical evidence which allows for inter-cultural and inter-temporal differences in behaviours; in particular, our approach does not require the labelling of actual human behaviours as *"irrational"* when these behaviours are inconsistent with the prevailing theoretical model [20]. As such, our framework therefore provides a more general model of human decision-making than the selfish utility-maximiser model. Because our model is qualitative, it is also more general than quantitative models which incorporate social welfare into the utility function of a decision-maker, as in [15]. Secondly, our argumentation-based approach can, unlike an approach based solely on quantitative utility, explain the framing effects seen in the cited experiments such as [5]. If the choices made by experiment participants depended only on estimating the utility of the state reached, we should expect the same participants to choose the same outcome whatever their initial position; in the argumentation approach, in contrast, the arguments available to a decision-maker depend on the initial state as well as on the target state.

Thirdly, the proposed argumentation approach is transparent, since reasoned justifications for selected decision-options are automatically generated: giving explanations and justifications of the choices in terms of arguments is more informative and more open to discussion and criticism than referring to a formula for the utility function which can only be obtained by fitting the function to the choices made. For software agents which need to explain their action-recommendations to their human principals and for human principals who wish to guide their subordinate software agents, this is a very important feature. This transparency also makes clear the tradeoffs involved when there are competing values, in a manner absent from multi-dimensional utility maximizing models. Finally, our approach allows for greater flexibility in agent design. A rational agent may be required to adopt a different value-order at different times, or for different types of interactions with other agents, or for interactions with different types of agents, or in different decision-making contexts. Our framework allows software engineers to create agents capable of such adaptability, in place of narrow-minded agents capable only of maximizing numerical values.

The approach here opens several lines of empirical investigation. One possibility - which should be of interest to experimental economists - would be to determine a value order appropriate to a particular culture, either by looking at their everyday way of life, or by direct elicitation, and simulate the playing of the Ultimatum Game by a group of such agents using this value ordering, with some probabilistic deviation. The results of the simulation could then be compared with that found in Ultimatum Game experiments within that culture. Such an experiment would confirm the explanatory possibilities of our approach. A more agent orientated line of enquiry would be to apply the three kinds of value order to examples of the three kinds of system described above. The hypothesis would be that the performance of, for example, a closed system would be optimised by agents with a selfish order, while a very open system would perform best with agents using an altruistic order. Such an experiment would go a long way to demonstrating the effectiveness of this approach to agent design. A third line of enquiry would be to apply this approach to further, more sophisticated experiments, such as the Coloured Trails Framework [12], employed in agent systems to investigate agent reasoning about other agents in negotiation contexts [9], [10]. Finally the transparency of the model permits investigation of adaptation of behaviour through reordering of values in the light of success and failure.

We believe that we have presented an approach to modelling behaviour in certain experimental economics scenarios that permits their ready and transparent simulation in agent systems. We believe that our model will facilitate the transfer of considerations determining human decisions to agent systems. Moreover we suggest that these considerations can help to identify the behaviour which is appropriate to different styles of multi agent system.

6. REFERENCES

[1] C. Aart, K. Marcke, R. Pels, and J. Smulders. International insurance traffic with software agents. In F. Harmelen, editor, *Proceedings of the Fifteenth European Conference on Artificial Intelligence (ECAI 2002)*, pages 623–627, Lyon, France, 2002.

[2] K. Atkinson and T. J. M. Bench-Capon. Practical reasoning as presumptive argumentation using action based alternating transition systems. *Artificial Intelligence*, 171(10–15):855–874, 2007.

[3] K. Atkinson and T. J. M. Bench-Capon. Value-based arguments in the dictator game. In *Proceedings of the Fourth Multidisciplinary Workshop on Advances in Preference Handling (M-PREF 2008)*, pages 1–6. AAAI Press, Technical Report WS-08-09, 2008.

[4] K. Atkinson, T. J. M. Bench-Capon, and P. McBurney. Computational representation of practical argument. *Synthese*, 152(2):157–206, 2006.

[5] N. Bardsley. Dictator game giving: Altruism or artefact. *Experimental Economics*, 11(2):122–133, 2008.

[6] T. J. M. Bench-Capon. Persuasion in practical argument using value based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–48, 2003.

[7] G. E. Bolton, E. Katok, and R. Zwick. Dictator game giving: Rules of fairness versus acts of kindness. *International Journal of Game Theory*, 27:269–299, 1998.

[8] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.

[9] S. G. Ficici and A. Pfeffer. Modeling how humans reason about others with partial information. In *Proceedings of AAMAS 2008*, pages 315–322, 2008.

[10] S. G. Ficici and A. Pfeffer. Simultaneously modeling humans' preferences and their beliefs about others' preferences. In *Proceedings of AAMAS 2008*, pages 323–330, 2008.

[11] R. Forsythe, J. L. Horowitz, N. E. Savin, and M. Sefton. Fairness in simple bargaining experiments. *Games and Economic Behavior*, 6(3):347–369, 1994.

[12] B. J. Grosz, S. Kraus, S. Talman, B. Stossel, and M. Havlin. The influence of social dependencies on decision-making: Initial investigations with a new game. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)*, pages 782–789, 2004.

[13] J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, and R. McElreath. In search of homo economicus. behavioral experiments in fifteen small-scale societies. *American Economic Review*, 91:73–78, 2001.

[14] G. Hofstede. *Cultures and Organizations: Software of the Mind*. McGraw-Hill, New York, USA, 1991.

[15] L. M. J. Hogg and N. R. Jennings. Socially intelligent reasoning for autonomous agents. *IEEE Transactions on Systems, Man and Cybernetics – Part A*, 31(5):381–399, 2001.

[16] R. Inglehart. Culture and democracy. In L. E. Harrison and S. P. Huntington, editors, *Culture Matters: How Values Shape Human Progress*. Basic Books, New York, USA, 2000.

[17] G. M. Jonker, F. P. M. Dignum, and J.-J. Meyer. Achieving cooperation among selfish agents in the air traffic management domain using signed money. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2007)*, pages 1258–1260. ACM Press, 2007.

[18] D. Kahneman, P. Slovic, and A. Tversky, editors. *Judgment under Uncertainty: Heuristics and biases*. Cambridge University Press, Cambridge, UK, 1982.

[19] N. C. Karunatillake, N. R. Jennings, I. Rahwan, and S. D. Ramchurn. Managing social influences through argumentation-based negotiation. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2006)*, pages 426–428, 2006.

[20] O. Lange. The scope and method of economics. *The Review of Economic Studies*, 13(1):19–32, 1945-1946.

[21] P. Mihailescu, J. Shepherdson, P. Marrow, L. Lee, and H. Lee. MAS platforms as an enabler of enterprise mobilisation: The state of the art. In *Proceedings of the IEEE International Conference on Systems, Man & Cybernetics (SMC 2004)*, volume 2, pages 1889–1894, The Hague, The Netherlands, 2004. IEEE.

[22] J. S. Mill. On the definition of political economy, and on the method of investigation proper to it. In *London and Westminster Review, October 1836. Essays on Some Unsettled Questions of Political Economy*. Longmans, Green, Reader and Dyer, London, 1874.

[23] M. A. Nowak, K. M. Page, and K. Sigmund. Fairness versus reason in the ultimatum game. *Science*, 289:1773–1775, 2000.

[24] H. Oosterbeek, R. Sloof, and G. van de Kuilen. Differences in ultimatum game experiments: Evidence from a meta-analysis. *Experimental Economics*, 7:171–188, 2004.

[25] S. Parsons, P. Gmytrasiewicz, and M. Wooldridge, editors. *Game theory and Decision Theory in Agent-based Systems*. Kluwer Academic, Boston MA USA, 2002.

[26] S. D. Ramchurn, C. Sierra, L. Godo, and N. R. Jennings. Negotiating using rewards. *Artificial Intelligence*, 171(10-15):805–837, 2007.

[27] D. N. Walton. *Argumentation Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates, Mahwah, NJ, USA, 1996.

[28] 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.