Autonomous Agents and Multi-Agent Systems manuscript No.
(will be inserted by the editor)

Using Argumentation to Model Agent Decision Making in Economic Experiments

Trevor Bench-Capon · **Katie Atkinson** · **Peter McBurney**

Received: date / Accepted: date

Abstract In this paper we demonstrate how a qualitative framework for decision making can be used to model scenarios from experimental economic studies and we show how our approach explains the results that have been reported from such studies. Our framework is an argumentation-based one in which the social values promoted or demoted by alternative action options are explicitly represented. Our particular representation is used to model the Dictator Game and the Ultimatum Game, which are simple interactions in which it must be decided how a sum of money will be divided between the players in the games. Studies have been conducted into how humans act in such games and the results are not explained by a decision-model that assumes that the participants are purely self-interested utility-maximisers. 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. In this paper we show how these interactions can be modelled in agent systems in a framework that makes explicit the reasons for the agents' choices based upon their social values. Our framework is intended for use in situations where agents are required to be adaptable, for example, where agents may prefer different outcome states in transactions involving different types of counter-parties.

Keywords Argumentation · Decision making · Empirical studies of behaviour

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

Trevor Bench-Capon and Katie Atkinson
Department of Computer Science,
University of Liverpool, UK.
E-mail: {tbc, katie}@liverpool.ac.uk

P. McBurney
Department of Informatics,
King's College London, UK.
E-mail: peter.mcburney@kcl.ac.uk

theory, is that humans act so as to maximise their satisfaction, well-being or utility. The idea was stated succinctly by John Stuart Mill [21]:

> *"[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* (see e.g. [9]) and the *Ultimatum Game* (see 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 [9], 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.* [12], carried out over fifteen small-scale societies in twelve countries over 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, may confer 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 [5, 7, 12, 23, 24]. Indeed, there may even be experimental effects in which experiment participants enact the roles they expect participants of social science experiments to play [5]. A number of experiments have attempted to isolate or control for these various factors in an effort to confirm or disconfirm their influence, e.g., [5, 7].

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 [12, 24]. 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 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 provides a computational explanation of the evidence from these empirical studies in terms of general societal values discussed in the papers, which can thus be directly related to the way in which economic decisions are debated in press and politics. 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 self-interested 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., [27]. However, human beings are not always utility-maximisers, as revealed by evidence in experimental and behavioural economics [17]. The invariable deviance from the canonical model in these experiments suggests that there must be beneficial effects from so deviating. While there are considerable variations, suggesting that there is room, or need, for diversity, it is striking that the canonical model is never followed in these experiments. 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 one such model in this paper.

In addition, human principals owning or directing intelligent software agents may well desire their agents to be able to explain their choices and actions, either prospectively or retrospectively. This requirement for software systems to provide explanations for their behaviors is qualitatively different, and vastly more challenging, than merely asking that a system provides a good model for empirical reality. A common view within economics, due originally to Milton Friedman [11], is that goodness-of-fit of models should be the sole criterion for their assessment and deployment, and not their human comprehensibility, their plausibility, or even the reasonableness of their assumptions and internal logic. Whether this viewpoint is appropriate for Economics or not is not for us to say. However, the goodness-of-fit criterion is certainly insufficient for Artificial Intelligence — the creation of intelligent machine entities — since machine entities unable to explain their actions will provide little confidence to their human owners, or to those others (whether human or machine) who may have to interact with them. The argumentation-based framework we present in this paper provides such explanatory capabilities.

In section 2 we discuss a classical approach to decision making in economics, the maximum-expected utility (MEU) model, and some of the criticisms that have been put forward against this. This discussion motivates our argumentation-based approach to modelling agent decision making that we present in section 3, and in section 4 we show how

[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.

the Dictator Game can be formulated in terms of our approach. In section 5 we discuss how the presentation of the game can affect the results, and we show how this can be explained in terms of our model of the problem. In section 6 we use our approach to model the Ultimatum Game, and in section 7 we discuss how our findings might be used in the design of multi-agent systems. In section 8 we make some remarks about how cultural differences could be used to explain different behaviours in Prisoner's Dilemma type games, and we conclude the paper in section 9.

2 Beyond Simple Utility Maximisation

The most commonly-taught model of decision making in economics and business studies is the maximum-expected utility (MEU) model, developed by Leonard Savage and others in the middle of the twentieth century (see, for example, [30, 19]). This model assumes that a decision-maker knows a list of alternative action options, and can determine the consequences of these actions, along with his or her personal utility for each consequence. Which consequences accrue to which actions may depend on various uncertain events, which are also assumed known by the decision-maker and for which probabilities of occurrence can also be determined by him or her. The decision-maker then calculates the utility which will arise from the successful execution of each action, weighted by the probabilities of those uncertain events believed to impact the consequences of that action. The decision-maker then chooses that action whose expected utility (i.e., utilities weighted by probabilities) is greater than the expected utilities of all other action options.

From its inception, the MEU model has been criticised on many grounds, for example: that it makes assumptions about human behaviour which have been empirically refuted (e.g., Shafer [33]); that it is inaccurate as a description of the processes human decision-makers actually use to make decisions (Shackle [32]); that it does not provide a basis in terms of reasons for the selection of action-options (Anand [2]); that it assumes decision-makers have infinite processing capabilities, time and other resources (Simon [35]); that it assumes decision-makers consume no resources when applying it, such as the so-called *costs of thinking* (Shugan [34]); that decision-makers typically lack the quantitative information (the probabilities and utilities) needed to apply it; and that, by not treating catastrophic outcomes separately from averaged (expected) outcomes, it potentially leads decision-makers greatly astray (Shackle [32]). In response, refinements of the MEU model have been proposed, for example to allow for the limited processing capabilities of decision-makers; the various *Consideration Set* models in consumer marketing are examples of these [29]. A widespread response to the typical absence of quantitative probability and utility data has been to develop various *qualitative* models of decision-making, such as that in [10]; see Parsons [26] for a review of these qualitative models. Foundational axioms for decision theory alternative to those underlying Savage's MEU model have also been proposed, for example, by Paul Anand [2].

As it stands, the MEU model does not explain the results of the Dictator and Ultimatum Game experiments, since some participants seem to be selecting actions which do not maximise their individual expected utilities. One way of accommodating these results from Dictator and Ultimatum Game experiments would be to retain the idea that each agent maximises their individual utility, but to assume that some part of this individual utility arises from the welfare of other agents, perhaps in accordance with various cultural norms. For example, the model proposed by Bolton and colleagues [7, Section 7.1] for the Dictator Game assumes that participants have in mind a "fair share" allocation for a dictator, and only maximize their individual welfare within this share; the non-zero remainder not allocated to themselves is assigned to other agents. In another domain, Hogg and Jennings [14] implement such social utility functions, comprising both individual and social welfare components, for agents engaged in civil emergency management. However, the alternative approach we will use here is a qualitative one based the use of arguments that justify and challenge proposals for action. 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 [6] developed in the argumentation in AI community based on a ranking of the various motivating factors, which can of course vary across individuals and cultures.

There are a number of advantages to our approach. To begin with, 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. Our approach provides a reasoned justification for the options selected by a decision-maker, and thus enables analysis of the underlying causes of the differences between different decision-makers. Reliance on quantitative utility calculations, in contrast, enables an observer only to conclude that one decision-maker had a higher utility for a certain outcome than did another decision-maker, not *why* he or she had this higher utility, nor how such differences between decision-makers could be eliminated or resolved through discussion or analysis. Because of this explanatory capability our approach satisfies the requirement that a decision-making model support reasoned decision-selection as specified by, for example, the *Minimal Prudence Axiom* of Anand's foundational axioms for rational decision theory [2, Chapter 8, Axiom 4]. As mentioned in the previous section, the explanatory capability of our framework is a key feature for the design of intelligent software entities intended to work for or with human beings. Furthermore, and importantly, the argumentation framework can explain the framing effect. If the choice only depended on estimating the utility of the state reached, we should expect the same individuals to choose the same outcome whatever the initial position. In contrast, in the argumentation approach, the available arguments depend on the initial state as well as the state that is reached. 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. Note that what we provide is a mechanism which can be directly related to the kinds of explanations that people actually advance during debates about decisions concerning what to do. Additionally, using techniques to compare and evaluate arguments based on values (e.g. [6], [22]) allows arguments about values preferences to be explicitly stated.

3 The Argumentation Approach

Our 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 [36] 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 [3] and describe the problem as an Action-based Alternating Transition System (AATS) [37]. 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; the full formal details can be found in [3]:

- 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 R) 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 (i.e. performing 'ab', where agent 1 does *a* and agent 2 does *b*) so as to move to q2 promotes V1, and acting (i.e. performing 'ba', where agent 1 does *b* and agent 2 does *a*) 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.

Instantiating the argument scheme and the critical questions gives rise to a set of conflicting arguments. Counter-arguments respond to arguments, seeking to undercut, rebut, or

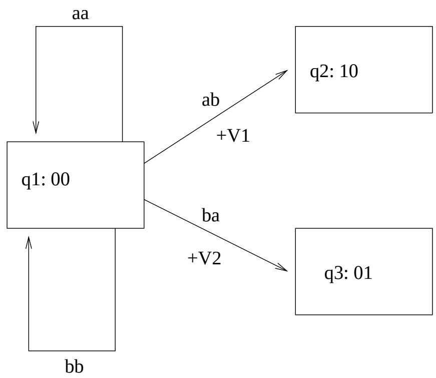

Fig. 1 Simple AATS

otherwise attack the argument being countered. The distinction is standard within the argumentation literature, and corresponds to ordinary usage of these words. Once the conflicting set of arguments 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. The key feature of VAFs is that they allow a distinction to be made between successful attacks (defeats) and unsuccessful attacks, on the basis of societal values associated with the arguments. VAFs, and some associated notions, are formally defined in Definitions 1 and 2.

Definition 1 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, where, as in [8], X is a set of arguments and A is a binary attack relation between pair of arguments, $\nu = v_1, v_2, ..., v_k$ a set of *k values*, and η :X→ ν a mapping that associates a value $\eta(x) \in \nu$ with each argument $x \in X$. A *specific audience*, α, for a VAF $\langle H(X,A), \nu, \eta \rangle$, 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 Let $\langle H(X,A), V, \eta \rangle$ 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: $\langle x,y \rangle \in A$ and it is not the case that $\eta(y) \succ_\alpha \eta(x)$.
b. An argument x is *acceptable* w.r.t. 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 $\langle x,y \rangle \in R \times R$, either $\langle x,y \rangle \notin A$ 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. As indicated above, 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 resulting framework evaluated as a standard AF, the preferences of the particular audience having been accounted for by the removal of the unsuccessful attacks. 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.

In the following sections we will apply this approach to the Dictator Game and the Ultimatum Game to give fully worked examples of how arguments for action within these games can be generated and evaluated.

4 Modelling the Dictator Game

In this section we show how we can use the approach to argumentation presented in the previous section to model the Dictator Game. We begin by considering the problem formulation and the states in the set Q that can be reached. We will represent a limited number of options that comprise the set of actions A and assume 100 needs to be divided. The set A thus comprises the following five actions corresponding to the five different divisions of the money, as shown in Table 1.

Table 1 Actions corresponding to the division of money

Label and action
a1 = give(70%)
a2 = give(100%)
a3 = give(50%)
a4 = give(0%)
a5 = give(30%)

We also include, to begin with, as the initial state *q0*, the state where neither have any money, as this is the situation of the subjects at the start of the experiment. The dictator thus starts in the state {0,0} and can move to any of the other states, as shown in Figure 2. We then need to attach values to the transitions. We will use a range of values representing the various motivations that have been discussed in the literature.

Money: Most obvious is the value of money. This is what economic man is supposed to maximise. Because, however, we need to recognise the idea that the other player having money may be considered positively by the dictator, we distinguish money for the dictator himself (MS) from money for the other (MO).

Giving (G): It can be held that giving a gift is a source of pleasure, and this is what motivates the dictator to share.

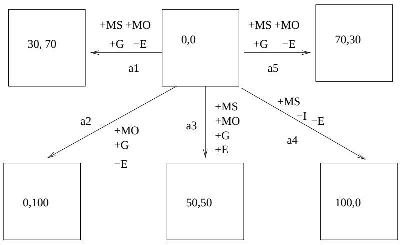

Fig. 2 Values promoted or demoted by transitions from $\{0{,}0\}$

Image (I): Another consideration is the desire not to appear mean before the experimenter that motivates sharing. It could even be that one does not want to appear mean to oneself.
Equality (E): Equality, as defined by an equal distribution, characterises a sense of fairness.

These values will be promoted and demoted as follows: MS will be promoted if the dictator's money increases and demoted if it is reduced. MO will be promoted if the other's money increases and demoted if it is reduced. G will be promoted if the other's money increases as a result of the dictator's decision. Value I will be demoted if the dictator keeps all the money. E will be promoted if the distribution becomes more equal, and demoted if a distribution becomes less equal. These transitions are also shown in Figure 2. Note that in the case where the dictator chooses not to give any money, MO is not demoted (rather it remains neutral) since the second player is not losing any money (similarly for G). However, as we discuss in section 5, the values promoted and the results of the experiment can change depending upon the set up of the game, for instance, when a player can take money from the other, rather simply receiving some.

At this point we wish to distinguish between two kinds of value. With some values more is always better: the object is to maximise the promotion of the value. For other values, however, more is not always desired. For example there may be pleasure in giving, but this pleasure may not be increased if the gift is larger. Thus with respect to each of the values the agent may be what we will call a *maximiser* or a *satisficer*. In our usage, which deviates slightly from standard usage, the difference between these is that if we have an argument "I should do a1 to promote V1" then the counter argument "a2 would promote V1 to a greater extent" is possible for a value to be maximised, but not for a value to be satisficed.

We can now consider the arguments for the various actions in Figure 2. We omit reference to the current and target states, and the goal, as these are obvious from the context. The promotion of the various values gives rise to the following instantiations of AS1[2] in favour of the possible actions.

Arg1.1: We should do a1 to promote MS
Arg1.2: We should do a1 to promote MO
Arg1.3: We should do a1 to promote G
Arg2.1: We should do a2 to promote MO
Arg2.2 :We should do a2 to promote G

[2] We assume that the single reasoning agent in this scenario has perfect information about the representation of the problem scenario and the effects of actions. Thus, instantiations of AS1 need not consider doubts as to the current circumstances, nor the goal states achieved by performing the actions, since these are not under dispute. Hence, the following arguments comprise only the action choices and the values promoted by the actions.

Arg3.1: We should do a3 to promote MS
Arg3.2: We should do a3 to promote MO
Arg3.3: We should do a3 to promote G
Arg3.4: We should do a3 to promote E
Arg4.1: We should do a4 to promote MS
Arg5.1 We should do a5 to promote MS
Arg5.2: We should do a5 to promote MO
Arg5.3: We should do a5 to promote G

If we were to choose simply on the basis of number of arguments in favour, we would choose a3. However, we need to consider the counter arguments. Let us first assume that the agent will wish to maximise rather than satisfice money, but wishes to satisfice giving rather than maximise it[3]. Note that the question does not arise with equality or image which are each taken to be an all or nothing thing here. The critical questions that are applicable here are:

1. Would a different action promote the value (of the maximiser) to a greater extent?
2. Would a different action also promote the value (of the satisficer)?
3. Would the action demote some other value?
4. Would the action preclude the promotion of some other value?

We then have the following objections giving rise to counter arguments to a3:

Obj3.1 a4 would promote MS more than a3
Obj3.2 a5 would promote MS more than a3
Obj3.3 a2 would promote MO more than a3
Obj3.4 a1 would promote MO more than a3
Obj3.5 a5 is as good as a3 with respect to G
Obj3.6 a2 is as good as a3 with respect to G
Obj3.7 a1 is as good as a3 with respect to G

Similarly we have arguments to counter the arguments for other actions:

Obj1.1 a3 would promote MS more than a1
Obj1.2 a4 would promote MS more than a1
Obj1.3 a5 would promote MS more than a1
Obj1.4 a2 would promote MO more than a1
Obj1.5 a1 would demote E
Obj1.6 a5 is as good as a1 with respect to G
Obj1.7 a3 is as good as a1 with respect to G
Obj1.8 a2 is as good as a1 with respect to G
Obj2.1 a2 would demote E
Obj2.2 a5 would promote G as well as a2
Obj2.3 a3 would promote G as well as a2
Obj2.4 a1 would promote G as well as a2
Obj2.5 a2 precludes the promotion of MS
Obj4.1 a4 would demote E
Obj4.2 a4 would demote I
Obj4.3 a4 precludes the promotion of MO

[3] This seems to conform with experiments in [7], in which dictators would give trivial amounts rather than nothing.

Obj4.4 a4 precludes the promotion of G
Obj5.1 a4 would promote MS more than a5
Obj5.2 a1 would promote MO more than a5
Obj5.3 a2 would promote MO more than a5
Obj5.4 a3 would promote MO more than a5
Obj5.5 a5 would demote E
Obj5.6 a1 would promote G as well as a5
Obj5.7 a2 would promote G as well as a5
Obj5.8 a3 would promote G as well as a5

We can now determine the attack relation between these arguments. Note that some of the objections urge the performance of a particular action. These are thus attacked by arguments against that action. Note also that we characterise as attacks objections that point out the promotion of values by other actions. As described in [4], this is in line with reasoning embodied in the critical questions that raise objections to instantiations of AS1 intended to raise for consideration alternative potential actions to perform. The attack relation between the above arguments is summarised in Table 2. Arguments are annotated by their values (S for MS and O for MO).

Table 2 Attack relation applicable in {0,0}

Action	For	Against
a1	Arg1.1S, Arg1.2O, Arg1.3G, Obj2.4G, Obj3.4O, Obj3.7G, Obj5.2O, Obj5.6G	Obj1.1S, Obj1.2S, Obj1.3S, Obj1.4O, **Obj1.5E**, Obj1.6G, Obj1.7G, Obj1.8G
a2	Arg2.1O, Arg2.2G, Obj1.4O, Obj1.8G, Obj3.3O, Obj3.6G, Obj5.3O, Obj5.7G	**Obj2.1E**, Obj2.2G, Obj2.3G, Obj2.4G, **Obj2.5S**
a3	Arg3.1S, Arg3.2O, Arg3.3G, Arg3.4E, Obj1.1S, Obj1.7G, Obj2.3G, Obj5.4O, Obj5.8G	Obj3.1S, Obj3.2S, Obj3.3O, Obj3.4O, Obj3.5G, Obj3.6G, Obj3.7G
a4	Arg4.1S, Obj1.2S, Obj3.1S, Obj5.1S	**Obj4.1E**, **Obj4.2I**, **Obj4.3O**, **Obj4.4G**
a5	Arg5.1S, Arg5.2O, Arg5.3G, Obj1.3S, Obj1.6G, Obj2.2G, Obj3.2S, Obj3.5G	Obj5.1S, Obj5.2O, Obj5.3O, Obj5.4O, **Obj5.5E**, Obj5.6G, Obj5.7G, Obj5.8G

Note that some objections do not urge another action. These, emboldened in Table 2, are interesting because as unattacked arguments in the argumentation framework they cannot be defeated, so their attack can fail only in virtue of a value preference. In the other cases, the attacks are symmetric.

Now let us consider what this says about the preference ranking of agents who choose particular actions. Consider first the simplest case, a4. In this case all the arguments in favour are based on MS, and each of the other four values give rise to a distinct argument against. Hence an agent selecting a4 must have a preference ranking corresponding to the partial order in which MS is the most preferred value. To justify a2, the preference order must include MO greater than both MS and E. G cannot justify doing a2, since other actions are

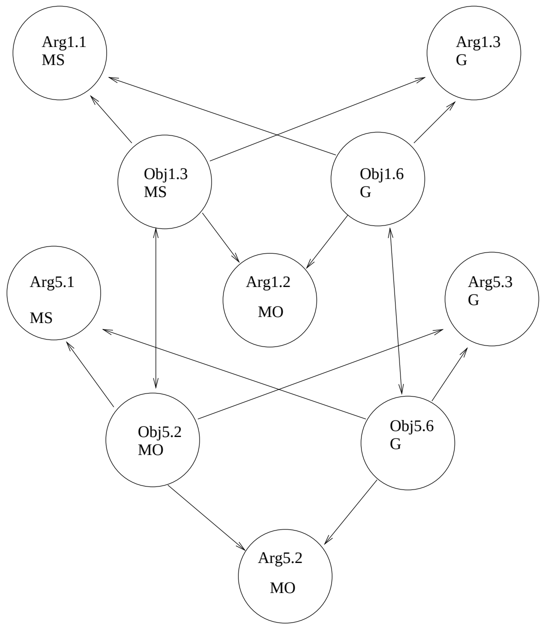

Fig. 3 VAF for arguments relating to a1 and a5.

equally good with respect to this satisficer. Turning to a3, we can see that E must be preferred to all of {MS, MO and G}, since these three values would be better or equally well promoted by other actions. This leaves a1 and a5. Clearly in both cases one of MS, MO and G must be preferred to E. This already means that a3 is defeated. Similarly a preference for G, I or MO over MS will defeat a4. So we need only consider the conflict between a1 and a5. The VAF for the arguments relevant to the choice between a1 and a5 is shown in Figure 3. The VAF is depicted as a directed graph with nodes denoting arguments and edges showing the attack relations between the arguments. The nodes are annotated with the identifying label for the arguments and objections, as well as the value promoted by each of the arguments. Where arrows are doubled headed this denotes a symmetric attack between two arguments. The VAF in Figure 3 contains the three arguments in favour of a1, the three arguments in favour of a5, the two objections against a1 based on a5, and the two objections against a5 based on a1. The nodes contain the labels of the arguments and objections as introduced above and the values associated with the arguments and objections.

We must now evaluate the arguments to determine which of them survive the attacks on them and hence which action is selected. Note first that a preference for G cannot discriminate between the actions a1 and a5. This is because the two cycle between Obj1.6 and Obj5.6 means that these two arguments will each give rise to a distinct preferred extension. It is desirable that G cannot be used to discriminate since both the actions promote G and it is only a satisficer. Thus if MS $>$ MO we will perform a5 and if MO $>$ MS we will select a1. G may be the most preferred value, but the preference between MS and MO will determine the action selected.

Recall now that one of MO, G and I must be preferred to MS in order to defeat a4. Suppose we now make the reasonable assumption that MS is preferred to MO. This means that if MS is the most preferred value, a4 will be selected. If I or G is preferred to MS but

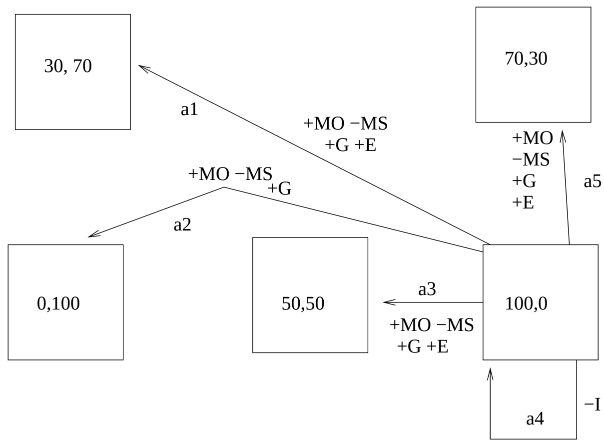

Fig. 4 Values promoted and demoted by transitions from {100,0}

MS is preferred to E, then a5 will be chosen. Otherwise a3 will be chosen. Note that one of these three choices is made by almost all subjects in the experiments.

5 Framing

One claim made about the Dictator Game is that the results are highly sensitive to the way the task is presented to the subjects [5]. Bardsley writes:

> *"Experimental dictator games have been used to explore unselfish behaviour. Evidence is presented here, however, that subjects' generosity can be reversed by allowing them to take a partner's money. Dictator game giving therefore does not reveal concern for consequences to others existing independently of the environment, as posited in rational choice theory. It may instead be an artefact of experimentation."*

If the framing is to have an impact, the choice of action cannot be determined by the expected utility of the target state, since the utilities of the states are unchanged and the effects of action certain. In our approach, however, the justification of an action does also depend on the current state.

In the above previous section we took {0,0} as a starting point. This can be rationalised because the subjects come to the game with nothing. But the initial state could be seen as {100,0}, or as explicitly presented in the Taking Game[4], as {0,100}. Let us consider this from our argumentation perspective. Figure 4 has {100,0} as the initial state.

Now the arguments change significantly. In particular notice that in Figure 4 E can now be promoted by a5 and a1, as well as a3. In particular, if E is a satisficing rather than a maximising value, this should lead to more agents selecting a5. Note also that there are now no positive arguments in favour of a4: however, the demotion of MS provides arguments against all the actions other than a4 based on actual demotion of MS rather than just not maximising it.

Now consider the Taking Game, as shown in Figure 5.

[4] The Taking Game presented here is a simplification of the game used in [5].

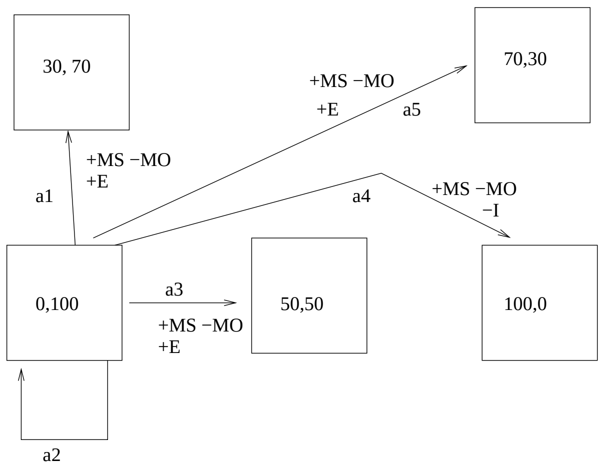

Fig. 5 Values promoted and demoted by transitions from {0,100}, the Taking Game.

The most significant factor here is that G can no longer be promoted at all. This may help to explain Bardsley's finding "that subjects' generosity can be reversed by allowing them to take a partner's money". Notice also that E can be promoted by both a1 and a5, giving a positive non-selfish motivation, and hence more acceptable justification, for taking at least some of the other's allocation. Both of these aspects would lead to the less generous behaviour observed.

Another aspect of the framing problem is that it might affect the preference ranking on values. For example, in the Taking Game, the dictator is very likely to perceive the initial situation as unfair, which could give greater prominence to, and result in a greater preference for, E. In the game starting from {100,0} the sense of giving is emphasised, and thus G may well be ranked more highly by subjects with this perception of the initial situation. Such influences on value preferences can be modelled using arguments about value preferences, as in [22], thus providing an explanatory justification rather than reference to numerical values.

In order to validate the hand-drawn model of the Dictator Game and the Taking Game that we have produced using our argumentation techniques, an implementation has been completed to enable simulation of the reasoning in the scenario. The implementation is a Java program that enables an AATS to be constructed and instantiated to model particular scenarios. Given an instantiated AATS, the program can then automatically generate all the arguments and critical questions that can be produced, and subsequently organise them into a VAF and evaluate the arguments to produce the preferred extensions for given audiences. This implementation has validated our representation of the scenarios using an AATS and shown that the reasoning produced matches our hand-drawn analyses. Full technical details of the implementation can be found in [25]. Currently, the implementation is only able to simulate scenarios where the reasoning involves a single agent, though work is presently underway to produce an implementation that can handle scenarios where the reasoning of a second agent is also modelled within the AATS, as in the Ultimatum Game, which we discuss in detail in the next section.

6 Modelling the Ultimatum Game

In this section we apply our argumentation approach to the Ultimatum Game which differs from the Dictator Game in that the second player has an option of whether to accept or reject the offer made. Again, we must begin by constructing 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 3.

Table 3 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 6.

Now we must identify some values and the transitions which promote and demote them. First there is 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 again take from the

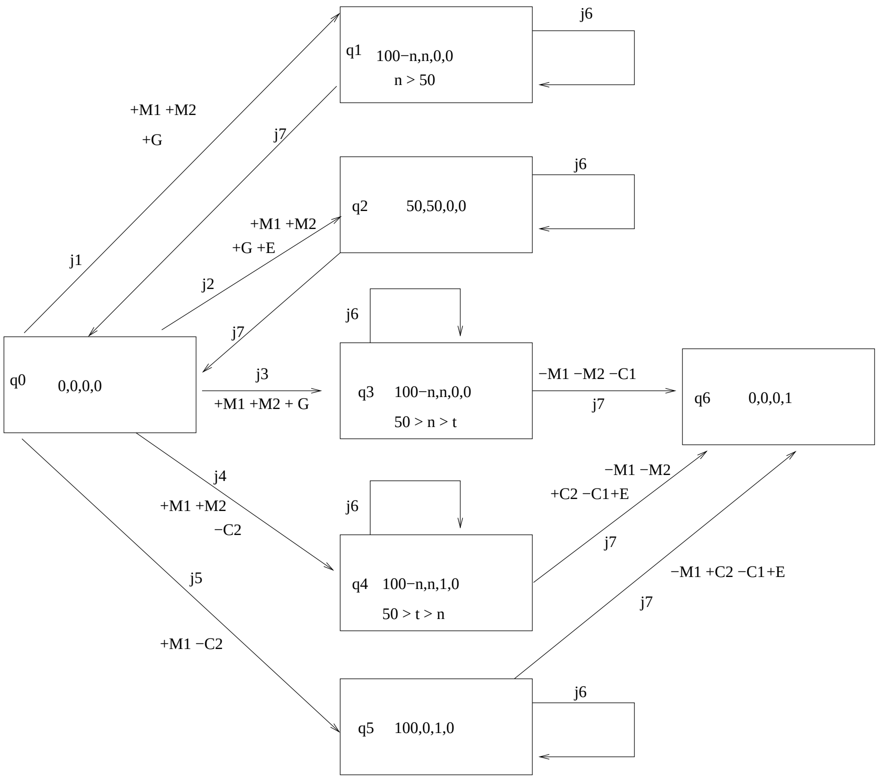

Fig. 6 AATS for Ultimatum Game

literature that some people seem to value fairness, which we shall again 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? As noted previously, 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 M2. Finally either player may be content with the outcome, and we represent this as C1 and C2. Again we will not model degrees of contentment. Labels indicating the promotion and demotion of these values are shown on the AATS in Figure 6.

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. If the second player accepts, the state remains the same, and does not promote any values, and so will be chosen only if the objections to rejecting are preferred to the justification for rejection. The arguments justifying rejection and the objections to them are

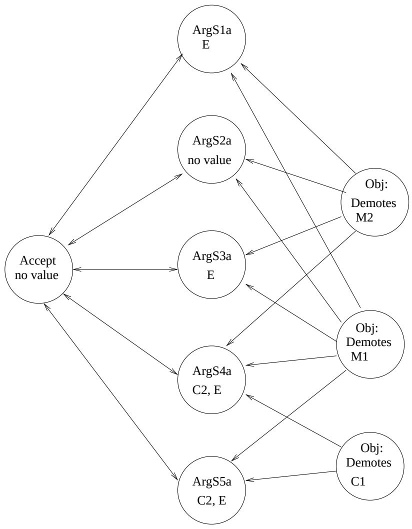

Fig. 7 VAF for acceptance or rejection

shown in Table 4. 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 4 Arguments for the Second Player

ID	Argument	Objections
ArgS1a	Reject in q1 to promote E	Demotes M1, M2
ArgS2a	Reject in q2 (no reason)	Demotes M1, M2
ArgS3a	Reject in q3 to promote E	Demotes M1, M2
ArgS4a	Reject in q4 to promote C2 and E	Demotes M1, M2, C1
ArgS5a	Reject in q5 to promote C2 and E	Demotes M1, C1

This gives rise to the VAF shown in Figure 7.

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 $>$ {M1,M2}. 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 M2 $>$ C2 $>$ {M1,C1} for acceptance, or C2 $>$ M2 $>$ {M1,C1} for rejection. If E is preferred to both M2 and C2 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: {C1, M1} > C2. 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 5. 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.

Table 5 Arguments for First Player

ID	Arguments	Objections
ArgF1a	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
ArgF1b	a1 to promote M2	
ArgF1c	a1 to promote G	
ArgF1d	a1 to avoid demoting C1	
Arg F1e	a1 to avoid demoting C2	
ArgF2a	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
Arg F2b	a2 to promote M2	
ArgF2c	a2 to promote G	
ArgF2d	a2 to avoid demoting C1	
ArgF2e	a2 to avoid demoting C2	
ArgF2f	a2 to promote E	
ArgF3a	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
ArgF3b	a3 to promote M2	
ArgF3c	a3 to promote G	
ArgF3d	a3 to avoid demoting C2	
ArgF5a	a5 to promote M1	a1 demotes C2; a1 may demote C1

If the first agent is highly altruistic, so that M2 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 M2. 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 C2 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 M1 if the second agent prefers M1 to its own contentment, this action will probably be rejected, even by an agent who values only M1, 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 M1 while not demoting C2, and risking the demotion of C1. 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 M1 and C1. If it prizes C1 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 M1 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 C1 or both, then the offer might even rise to 50%, since this will ensure that G is promoted and C1 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 those from the small scale cultures used in [12] can form the basis of an explanation of diverse behaviour (although there are other factors, as evidenced by the variations within cultures).

6.1 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 [12], 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, since both of these are in the South America group and they 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. The authors of [24] were, however, unable to come up with any explanation of this. Four hypotheses were explored based on notions taken from Hofstede [13] 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 [15]. 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 continents rather than across them.

The results from Heinrich *et al* [12], which looked at smaller, more homogeneous, 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 in [12]: 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. Using these dimensions, [12] explains how 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. In [12] 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 by the Lamelara themselves). 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 [12]: 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. We have not included this value in our discussion above, but it would be straightforward to include this, or any other value the effect of which was seen as worth exploring. 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 perceived 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 examples 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 [12] 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.

7 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. Moreover, because this preference order can be reasoned about, the agent can readily adapt its behaviour according to the situation in which it finds itself.

The role of the system's 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 infrequently. 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 *self-interested 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 and the absence of transaction costs, are more easily satisfied. Closed, single-company, multi-agent systems, such as the telecommunications services maintenance system described in [20], 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; in such an extension, agents would often have to interact with other agents with whom they had no prior relationships.

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 [16].

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 [28]. 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 and/or the evolution of sub-populations of agents sharing similar value orderings. A similar idea, although essentially based on an underlying utility function, can be found in [18]. We believe that our approach will be able to achieve similar benefits, but in a more straightforward and transparent manner.

8 Relation to Game Theory

Experiments such at the Dictator and Ultimatum Games, are, of course, the kind of situation which has been thoroughly treated by Game Theory, and a considerable body of knowledge has been developed with respect to these situations. For example, consider the well known Prisoner's Dilemma. This has the payoff matrix shown in Table 6, where T stands for Temptation to defect, R for Reward for mutual cooperation, P for Punishment for mutual defection and S for Sucker's payoff. To be defined as prisoner's dilemma, the following inequalities must hold: $T > R > P > S$

Table 6 Attack relation applicable in $\{0,0\}$

	Cooperate	Defect
Cooperate	R,R	S,T
Defect	T,S	P,P

It has been shown that in theory the optimal strategy in a single shot game is to defect (the Dictator Game and Ultimatum Game are both single shot games), since this will maximise a player's payoff whatever the choice of the other player. But suppose we were describing not a game, but a genuine economic choice in one of the small scale cultures of the sort investigated in [12]. We could imagine that people would have a choice between cooperation and defection, with similar payoffs, but now there would be some additional context. In particular the actual amounts of the payoffs would become important. Suppose first that P was not sufficient to sustain life, just as the whalers of Lamelara would be unable to eke out an existence by individual hunting or farming. Now, it would not be rational to defect since if everyone defected the community would starve. In contrast, if P were sufficient to maintain an acceptable living, but S was not, defection would be rational since to try to gain by cooperation would not be a sensible risk. This would be reinforced further if R was only slightly greater than P. Finally if even S were enough for a living, then cooperation could be risked since being taken advantage of would not be disastrous. Again this would be reinforced if the difference between R and P were substantial, or the difference between P and S small. The relative differences between the amounts would also be likely to influence attitudes to whether defection should be punished. If T was only slightly higher than R, cooperation was necessary, and S were only slightly lower than P, the temptation to defect would be small and any defections might well be overlooked for the sake of securing ongoing cooperation. If, however, T was substantially more than R, the society might well feel the need to impose punishments to enforce cooperation, especially when S was substantially less than P. In cases where defection was the cultural norm, questions of punishment would rarely arise, since no one would be putting themselves in a position to receive S, and the defector could not really be said to have done anything wrong. Thus while in an abstract setting the absolute and relative amounts associated with the payoffs is not relevant beyond establishing the characteristic inequalities, when transferred to a concrete setting, both the absolute amounts relative to what is needed, and the relative amounts should have important effects on both the choice to cooperate or not, and the choice to punish or not. Moreover the cultural norms established in this way might be expected to carry over into the behaviour in experimental situations.

9 Concluding Remarks

In this paper we have shown how our model of an agent choosing an action on the basis of justifications in terms of arguments and a preference ranking on personal values can account for the behaviours of subjects in games from experimental economics, and which, in particular, emphasises the importance of how the scenario is presented to the subjects. It could, however, be argued that the machinery is unnecessary and the behaviour could equally be accounted for in terms of subjects maximising their expected utilities. Such a utility function would, of course, relate not only to the subject's money, but to other factors corresponding to our values, and individual preferences would be modelled by attaching different weights to these factors.

While this may be true, the following points should be noted. First, such a function cannot be simply be applied to states, unlike classical decision theory in which expected utility is calculated on the basis of the likelihood of reaching various states. If so, in the Dictator Game, the same state would be chosen regardless of how the problem were framed. Thus the utility function would need to be applied to the transitions, recognising that actions can have intrinsic utility. Second, any such function would need to be complicated to distinguish between maximisers and satisficers, whereas our account handles this distinction rather elegantly by using different critical questions and so identifying different attacking arguments. Third, our framework provides an explanatory account of the reasoning process of the participants in terms of arguments, which we regard as more instructive than reference to a formula and expected utility calculations. As we have said above, this explanatory capability is also a key requirement for any conceptual framework intended to support the creation of software entities empowered to make decisions, as distinct from a framework intended merely to model human decision-making. Fourth, in extreme situations we cannot trade off one benefit for another; some risks are simply too great to run for a small benefit, however, unlikely they may be. The value-based account naturally represents this absence of trade-off. Fifth, to be usefully deployed, we need a way of identifying the utility function in advance: this is psychologically implausible. In many situations we are really rather bad at attaching numerical weights to our various interests and aspirations, making a qualitative account more plausible. Finally, another feature of practical reasoning identified in [31] is that our preferences are typically determined as a product of practical reasoning rather than given as an input to it. On our account this process is seen when the agent needs to choose actions on the basis of value preferences: we believe that considering the issue that "if you prefer V1 then do A with these consequences but if you prefer V2 do B with these other consequences" gives a more plausible basis for arriving at these preferences than being asked to assign relative weights to factors V1 and V2 so as to determine a calculation of utility. For all these reasons, we believe that the approach described here is worthy of consideration, at least in some situations, as an alternative to utility theory.

To summarise, we have given an account of reasoning in simple scenarios from experimental economics using preferences over values which provides transparent justification of actions, does not require estimation of utilities and weights for different factors, can allow preferences to emerge from reasoning, and which can be used to explain observed behaviours, in particular the framing effect present in these experiments.

As for future work, we should experiment to determine whether transactions are more successful in homogeneous rather than heterogeneous societies. It would also be worthwhile to apply this approach in the context of the various types of application discussed in Section 6, to verify that there were indeed efficiency gains. Another avenue would be to investigate the factors influencing the co-evolution over time of shared value orderings among subpopulations in a population of agents.

Acknowledgments

We are grateful to Ricardo Parreira Duarte for his contributions to the implementation discussed in this paper. We also thank the anonymous reviewers for their helpful comments.

References

1. Aart, C., Marcke, K., Pels, R., Smulders, J.: International insurance traffic with software agents. In: F. Harmelen (ed.) Proceedings of the Fifteenth European Conference on Artificial Intelligence (ECAI 2002), pp. 623–627. Lyon, France (2002)
2. Anand, P.: Foundations of Rational Choice under Risk. Clarendon Press, Oxford, UK (1993)
3. Atkinson, K., Bench-Capon, T.J.M.: Practical reasoning as presumptive argumentation using action based alternating transition systems. Artificial Intelligence **171**(10–15), 855–874 (2007)
4. Atkinson, K., Bench-Capon, T.J.M., McBurney, P.: Computational representation of practical argument. Synthese **152**(2), 157–206 (2006)
5. Bardsley, N.: Dictator game giving: Altruism or artefact. Experimental Economics **11**(2), 122–133 (2008)
6. Bench-Capon, T.J.M.: Persuasion in practical argument using value based argumentation frameworks. Journal of Logic and Computation **13**(3), 429–48 (2003)
7. Bolton, G.E., Katok, E., Zwick, R.: Dictator game giving: Rules of fairness versus acts of kindness. International Journal of Game Theory **27**, 269–299 (1998)
8. Dung, P.M.: 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. Forsythe, R., Horowitz, J.L., Savin, N.E., Sefton, M.: Fairness in simple bargaining experiments. Games and Economic Behavior **6**(3), 347–369 (1994)

10. Fox, J., Parsons, S.: Arguing about beliefs and actions. In: A. Hunter, S. Parsons (eds.) Applications of Uncertainty Formalisms, *Lecture Notes in Artificial Intelligence*, vol. 1455, pp. 266–302. Springer, Berlin, Germany (1998)
11. Friedman, M.: The methodology of positive economics. In: M. Friedman (ed.) Essays in Positive Economics, pp. 3–43. University of Chicago Press, Chicago, IL, USA (1953)
12. Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., McElreath, R.: In search of homo economicus. Behavioral experiments in fifteen small-scale societies. American Economic Review **91**(2), 73–78 (2001)
13. Hofstede, G.: Cultures and Organizations: Software of the Mind. McGraw-Hill, New York, USA (1991)
14. Hogg, L.M.J., Jennings, N.R.: Socially intelligent reasoning for autonomous agents. IEEE Transactions on Systems, Man and Cybernetics – Part A **31**(5), 381–399 (2001)
15. Inglehart, R.: Culture and democracy. In: L.E. Harrison, S.P. Huntington (eds.) Culture Matters: How Values Shape Human Progress. Basic Books, New York, USA (2000)
16. Jonker, G.M., Dignum, F.P.M., Meyer, J.J.: 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), pp. 1258–1260. ACM Press (2007)
17. Kahneman, D., Slovic, P., Tversky, A. (eds.): Judgment under Uncertainty: Heuristics and biases. Cambridge University Press, Cambridge, UK (1982)
18. Karunatillake, N.C., Jennings, N.R., Rahwan, I., Ramchurn, S.D.: Managing social influences through argumentation-based negotiation. In: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2006), pp. 426–428 (2006)
19. Lindley, D.V.: Making Decisions, second edn. John Wiley and Sons, London, UK (1985)
20. Mihailescu, P., Shepherdson, J., Marrow, P., Lee, L., Lee, H.: 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), vol. 2, pp. 1889–1894. IEEE, The Hague, The Netherlands (2004)
21. Mill, J.S.: 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)
22. Modgil, S.: Reasoning about preferences in argumentation frameworks. Artificial Intelligence **173**(9–10), 901–934 (2009)
23. Nowak, M.A., Page, K.M., Sigmund, K.: Fairness versus reason in the Ultimatum Game. Science **289**, 1773–1775 (2000)
24. Oosterbeek, H., Sloof, R., van de Kuilen, G.: Differences in Ultimatum Game experiments: Evidence from a meta-analysis. Experimental Economics **7**(2), 171–188 (2004)
25. Parreira Duarte, R.L.: Modelling arguments in the Dictator Game. Tech. Rep. ULCS-10-006, Department of Computer Science, University of Liverpool (2009)
26. Parsons, S.: Qualitative Methods for Reasoning Under Uncertainty. MIT Press, Cambridge, MA, USA (2001)
27. Parsons, S., Gmytrasiewicz, P., Wooldridge, M. (eds.): Game theory and Decision Theory in Agent-based Systems. Kluwer Academic, Boston MA USA (2002)
28. Ramchurn, S.D., Sierra, C., Godo, L., Jennings, N.R.: Negotiating using rewards. Artificial Intelligence **171**(10-15), 805–837 (2007)
29. Roberts, J.H., Lattin, J.M.: Development and testing of a model of consideration set composition. Journal of Marketing Research **28**, 429–440 (1991)
30. Savage, L.J.: The Foundations of Statistics. Wiley, New York, NY, USA (1950)
31. Searle, J.R.: Rationality in Action. MIT Press, Cambridge, MA, USA (2001)
32. Shackle, G.L.S.: Decision, order and time in human affairs. Cambridge University Press, Cambridge, UK (1961). 2e 1969
33. Shafer, G.: Savage revisited. Statistical Science **1**(4), 463–485 (1986)
34. Shugan, S.M.: The cost of thinking. Journal of Consumer Research **7**, 99–111 (1980)
35. Simon, H.A.: Rational choice and the structure of the environment. In: Models of Bounded Rationality, Volume 2, pp. 259–268. MIT Press, Cambridge, MA, USA (1957)
36. Walton, D.N.: Argumentation Schemes for Presumptive Reasoning. Lawrence Erlbaum Associates, Mahwah, NJ, USA (1996)
37. Wooldridge, M., van der Hoek, W.: On obligations and normative ability: Towards a logical analysis of the social contract. Journal of Applied Logic **3**, 396–420 (2005)