Emotions in Rational Decision Making

Fahd Saud Nawwab, Trevor Bench-Capon, and Paul E. Dunne

The University of Liverpool

Abstract. Emotions play an important role in human decision making. This often has a beneficial effect, and so it is desirable to explore whether emotions can also play a role in agent decision making. This paper builds on several pieces of earlier work which resulted in a formalization of a model of emotions. We show how, in particular, the emotions of gratitude and displeasure and related emotions can be beneficially integrated into a model of decision making. The model of decision making used is based on a notion of choosing between competing justifications according to the agent's preferences over the social values promoted by the available actions. Emotions of gratitude and displeasure are generated according to whether other agents relevant to the situation support or frustrate the agent's goals and values, and the emotional attitude towards the other agents then influences the ranking of these values, and so affects future choices. The paper summarizes the previous work on which we build, describes the decision making model we use, and explains how emotions are generated and used within that model.

1 Introduction

That emotions do in practice have an influence on human decision making cannot be denied. We should not, however, consider this to be something to be deprecated: if there were no beneficial effects of emotions, it seems unlikely that evolution would have left them with such a widespread role. Damasio [5] writes:

> "Emotions and feelings can cause havoc in the process of reasoning under certain circumstances. Traditional wisdom has told us that they can, and recent investigation of the normal reasoning process also reveal the potentially harmful influence of emotional biases. It is thus even more surprising and novel that the absence of emotion and feeling is no less damaging, no less capable of compromising the rationality that makes us distinctively human and allows us to decide in consonance with a sense of personal future, social convention, and moral principle."

So, what is the beneficial role of emotions, and how might we gain these beneficial effects in agent reasoning? One aspect identified by Damasio and developed in [13] is that emotions can act as heuristics and prevent excessive deliberation. Thus in [13] the authors show how the emotions of hope and fear with respect to plans can be used to determine whether replanning is necessary, or whether the current plan can still be followed. There are, however, other sorts of emotion and

other roles. In this paper, we will consider the emotions of gratitude, displeasure, like, dislike, joy and distress to show how these emotions can impact on practical reasoning. We argue that these emotions play an important role in social interaction: humans adopt emotional attitudes towards one another, and this seems to play an essential role in developing and maintaining cooperation, and behaving towards one another in a consistent and appropriate way. These emotions also seem to act as tie-breakers to enable a reasoned choice between two alternatives which are equally acceptable on purely rational grounds. We will show how these effects can be realised in a computational model of practical reasoning.

There are several approaches to modeling emotions in artificial agents. One example is the EBDI model built by Jiang et al [7] based on the famous BDI (Belief - Desire - Intention) model of Bratman [4] which incorporates an emotional function into that architecture. Another is that developed by Reilly [12] which aims to express emotions rather than influence decisions. Padgham and Taylor [11] introduced a system that is designed to treat emotions with goal-oriented behaviors trying to model the personality aspect of agents. The work of Steunebrink et al [13], which takes the Ortony, Clore and Collins (OCC) model [10] as its basis, presented a formalization of a number of emotions, and as mentioned above, focused on the use of hope and fear in particular to guide replanning. We will take [13] as our starting point, and show how some of the other emotions impact on the practical reasoning mechanism described in [9].

In Section 2, we will briefly describe the work in emotional agents on which we base our own, and identify the specific emotions on which we will focus. Section 3 will describe the model of decision making we use, and Section 4 will show how emotions can be generated in this model and the ways in which they can influence decision making, using a simple two agent example. Section 5 walks through a more extended example and Section 6 offers some concluding remarks.

2 Emotional Agents

The original psychological basis for the work on which we build is provided by the model of (OCC) [10]. This model, which identifies twenty-two emotions organised into a concise hierarchy, is particularly attractive to those wishing to build computational models as it provides a specification of the conditions which give rise to each distinct emotion in terms of the computationally familiar notions of objects, events and actions. The OCC model includes notions of intensity of emotions, and identifies a number of variables which influence the intensity of the emotions experienced. The OCC model formed the basis of the implementation of a model of emotional agents developed by Reilly [12]. Reilly, whose aim was to construct believable emotional agents, gives a system capable of generating emotions, storing them, and expressing them through behavior. Reilly gives a simplified method of determining intensity, using only a subset of the variables from the OCC model, most importantly the *importance* and *unex-*

pectedness of the triggering event. Under storage, Reilly includes discussions of, and options for, the mechanisms for combining emotions of a similar type and enabling them to decay over time. Reilly recognises that different combinations and decay mechanisms may be appropriate for different emotions. As far as expression goes, Reilly's focus is on believability and he does not make any strong claims for how it interacts with rationality in decision making. We adopt Reilly's approach to the quantitative aspects of emotions.

The immediate precursor of our approach is the work done by the agents' group at Utrecht e.g. [8, 13]. They have provided a formalisation of the emotions of the OCC model, and have shown how hope and fear in particular can play a role in decision making. In [13], the twenty-two emotional fluents of the OCC model are defined. We largely follow their definitions, but subscript states of affairs with the name of an agent where appropriate. We also change the names of some of the emotions: we regard love, hate and anger as suggesting more intensity than is necessary, and so will use *like*, *dislike* and *displeasure* instead.

Definition 1 (Emotional Fluents); The set Emotions is the set of emotional fluents, which is defined as follows:
Emotions =

$joy_i(k_i)$,	$distress_i(k_i)$,
$hope_i(\pi, k_i)$,	$fear_i(\pi, \neg k_i)$,
$satisfaction_i(\pi, k_i)$,	$disappointment_i(\pi, k_i)$,
$relief_i(\pi, \neg k_i)$,	$fears-confirmed_i(\pi, \neg k_i)$,
$happy-for_i(j, k)$,	$resentment_i(j, k)$,
$gloating_i(j, k)$,	$pity_i(j, k)$,
$pride_i(\alpha_i)$,	$shame_i(\alpha_i)$,
$admiration_i(j, \alpha_j)$,	$reproach_i(j, \alpha_j)$,
$like_i(j)$,	$dislike_i(j)$,
$gratification_i(\alpha_i, k_i)$,	$remorse_i(\alpha_i, k_i)$,
$gratitude_i(j, \alpha_j, k_i)$,	$displeasurei(j, \alpha_j, k_i)$,

Where i and j are distinct agents, α_i and α_j are actions available to these agents, π a plan of agent i expressed as a sequence of actions, and k is a partial state of affairs, so that k_i is a goal of agent i.

From this definition, we can see that the twenty-two emotions can be divided into eleven pairs. For example, distress is the negative correlate of joy, and so on down the list. We can also note that emotions form a number of distinct types, with respect to their arguments.

- Those relating to a plan: joy, distress, hope, fear, satisfaction, dissatisfaction, relief and fears-confirmed. We will not consider these further as we do not consider plans in this paper.
- Those directed towards another agent. Gratitude and displeasure are directed towards another agent in respect of that other agent's action bringing about or frustrating a state of affairs desired by the emotional agent. Admiration and reproach are also directed towards another agent in respect of that other agent's action, but here there is no connection with the emotional

agent's goals. Like and dislike are directed towards another agent, but in respect of no particular event. Gloating and pity, happy-for and resentment are directed towards another agent in respect of a state of affairs which need not be a goal of the emotional agent
- Gratification and remorse correspond to gratitude and displeasure, but this time are directed by the agent towards itself in respect of its own action and its effect on the agent's goals.
- Joy and distress, which are experienced when goals are achieved or fail to be achieved.

In this paper our focus will be on three pairs in particular: gratitude and displeasure; like and dislike; and joy and distress. In general, terms we are considering an agent which selects an action hoping to realise some goal. Whether it is successful or not, depends in part on the actions of other agents. Where other agents cooperate in realising the goal; the emotional agent will feel gratitude towards them, and conversely, if they frustrate the goal the emotional agent will feel displeasure. Over a series of actions, gratitude will accumulate into liking and displeasure into disliking. Similarly a series of successful actions will lead to joy and unsuccessful actions to distress. Of course, actions giving rise to admiration and respect will also influence like and dislike, and gratification and remorse will influence joy and distress. Here, however, we shall consider only gratitude and displeasure and the emotions influenced by them. Widening the scope into other related emotions will be left for future work. In the next section, we will present our model of decision making before going on to explain how the emotions of gratitude and displeasure are generated in this context.

3 A Model of Agent Decision Making

The model of agent decision making is based on the approach of [1], which has been further developed in [9]. The approach relies on argumentation techniques: candidate actions are identified by providing *prima facie* justifications for them by instantiating a particular argumentation scheme. This justification can then be critiqued by a set of characteristic counterarguments. The decision is then made by choosing a defensible set of action justifications. In [9] five steps were identified:

1. *Formulating the Problem*: A representation of the problem scenario in terms of an Action-Based Alternating Transition System (AATS) [14].
2. *Determining the Arguments*: An instantiating of an argument scheme justifying an action in terms of the AATS, and identifying counter arguments based on critical questions appropriate to that scheme. A formal description of the argument scheme and critical questions can be found in [1].
3. *Building the Argumentation Framework*: The arguments and counter arguments identified in the previous set are formed into a Value Based Argumentation Framework (VAF) [3].

4. *Evaluating the Framework*: The particular value ordering of the agent concerned is now used to identify the preferred extension of the framework for that agent. This identifies the arguments that withstand the critique from the perspective of the decision making agent.
5. *Sequencing the Actions*: The set of actions justified in the previous step are now sequenced for execution.

Here we are especially interested in the fourth step: our contention is that the emotional state of the agent will affect its ordering of values, and so lead to a preferred extension different from that of an agent unaffected by emotions. We will, now, give some necessary detail of the argument scheme, the AATS, and VAFs.

The particular argument scheme and its associated critiques used in Step 2 was introduced in [2]. The argument scheme to justify an action is as follows:

AS1: In the current circumstances R, action A is justified, since it will bring about new circumstances S, which realise goal G, which promotes value V.

What is distinctive about this scheme is that it distinguishes three aspects of the results of an action: the state of affairs which follows from the action; desirable elements of that state of affairs, the goal, and the reasons why those elements are desirable, the value. Values, which are social values, such as liberty and equality, are central here. Suppose we wish to justify an increase in taxes, then our goal may be redistribution of income, justified by our desire to promote equality, or our goal may be to have more money to spend on national defense, promoting the value of national security. Different justifications will appeal to different audiences, depending on their values and the order in which they prioritities them. A fuller discussion of values and their role can be found in [2]. What is important here is that using this scheme associates arguments with values.

In order to generate the arguments, [1] uses a formulation of the problem as an Action-Based Alternating Transition System (AATS). AATS was introduced in [14] in order to represent the effect of a group of agents acting independently in a situation, so that the result of the action of a given agent is dependent on the actions of the other agents concerned. The states of the AATS represent possible states of affairs, and the transitions between them, that is actions composed from the individual actions of the relevant agents. Thus, if two agents can each independently choose one of three actions, there will be nine possible joint actions. Additionally, each of the transitions can be labeled with the values promoted by moving from the source state to the target state. In [1], the argument scheme for justifying actions was formally described by reference to an AATS, as were the possible counterarguments to such a justification. The formal definition of an AATS is given in Definition 2, and the argument scheme AS1 is defined in terms of an AATS in Definition 3. Each transition from the current state which promotes some value can be used to instantiate AS1 and so justifies the agent in performing its component of the corresponding joint action. This justification is, however, subject to counterarguments, and so the result is a set of conflicting arguments, which will justify a range of actions, some of which may be incompatible. These arguments can be ascribed a status by forming them into

a Value-Based Argumentation Framework (VAF) [3], an extension of standard Argumentation Frameworks (AF) [6], which are presented in Definitions 4 and 5 below. Whereas, in an AF, attacks always succeed, in a VAF an argument is defeated by its attacker, if that attacker is associated with a value of equal or greater preference. Note that AS1 associates arguments which instantiate it with a value. The agent may now, by applying its preference ordering over values, identify a set of acceptable arguments in the VAF. The actions justified by these arguments will be those that the agent will choose, in the light of its preferences.

Definition 2: AATS. A joint action j_{Ag} for a set of k agents Ag, is a tuple $\langle \alpha_1, .., \alpha_k \rangle$, where for each α_i $(j \leq k)$ there is some $i \in Ag$ such that $\alpha_j \in Ac_i$, the set of available actions for agent i. Sets of available actions are *disjoint*, i.e. $i \neq j \Rightarrow Ac_i \cap Ac_j = \emptyset$. The set of all joint actions for a set of agents Ag is denoted by J_{Ag}, so $J_{Ag} = \Pi_{i \in Ag} Ac_i$. Given an element j_{Ag} of J_{Ag} and an agent $i \in Ag$, i's action in j_{Ag} is denoted by j_i

An AATS is a $(2n{+}8)$ element tuple S=$\langle Q, q_0, Ag, Ac_1, Ac_n, Av_1, Av_n, \rho, \tau, \Phi, \pi, \delta \rangle$

Q is a finite, non-empty set of states

$q_0 = q_x \in Q$ is the initial state

$Ag = \{1, ..., n\}$ is a finite, non-empty set of agents

Ac_i is a finite, non-empty set of actions, for each $i \in Ag$ (where we recall that $Ac_i \bigcap Ac_j = \emptyset$ whenever$i \neq j$)

Av_i is a finite, non-empty set of values $Av_i \subseteq V$, for each $i \in Ag$

$\rho : Ac_{Ag} \rightarrow 2^Q$ is an action precondition function, which for each (joint) action $\alpha \in Ac_{Ag}$ defines the set of states $\rho(\alpha)$ from which α may be executed

$\tau : Q \times J_{Ag} \rightarrow Q$ is a partial system transition function, which defines the state $\tau_(q, j)$ that would result by performing j from state q - note that, as this function is partial, not all joint actions are possible in all states (cf. the precondition function above)

Φ is a finite, non-empty set of atomic propositions

$\pi : Q \rightarrow 2^\Phi$ is an interpretation function, which gives the set of propositions satisfied in each state: if $p \in \pi(q)$, then this means that the propositional variable p is satisfied (equivalently, true) in state q

$\delta : Q \times Q \times Av_{Ag} \rightarrow \{+, -, =\}$ is a valuation function which defines the status (promoted (+), demoted (−) or neutral (=)) of a value $v_u \in Av_{Ag}$ ascribed by the agent to the transition between two states: $\delta(q_x, q_y, v_u)$ labels the transition between q_x and q_y with one of $\{+, -, =\}$ with respect to the value $v_u \in Av_{Ag}$.

Definition 3 : Argumentation Scheme AS1 expressed in terms of AATS. The initial state q_0 is $q_x \in Q$; agent $i \in Ag$ should perform α_i, consistent with the joint action $j_{Ag} \in J_{Ag}$ where $j_{Ag_i} = \alpha_i$, so that $\tau(q_x, j_{Ag}) = q_y, p_a \in \pi(q_y) \backslash \pi(q_x)$ and for some $v_u \in Av_i$, $\delta(q_x, q_y, v_u) = +$.

Definition 4: Value-Based Argumentation Framework (VAF). VAF is a triple $\langle H(X, A), \nu, \eta \rangle$, where X is a set of arguments, A is a binary relation on

arguments, called the attack relation (i.e. $H(X, A)$ is a standard argumentation framework in the sense of [6]), $\nu = v_1, v_2, ..., v_k$ 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 5: **Admissibility in VAFs**. Let $\langle H(X, A), V, \eta \rangle$ be a VAF and α an audience. 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 $\eta(y) \succ_\alpha \eta(x)$. 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. 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 $\eta(y) \succ_\alpha \eta(x)$. 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. A subset R is a preferred extension for the audience α if it is a maximal admissible set with respect to α.

3.1 Example with Two Agents

In section 5 we give an extended example, but we here use a simpler example. Suppose two colleagues, Teddy and Karen, need to travel from Liverpool to Paris for a conference. They can choose to travel by plane or train. The plane is faster, but the train journey is much more comfortable. Their choice, therefore, comes down to whether they prefer the value of comfort or of speed. There is, however, a third consideration: it is more pleasant to travel in company than alone. The initial state has Karen and Teddy in Liverpool and there are four joint actions, depending on their independent choice of train or plane. These joint actions are:
J1: Teddy plane(Tp) Karen plane(Kp) **J2:** Teddy plane(Tp) Karen train(Kt)
J3: Teddy train(Tt) Karen train(Kt) **J4:** Teddy train(Tt) Karen plane(Kp)

The transitions will promote the values of Teddy's Comfort (Ct), Teddy's Speed (St), Karen's Comfort (Ck) and Karen's Speed (Sk), and a Pleasant Journey (P). Note that comfort and speed are promoted in respect of particular agents, whereas the pleasant journey is a common good. On the basis of the AATS shown in Figure 1 we can get a number of instantiations of AS1:
A1: Karen travel by train (Kt) to reach q1/q3 to promote her comfort (Ck).
A2: Karen travel by plane (Kp) to reach q2/q4 to promote her Speed (Sk).
A3: Teddy travel by train (Tt) to reach q1/q2 to promote his comfort (Ct).
A4: Teddy travel by plane (Tp) to reach q3/a4 promote his speed (St).
A5: Both travel by train (Kt, Tt) to reach q1 for a pleasant journey(P).
A6: Both travel by plane (Kp, Tp) ro reach q4 for a pleasant journey(P).

These arguments form a VAF as shown in Figure 2, the attack relations representing that a person cannot travel both by train and by plane. First, let us suppose that the agents consider only their own values. Suppose that Karen orders her values $Ck > P > Sk$. For her, A1 will resist the attacks of

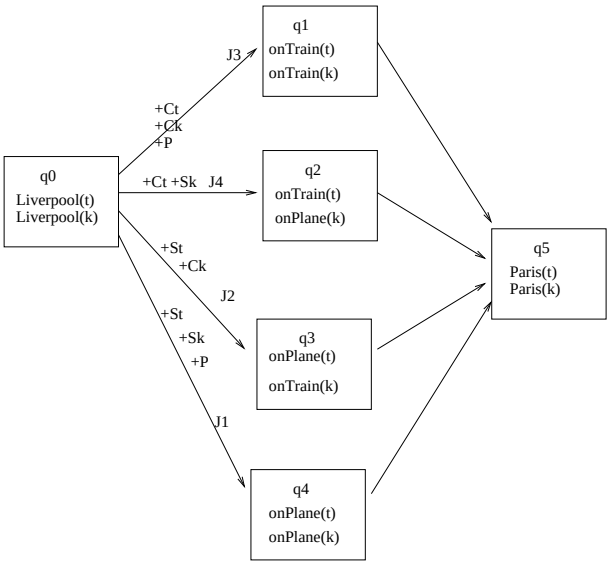

Fig. 1. The AATS of the example

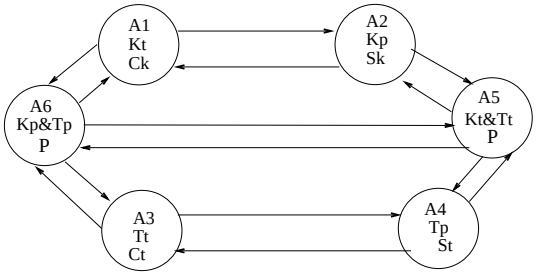

Fig. 2. VAF graph of the example

A2 and A6, and in turn defeat them, so that A5 is also justified for Karen, that is her preferred extension is (A1, A5). Thus, Karen will travel by train. But suppose that Teddy orders his values $P > St > Ct$. This preference gives rise to two preferred extensions: (A2, A4, A6) and (A1, A3, A5). Teddy can thus justify traveling by either method. Whereas Karen could ensure that her preferred extension was realised, since it does not matter what Teddy chooses, Teddy needs to predict Karen's actions in order to realise his goals. If he believed that Karen was equally likely to choose the plane and the train, he would choose the plane since $St > Ct$. But if he knows Karen's value order then he should choose the train since he should expect her to choose the train, so that J1 will not occur and choosing Tp will therefore lead to J2 rather than J1, reaching q3 rather than q4 and so failing to promote P.

Note, however, that since A1 subsumes two arguments, Karen's goal can be either q1 or q3. Since she also values P, this makes q1 preferred to q3, so her goal in A1 will be q1. Teddy on the other hand, because his choice is based on A5 or A6, needs to realise q1 or q4 respectively since Karen's choice leads to q2 or q3 P will not be promoted. In the next section, we will consider how the choices made by Teddy and Karen give rise to emotions, and how these emotions can affect subsequent choices.

4 Gratitude, Displeasure, Like, Dislike, Joy and Distress

4.1 Generation

Recall that gratitude and displeasure are emotions directed by one agent towards another in respect of an action which brings about a state of affairs. This maps very well to the AATS structure in which an agent selects an action to reach a particular state and promote particular values, but is not able to ensure success since the actual state reached, and values promoted, depend on the actions chosen by the other agents involved, the joint action. In so far as the other agent does not allow the desired joint action, the agent will feel displeasure and in so far as it does assist in the joint action being performed, the agent will feel gratitude.

Recall from section 2, that the intensity of emotions depend on how important it is to reach the goal and how unexpected the action is. In our example, suppose first that Karen chooses to travel by train with q1 as her desired state. It is not of vital importance to her whether q1 or q3, the result of other possible joint action when she chooses the train, is reached: both promote her important value of C. For Teddy, on the other hand, it is important that the intended joint action is performed, since the alternative will not promote his important value of P. So if Teddy chooses the plane and Karen the train, both will be displeased with one another, since they had intended to reach a different state, and the value of a pleasant journey is not promoted by the actual joint action. Teddy will feel more displeasure than Karen since he set more importance on promoting P. Suppose, however, that their preferences were known to one another in advance, and that Teddy was well aware of Karen's extreme dislike of flying. Then, Karen's choice should have been expected, and so his displeasure should be reduced to that extent. Teddy's choice is less predictable, since he has two preferred extensions: thus if he chooses the plane, Karen may well feel more intense displeasure with respect to this aspect.

In practice, however, if Teddy knew of Karen's preference, so that he could recognise the train as her only rational choice, Teddy should have performed Tt hoping to reach q1. In this case, both want to reach q1, and this is the state they will reach. They may both feel gratitude towards the other. However, because Kt was entirely expected, and because Ct was less important than the now unrealised St, Teddy's gratitude is likely to be minimal. On the other hand, Teddy's choice was less predictable, and enabled Karen's desired state to be reached, and so she should feel rather more intense gratitude towards him.

We need to make some attempt to quantify importance and unexpectedness. For our current purposes precision is less important than getting the qualitative aspects right: provided importance and intensity move in the right direction, we can explore their effects. We therefore, use a very simple method. More sophisticated methods are also possible, and can be explored in future work, but the following has the correct characteristics.

Let $Audience_i$ be the order given to a set V of n values agent $i = \langle v_n > v_{n-1} > ... > v_1 \rangle$. Let q be a state such that moving from q_0 to promote the set of values $V_v \subseteq V$. Now the *importance* of q to agent i, $importance_i(q)$ will be $\sum_{i=1}^{i=n} v_i$ such that $v_i \in V_v$. Let VAF_0 be the VAF formed by agent i in q_0 and m the number of preferred extensions of VAF_0 with respect to $Audience_i$. Now the unexpectedness of the action of agent i, $unexpectedness_{q_0}(i)$ in q_0 will be $1 - 1 \div m$. Note that when an agent i is considering the expectedness of an action of some other agent j, it will often only be possible to estimate a partial order, increasing the value of m considerably. If there is information available about the probability of an agent choosing an action, that can be used instead of the calculation in terms of preferred extensions.

In our example, the importance of q1 is 5 for Karen and 4 for Teddy, the importance of q2 is 1 for both of them, the importance of q3 is 3 for Karen and 2 for Teddy, and the importance of q4 is 4 for Karen and 5 for Teddy. The unexpectedness of Karen choosing the train in q_0 is 0, of Teddy choosing the train 0.5 and Teddy choosing the plane in 0.5.

Now suppose agent i in q_0 desires to reach q_d and, given the action chosen by agent j, α_j, q_a is the state actually reached. Where $q_d = q_a$, the intensity of $gratitude_i(j, \alpha_j, q_d)$ is $importance_i(q_d) \times unexpectedness_j(q_0)$. If $q_d \neq q_a$, then the intensity of displeasure is $(importance_i(q_d) - importance_i(q_a)) \times unexpectedness_j(q_0)$.

So in the example, the intensities of the emotions experienced are as follows. When both choose the train: Teddy feels neither gratitude nor displeasure, since Karen's action was entirely expected, while Karen feels gratitude to Teddy with intensity 2.5. Where Karen chooses the train and Teddy the plane, Teddy feels no displeasure because Karen's choice was entirely predictable, whereas Karen feels displeasure with intensity 0.5. If both choose the plane, Teddy feels considerable gratitude (5) to Karen, since her action was entirely unexpected. whereas Karen feels only mild gratitude to Teddy (2), since P was of less importance to her. Finally, if Teddy chooses the train and Karen the plane, Teddy feels displeasure towards Karen of 4 and Karen to Teddy of 1.5.

Definition 6: Gratitude and Displeasure. Formally, suppose agent i seeks a state q from the current state q_k to promote value v,ie $\delta(q_k, q, v) = +$. Let $J = \langle a_0, a_1, \ldots, a_n \rangle$ be an *intended* joint action whose effect achieves these, and $J' = \langle a_0, b_1, \ldots, b_n \rangle$ be the *actual* joint action performed. Agent j is *cooperative* w.r.t the joint action J if $b_j = a_j$; otherwise agent j is said to *frustrate* joint action J.

We can now model gratitude and displeasure felt by agent i towards other agents in terms of the outcome of joint action J' relative to the value promoted. We have the following possibilities where performing J' results in the state q'

1. If $\delta\{q_k, q', v) = +$ then $gratitude(i, b_j, g)$ for all cooperative agents.
2. If $\delta\{q_k, q', v\}\neg = -$ then $displeasure(i, b_j, g)$ for all frustrating agents.

Gratitude and displeasure relate to a single action. In practice, agents will, over time, participate together in many joint actions, each with the potential to evoke gratitude and displeasure. These experiences of gratitude and displeasure will combine, taking into account intensity and decay over time, to form an overall attitude towards the other agent, giving rise to the emotions of like and dislike. Again for illustrative purposes we suggest a very simple mechanism; more sophisticated work can be found in [12].

Let $L_i(j,t)$ be the intensity of like, where $L_i(j,t)$ positive, or dislike, where $L_i(j,t)$ negative, for agent i towards agent j at time t. Let $G_i(j,v,t)$ be the intensity of gratitude felt by agent i towards agent j at t and $D_i(j,v,t)$ be the intensity of displeasure felt by agent i towards agent j at t both with respect to v. Now $L_i(j,t+1) = L_i(j,t) + (G_i(j,v,t) - d_g) - (D_i(j,v,t) - d_d)$, where d_g and d_d are suitable decay factors for like and dislike respectively. In our examples, for simplicity, we will not consider decay, so that like and dislike are only altered by specific experiences.

Similarly the achievement, respectively failure, of the goals will contribute to the overall attitude of the agent, and so contribute to the emotion of joy, respectively, distress. We provide a mechanism similar to that for like and dislike.

Suppose there are n other agents in Ag. Let $J_1(v,t)$ be the intensity, with respect to some value v, of joy, where $J_1(v,t)$ positive, or distress, where $J_1(v,t)$ negative, for agent 1 at time t. Let $\sum_{j=2}^{j=n} G_1(j,v,t)$ be the intensity of gratitude felt by agent i towards the other agents at tin respect of v and $\sum_{j=2}^{j=n} D_1(j,v,t)$ be the intensity of displeasure felt by agent 1 towards the other agents at t with respect to v. Now $J_i(v,t+1) = J_i(v,t) + (G_i(j,v,t) - d_g) - (D_i(j,v,t) - d_d)$, where d_g and d_d are suitable decay factors. Again we will, for simplicity, not consider decay further in this paper.

4.2 Influence on Decision Making

In the last section, we indicated how gratitude and displeasure can be evoked in our decision making model, and how over a series of actions these could lead to the emotions of like, dislike, joy and distress. In this section, we will describe how these emotions can impact decision making.

The mechanism by which emotions will influence decisions is by their impact on the value ordering of the emotional agent. Recall, from Section 3 that the crucial element in the decision making model is the preference order which the agent uses to identify the acceptable arguments in the VAF. Had Teddy preferred Speed to a Pleasant journey, he would not have considered using the train, whatever Karen's preferences. Thus, if the emotions are to have an effect, they

must produce some change in the preference ordering. Recall, also, that some values are common goods and do not relate to any particular agent, whereas some values are promoted and demoted in respect of a specific agent.

An agent will have some initial value order. In the example, Teddy prefers pleasure to speed and speed to comfort. Thus far, our agents have considered only their own values and have decided independently of one another, but in practice the comfort and speed of Karen should have some importance for Teddy, and the speed and comfort of Teddy some importance for Karen. Suppose that they must come to a joint decision. Their two value orders can be merged on P to give either $Ck > P > Sk > St > Ct$ or $Ck > P > St > Sk > Ct$: plausibly Karen will favor the first and Teddy the second since that resolves the partial order by preferring their own value. Now, because Ck is the most preferred value, and because P is preferred to St, the preferred extension will be (A1, A3, A5), that is they both travel by train. A more interesting situation arises if we suppose that Karen orders her values $P > Ck > Sk$, and we assume that neither have any knowledge of the others preferences, so that each believe that there is an even chance of the other choosing plane or train. Now there are four possible merged orderings.

1-$P > St > Ck > Ct > Sk$, 2-$P > St > Ck > Sk > Ct$

3-$P > Ck > St > Ct > Sk$, 4-$P > Ck > St > Sk > Ct$

With these preferences, the situation is less clear, since there is no clearly preferred value. Order 1 and 2 give rise to (A2, A4, A6) and orders 3 and 4 to (A1, A3, A5). Suppose after some discussion they agree on the train, that is Teddy defers to Karen and they travel by train. The agents now consider the actual joint action, ranking states according to the method given above, that is using their own values. Teddy will feel some displeasure (0.5) towards Karen and Karen will feel gratitude towards Teddy (2.5). Accordingly, Karen will increase her liking for Teddy and Teddy will feel some dislike for Karen. Now our suggestion is that liking will influence the value order *by increasing the priority given to values promoted in respect of the liked agent* or *by demoting the values in respect of the disliked agent*. That is, as Karen's liking for Teddy increases, she is more likely to accept order 1. Similarly, Teddy is less likely to accept orders 2, 3 and 4. Like and dislike must pass some threshold to reorder the values, since otherwise the value order will be too volatile. Different agents may have different thresholds, and there may be different thresholds for different emotions. Suppose the liking of 2.5 is enough to dispose Karen to accept order 1. When the question arises again, they will travel by plane this time. Now the roles are reversed: Karen's liking for Teddy will decrease to 2, whereas Teddy will now like with intensity 2. The liking - disliking is asymmetrical because reaching the unfavored state still realises some of their values. Suppose, however, that 2.5 had been insufficient to get Karen to change her value order, but that Teddy had a very low threshold for dislike. Now Teddy may elevate St above P, and so choose the plane. This will, of course, displease Karen, and a mutual dislike will arise, so that P ceases to be so important for either of them, and they will habitually travel separately. We can see here how liking and displeasure can provide a mechanism for *reciprocation*: when one agent makes a concession to another, the influence of liking on the

part of agent who receives the concession and dislike on the part of the agent who makes the concession, influences the value order so that subsequently the roles are reversed. Where reciprocation occurs, mutual liking grows, making it more likely that these agents will value the interests of each other in subsequent decisions (a form of *bonding*). Where reciprocation does not occur, the influence of the value order is to diminish the importance of shared values, giving less importance to the resulting non-cooperation.

Thus far we have considered two agents with competing values relating to themselves. These emotions have a role also where several agents are involved, and the decision maker's own interests are not really of concern: this is typically the case, for example, where a manager is making decisions affecting a group of staff. Here we would expect attitudes of like and dislike to influence the way the manager orders values promoted in respect of his staff. For example, suppose that Teddy and Karen's travel arrangements were decided by their manager rather than by their mutual agreement: an indifferent manager might have the order $(St = Sk) > P > (Ct = Ck)$. Suppose that the manager liked Teddy better than Karen: now his ordering would become $St > Sk > P > Ct > Ck$. Normally, we would expect liking and disliking only to influence the order of the values with respect to particular agents *within* the overall value order. If, however, these emotions become very intense, the effect may be on the *order of values*, rather than merely the ordering of individuals within a value. Intense liking for Teddy might lead the decision maker to prefer Teddy's Comfort to the common good of Pleasure, and intense dislike for Karen might lead to preferring Pleasure to her Speed: $St > Ct > P > Sk > Ck$. At this point it may be that we think the decision making is being distorted to the point of bias: like and dislike have tipped over into love and hate (as they were termed in the original OCC model). Certainly such instances of illegitimate bias abound in human decision making, but we may wish to disallow them in an agent system, and to ensure that a value in respect of a particular agent is not promoted beyond a supposedly preferred value in this way.

We can also consider the effect of joy and distress. It is well known that the best time to ask a person for something is when they are in a good mood. A joyful agent is more likely to be amenable to requests than one experiencing distress. Thus, if the manager agent experiences a lot of joy-evoking events, it may revise its value order so that, for example, Pleasure becomes preferred to Speed. This seems natural: if a lot of time has been saved because Speed has been much promoted, Speed becomes less critical, whereas it may rise in importance if many delays have been suffered. Again, it can be seen that this can have some beneficial effects. A rigid and unchanging value order will lead to one-sided decisions, continually promoting one value at the expense of another. Experiencing joy will mean that a consistent run of success has been achieved, and so the more favored values will have been served. A mechanism which allows attention to be turned to promoting some more neglected value (remember that all the values are legitimate values worthy of promoting and that a different agent could equally well have chosen a different ordering to start with) will often be

desirable. Equally, the failure which gave rise to distress might have meant that the values of the decision-making agent were radically out of step with the other agents, and a more cooperative atmosphere may result from some reordering.

In this section, we have described how the emotions of gratitude, displeasure, like, dislike, joy and distress can arise in our decision-making model. We have also discussed how these emotions can influence decisions in ways which correspond to experience of human decision making, and which are likely to have beneficial effects on the decision making of a society of agents. In the next, section we will explore these points further with an extended example involving a manager and several staff.

5 Extended Example

Here our extended example is based on the example used to illustrate the decision making methodology in [9]. Suppose there is an academic Head of Department (HoD) who must make a series of decisions related to his staff. The HoD has four concerns: he wishes to increase the international reputation of the Department, increase the number and quality of the papers written by members of his Department, develop the careers of his staff and, so, enhance their experience, and to keep his staff happy. The HoD thus has four values: two are common goods, Reputation (R) and Publication (P); and two relate to individuals, Experience (E_1, E_2 and E_3) and Happiness (H_1, H_2 and H_3). The HoD will order these values according to his priorities. Suppose that he is most concerned for Reputation, then for Publication, then for Experience and, finally, for Happiness. Suppose also that, initially, the HoD ranks the Experience and Happiness of all his staff equally. The staff will have the values of Experience and Happiness. If they are ambitious and want to further their careers, they will value Experience over Happiness, otherwise they will prefer Happiness to Experience. We have four agents: The Head of the Department (HoD), and three staff (S1, S2 and S3)

We will use a simplified picture of the HoD's job, in which he can ask people to write papers, send them to conferences and give them places on training courses. In response, the members of staff may attend the training or not, may or may not succeed in writing a paper, and, if they go to a conference, may make an impact, enhancing the reputation of the department, or they may remain relatively anonymous. This gives the following actions, and consequent joint actions.

Possible Actions: $\alpha_1(n)$: HoD Send staff Sn to a conference,
$\alpha_2(n)$: HoD Ask staff Sn to write a paper
$\alpha_3(n)$: HoD Send staff Sn on a training course
$\beta_n, \beta_n\prime$: Staff Sn does well/poorly in the conference
$\gamma_n, \gamma_n\prime$: Staff Sn does /does not write a paper
$\delta_n, \delta_n\prime$: Staff Sn attends /does not attend the training

Possible Joint Actions:
$j1_n : (\alpha_1(n), \beta_n)$, $j2_n : (\alpha_1, \beta_n\prime)$, $j3_n : (\alpha_2, \gamma_n)$, $j4_n : (\alpha_2, \gamma_n\prime)$, $j5_n : (\alpha_3, \delta_n)$,

$j6_n : (\alpha_3, \delta_n\prime)$

Next we need to see how values can be promoted. Staff enjoy attending conferences, (irrespective of how they perform), but do not enjoy training courses. Staff experience is promoted by attending training courses. Publication is promoted whenever staff write papers (writing a paper includes getting it accepted), and Reputation is promoted when a member of staff performs well at a conference. Thus: Values (V) are:
R is promoted by $j1_n$ for any n, **P** is promoted by $j3_n$ for any n
H(n) is promoted by $j1_n$, $j2_n$, and demoted by $j5_n$, **E(n)** is promoted by $j5_n$.

We next need to define how emotions will affect the ranking of values for the HoD. Let us suppose that if the particular HoD likes a member of staff he will give more value to the happiness of that member of staff. Dislike will cause him to value the experience of the member of staff less. If the HoD experiences joy, the happiness of his staff will become more important, and the value with respect to which it is felt will diminish in importance, and conversely with distress, it will increase in importance. The effect on the value order of these changes will depend on the particular agent: some agents will be volatile and react to quite small emotional stimuli, whereas others will respond only to strong stimuli. We can represent this by associating initial weight with each value. The ordering of these weights will give the value order, and the difference between adjacent weights will represent the volatility of the agent with respect to those values. Suppose the initial weights, uninfluenced by emotions, for R, P, E and H are r, p, e and h. Now to give the effects of the emotions in our scenario we can say:

$r_{t_n} = r - J(R, t_{n-1})$. That is, the weight of R decreases as joy with respect to R is experienced, and rises as distress is experienced.

$p_{t_n} = p - J(P, t_{n-1})$. That is, the weight of P decreases as joy with respect to P is experienced, and rises as distress is experienced.

$e_{i_{t_n}} = e - J(E, t_{n-1}) + \min(0, L(i, t_{n-1})))$ That is, the value placed on an agent's experience decreases with joy in respect of experience and dislike for the particular agent.

$h_{i_{t_n}} = h + J(R, t_{n-1}) + J(P, t_{n-1}) + J(E, t_{n-1}) + \max(0, L(i, t_{n-1}))$. That is, weight placed on an agent's happiness increases with joy with respect to the other values and liking for the agent.

In what follows we will subscript the values with weights. We want the agent to initially order the values $R > P > (E_1 = E_2 = E_3) > (H_1 = H_2 = H_3)$. We also want the agent to rank R and P quite closely, and E and H quite closely, but to always prefer both R and P to H and E. The weights of P and R must therefore be substantially higher than those of H and E. For illustrative purposes we can use any weights with the correct properties: we arbitrarily choose $r = 240$, $p = 239$, $e = 7$ and $h = 6$ for our example. Thus the initial Value Order of the HoD, Value Order at t = 0 is:
VO_0: $R_{240} > P_{239} > (E1_7 = E2_7 = E3_7) > (H1_6 = H2_6 = H3_6)$.

The extent of the impact of events on the emotions of gratitude and displea-

sure depends, as discussed in the previous section, on how unexpected the event is. Since he assumes that his staff are ambitious, the HoD will expect that staff attend training courses when they are offered, will hope that they write a paper, but since this is not easy, recognizes that they may fail. Making a significant impact at a conference is quite unusual, and so the accomplishment of j1 will be rather unexpected. Thus, the HoD will, as opportunity arises, given V_0, try to bring about j1, j3 and j5: j5 is expected, (0.9) j3 is thought likely (0.6) and j1 is very unexpected (0.1).

In this paper, we will only show how emotions will affect the value ordering, the other steps of the decision making process are fully discussed in [9].

Suppose, on taking over, the HoD has the chance to send all three members of staff on training courses. He will do so and expect them to attend. The state reached by $j5_n$ will promote the HoD's fourth most important value (remember the experience of each member of staff is ranked equally, so their position is averaged), and demote his seventh most important value, since the staff do not enjoy the training, giving the state a value of 3, whereas $j6_n$ will have a value of 2. Thus the HoD will feel gratitude with an intensity of 0.3 if a member of staff attends and displeasure with intensity 1.8 if the staff member does not attend.

Suppose S1 and S2 are suitably ambitious and so do attend, but S3 prefers happiness to experience and therefore chooses to miss the course. Because attendance was expected, the HoD will feel a little gratitude towards S1 and S2, but considerable displeasure towards S3. This also means that the HoD will feel distress rather than joy. The HoD emotional state is now:
$L_{HoD}(S1, 1) = 0.3; L_{HoD}(S2, 1) = 0.3; L_{HoD}(S3, 1) = -1.8; J_{HoD}(E, 1) = -1.2$

His value order at time 1 then becomes

VO_1: $R_{240} > P_{239} > (E1_{8.2} = E2_{8.2}) > E3_{6.4} > (H1_{5.1} = H2_{5.1}) > H3_{4.8}$.

Now the HoD will be inclined to send S1 and S2 to conferences in preference to S3. Next, suppose a conference does occur, to which only two people can be sent. Since the HoD now prefers the happiness of S1 and S2 to S3, he will choose to send them. At the conference, S1 performs very well and impresses a number of leading academics, whereas S2 is an undistinguished participant. Thus j1 is achieved in respect of S1 and j2 in respect of S2. With respect to S1 R and H_1 are promoted. The state reached by S1 has an importance of 10.5 to the HoD. Since the unexpectedness was 0.9, the intensity of gratitude towards S1 is 9.45. With regard to S2, the state reached has importance 5.5 and unexpectedness 0.1: thus the intensity of displeasure towards S2 is 0.55. This gives the HoD emotional state as:
$L_{HoD}(S1, 2) = 9.75; L_{HoD}(S2, 2) = -0.25; L_{HoD}(S3, 2) = -1.8; J_{HoD}(E, 2) = -1.2; J_{HoD}(R, 2) = 9.5$

Now the HoD has experienced considerable joy: the success of S1 at the conference leads him to raise the importance of happiness for his staff. Also it lowers the importance of R : VO_2: $P_{239} > R_{230.5} > H1_{24.05} > (H2_{14.3} = H3_{14.3}) > E1_{8.2} > E2_{7.95} > E3_{6.4}$.

Now the new importance of P means that the HoD turns his focus to publication and asks all three staff to write papers. Suppose that S1 and S3 succeed, but

S2 is unable to produce a worthwhile paper. The importance of the paper being written was 8, and the importance of failure was 8, since no other values were promoted by reaching the alternative state. Since the unexpectedness of success was 0.4, this gives rise to gratitude with respect to S1 and S3 with intensity 3.2 and displeasure towards S2 with intensity 4.8. The HoD's value emotional state becomes:
$L_{HoD}(S1,3) = 12.95; L_{HoD}(S2,3) = -5.05; L_{HoD}(S3,3) = 1.4; J_{HoD}(E,3) = -1.2; J_{HoD}(R,3) = 9.5; J_{HoD}(P,3) = 1.6$

Now S2 has failed to live up to his initial good impression: successive failures have obliterated it. S3 meanwhile has redeemed himself, and is now in favor.

VO_3: $P_{237.4} > R_{230.5} > H1_{28.85} > H3_{17.3} > H2_{15.9} > (E1_{8.2} = E3_{8.2}) > E2_{1.95}$).

This is, necessarily, a simplified and somewhat idealised example, but nevertheless it does serve to illustrate some important points. First, we can see that when it comes to ordering agents within a value even quite small changes can have an impact, but something rather significant is required to affect the ordering of the values themselves. Where values are ranked equally, a slight change will make a difference: this is useful in that equally ranked values impede the ability of the decision making approach to reach definite conclusions, meaning that decisions may be arbitrary. Second, we can note that the short-term promotion of happiness by staff 3 when he missed the training course, was not in his long-term interests since, subsequently, he was overlooked for conference attendance. Thus, the mechanism fosters cooperation with the wishes of the HoD. Thirdly, we can see that the effect of joy is to enable a shift in focus: the success of S1 at the conference meant that reputation did not need to be given such a high priority. Finally, the example suggests that some decay on joy may be required: although the short-term shift in focus was desirable, we might prefer to see the standard preferences reasserting themselves after a while.

6 Concluding Remarks

In this paper, we have shown how we can integrate emotions into a rational decision making process by which an agent chooses between alternative courses of action, building on previous work, especially [13] for emotions and [9] for decision making. Whereas [13] focused on the role of hope and fear, however, we have in particular investigated the emotions of gratitude and displeasure, like and dislike, joy and distress. These emotions seem to play a beneficial role in decision making in that:

1. Emotions can motivate choice between actions which seem equally desirable from a purely rational perspective
2. The motivation in 1 provides a mechanism which fosters cooperation, since cooperative agents will be favored over less cooperative agents
3. Emotions provide a mechanism by which an agent can re-evaluate its priorities, avoiding a single-minded pursuit of one goal to the exclusion of other desirable goals.

Although we have described the mechanism, and characterised formally in terms of our decision making model, when these emotions arise, we have only sketched the more quantitative aspects. Currently, we have implemented the decision making model without emotions. Our next step will be to add emotions, so that we can get a better understanding of the acceptable settings for the crucial parameters which govern the intensity of emotions and the extent of its effect on the decisions of the agent.

References

1. K. Atkinson and T. J. M. Bench-Capon. Practical reasoning as presumptive argumentation using action based alternating transition systems. *Artif. Intell.*, 171(10-15):855–874, 2007.
2. K. Atkinson, P. McBurney, and T. Bench-Capon. Computational representation of practical arguments. *Knowledge, rationality and action*, pages 191–240, 2006.
3. T. BenchCapon. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.
4. M. E. Bratman, D. J. Israel, and M. E. Pollack. Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4:349–355, 1988.
5. A. Damasio. *Descartes' Error*. G P Putnams Sons, 1994.
6. P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming. In *Artificial Intelligence*, 1995.
7. H. Jiang, J. M. Vidal, and M. N. Huhns. Ebdi: an architecture for emotional agents. In *Proceedings of AAMAS 2007*, pages 1–3, New York, USA, 2007. ACM.
8. J. J. Meyer. Towards a quantitative model of emotions for intelligent agents. In *Emotions and computing*, 2007.
9. F. S. Nawwab, T. J. M. Bench-Capon, and P. E. Dunne. A methodology for action-selection using value-based argumentation. In P. Besnard, S. Doutre, and A. Hunter, editors, *Proceedings of COMMA 2008*, volume 172 of *Frontiers in Artificial Intelligence and Applications*, pages 264–275. IOS Press, 2008.
10. A. Ortony, G. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.
11. L. Padgham and G. Taylor. A system for modelling agents having emotion and personality. In L. Cavedon, A. S. Rao, and W. Wobcke, editors, *PRICAI Workshop on Intelligent Agent Systems*, volume 1209 of *Lecture Notes in Computer Science*, pages 59–71. Springer, 1996.
12. W. S. Reilly. *Believable social and emotional agents*. PhD thesis, Carnegie Mellon University (CMU), 1996.
13. B. R. Steunebrink, M. Dastani, and J.-J. C. Meyer. A logic of emotions for intelligent agents. In *AAAI*, pages 142–147. AAAI Press, 2007.
14. M. Wooldridge and W. van der Hoek. On obligations and normative ability: Towards a logical analysis of the social contract. *J. Applied Logic*, 3(3-4):396–420, 2005.