

## Emotional Consensus in Group Decision Making

Paul Thagard and Fred W. Kroon

University of Waterloo

[pthagard@uwaterloo.ca](mailto:pthagard@uwaterloo.ca)

Thagard, P., & Kroon, F. W. (forthcoming). Emotional consensus in group decision making. *Mind and Society*.

**Abstract.** This paper presents a theory and computational model of the role of emotions in group decision making. After reviewing the role of emotions in individual decision making, it describes social and psychological mechanisms by which emotional and other information is transmitted between individuals. The processes by which these mechanisms can contribute to group consensus are modeled computationally using a program, HOTCO 3, which has been used to simulate simple cases of emotion-based group decision making.

### 1. Introduction

How do you and your friends decide what movies to attend together? Do you create a chart of the movies available, each rate them numerically, and then sum the ratings to make the group decision? Or do you discuss the available movies until everyone has a good feeling about the best movie to go to? Similarly, if you are in a North American or other university department that does its own faculty hiring, how does your department proceed? Does it use either of the following procedures? Procedure A: jointly produce a rating of candidates on a scale of 1 to 10 for their potential contributions to research, teaching, and administration; jointly agree on a weighting of the importance of research, teaching, and administration; multiply the ratings time the weightings to produce a score for each candidate, and make an offer to the first one. Procedure B: argue heatedly about the strengths and weaknesses of different candidates and the value of different kinds of research and teaching; gradually reach some

consensus, or at least take a vote on which candidate should be made an offer; try to mollify those who did not get the candidate they wanted. Procedure A sounds rational, but we have never heard of it being conducted, whereas something like procedure B is very common.

Movie going and academic hiring are just two examples of the pervasive phenomenon of group decision making, which abounds in organizations as diverse as families, businesses, juries, and political parties. Unless there is an autocrat empowered to make decisions for the whole group, a joint decision will often require extensive discussion and negotiation before a consensus or majority opinion is reached. In ideal cases, the discussion produces sufficient communication that all members of the group come to share the same view of what to do, for example all agreeing about which job candidate is the best.

Group decisions are often highly emotional. In academic hiring, different members of a department are usually enthusiastic about some candidates and scornful about others. These emotional attitudes partly reflect the records of the candidates, but also reflect the values of the professors doing the hiring, who often have strong opinions about the value of particular aspects of academic work, for example teaching versus research, or research on some topics rather than others. It is not coincidental that professors tend to most admire the kinds of work that they know best, so conflict arises in a department because everyone wants to hire replicates of themselves. Ideally, conflict is resolved by acquiring and exchanging sufficient information that the members of the department arrive at similar emotional reactions toward the different candidates.

This paper presents a theory and computational model of group decision making understood as emotional consensus. It proposes that individual decision making is inherently emotional, and that group decision making requires communication of emotional as well as factual information. In the ideal case, sufficient communication of both facts and emotions will lead to cognitive/emotional consensus about what to do, a consensus that consists of at least partial convergence of both beliefs and emotional values. After reviewing the emotional character of individual decision making, we describe a set of mechanisms for transmitting emotional values between people. Some of these mechanisms are implemented in a computational model that simulates how emotional communication can lead to consensus in group decisions.

## **2. Emotions in Individual Decision Making**

In the past decade, psychologists and neuroscientists have increasingly recognized the inherently emotional nature of decision making (e.g. Damasio, 1994; Finucane, et al., 2000; Lerner, Small, and Loewenstein, 2004; Loewenstein, et al. 2001). For most people, decisions are not the result of a cognitive calculation of the sort described in normative models such as multiattribute utility theory, but rather the result of arriving at emotional reactions to different situations. Preferences arise in favor of options associated with strong positive emotions, and against options associated with strong negative emotions.

Thagard (2000) proposed a theory of emotional coherence that we will now review, because it provides the individual basis for the theory of group decision making to be developed in the next section. On this theory, the elements in a decision include representations of competing actions and goals that they potentially accomplish. A pair

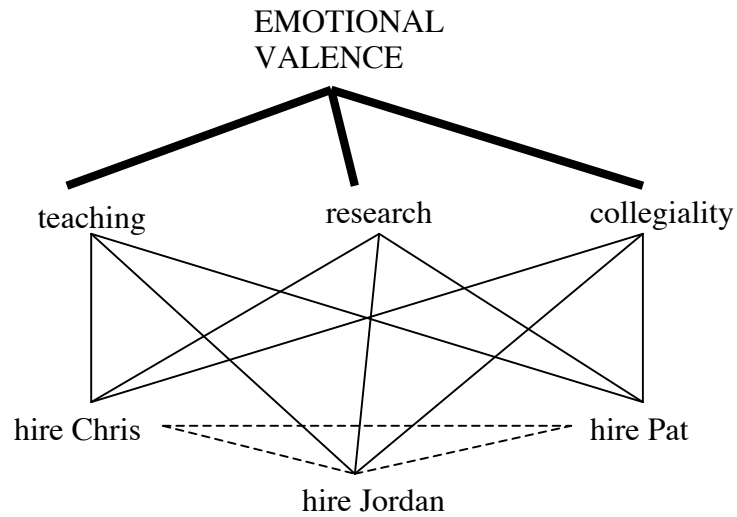
of elements may cohere, as when an action and a goal fit together because the action facilitates the goal. If two elements cohere, then there is a positive constraint between them that can be satisfied either by accepting both elements or by rejecting both. On the other hand, two elements may “incohere”, as when two actions resist fitting together because they are incompatible with each other. Between incoherent elements, there is a negative constraint that can be satisfied only by accepting one and rejecting the other. In a purely cognitive coherence problem, the elements are divided into ones that are accepted and rejected in such a way as to satisfy the most constraints. But in an emotional coherence problem the elements also have positive and negative *valences*, which reflect the emotional attitudes associated with the different representations. Moreover, elements can have positive or negative emotional connections to other elements, so that the valence of one element can influence the valence of the other. The valence of an element is determined by the valences and acceptability of all the elements to which it is connected. Thagard (forthcoming) describes how the adoption and maintenance of religious beliefs and practices can be affected by emotional coherence.

This characterization of emotional coherence is rather abstract, but can be made clearer by means of a concrete example and a computational model. Suppose you are trying to decide whom your department should hire and you are focusing on three candidates: Chris, Jordan, and Pat. You have several main goals that you are trying to accomplish, including acquiring a colleague who will teach well, do lots of interesting research, and be reasonably easy to get along with. Then the elements of your decision are the representations: hire Chris, hire Jordan, hire Pat, good teaching, productive research, and collegiality. The positive constraints among these elements reflect the

extent to which the three candidates accomplish your different goals. The negative constraints arise because you only have one position, so that hiring one candidate means not hiring the others. See the coherence theory of decision developed by Thagard and Millgram (1995; Millgram and Thagard, 1996).

So far, this just sounds like standard multiattribute decision making with different jargon, but the differences will become apparent when the emotional and computational aspects of the decision come into view. It is possible that your decision about whom you want to hire is coldly cognitive, but if you are like most people you will have definite emotional reactions to the candidates, ranging from enthusiasm to disgust. How do these reactions arise?

According to the theory of emotional coherence, the reactions result from emotional valences attached to the goals that affect the decision making. If you think that research is extremely important and care little about teaching and social interaction, then you will end up with a positive emotional reaction toward the candidate who is the best researcher. The positive valence of the goal thus transfers to the action, so you feel good at the prospect of hiring the candidate who excels in research. Normally, however, decisions involve a trade-off among different goals, so you may have to balance a fairly good researcher off against a superb teacher. Figure 1 shows the structure of the constraint network involving the actions, goals, and their emotional inputs.



**Figure 1.** Structure of the constraint network for the hiring decision.

Thick lines indicate emotional strength of various goals, and thin solid lines indicate the extent to which the different hiring options contribute to those goals. Dotted lines indicate negative constraints between pairs of options. Not shown are weights that indicated the strength of the various constrains.

Decision problems like those in figure 1 are naturally solved in a computationally powerful and psychologically natural way by artificial neural networks. HOTCO (for HOT COherence) is a computational model of emotional inference that has been applied to various kinds of thinking, including trust, legal reasoning, and emotional analogies (Thagard, 2000, 2003; Thagard and Shelley, 2001). HOTCO incorporates and expands a group of connectionist models of high-level inference that perform such tasks as explanatory inference (ECHO), analogical mapping (ACME), decision making (DECO), and conceptual coherence (Kunda and Thagard, 1996). In each of these models, elements

are represented by units, which are very simple artificial neurons whose activation is interpreted as a measure of their degree of acceptability. In HOTCO, each unit also has a *valence*, a positive or negative numerical measure of the emotional appeal of what is represented by the unit. Acceptability and valence can be very different: one has to accept the proposition that George W. Bush is president of the United States in 2004, but may be unhappy about it. In the original version of HOTCO (Thagard, 2000), activations were allowed to influence valences but not vice versa, but in the revised version, HOTCO 2, valences were also allowed to influence activations. For belief evaluation, this influence is psychologically common, although it can lead to irrational judgments corresponding to motivated inference and wishful thinking, in which beliefs are adopted because they fit with personal goals and emotions rather than evidence. In decision making, however, acceptability and valence usually coincide: the most acceptable action is usually the one that has the most emotional support, although there are unusual circumstances in which people may decide that an action should be performed even though they do not like it. Mathematical details of how HOTCO 2 works are provided in the appendix to Thagard (2003).

HOTCO is not very neurologically realistic, in that it uses single artificial neurons to stand for complex representations such as “hire Chris”, and it does not take into account how cognitive and emotional processing is carried out by different brain areas. However, HOTCO is compatible with a much more detailed neurocomputational model, GAGE, which models decision making using distributed representations over spiking neurons that are organized into anatomical groups, including the prefrontal cortex, hippocampus, amygdala, and nucleus accumbens (Wagar and Thagard, 2004). Each

HOTCO unit that stands for a high-level representation can be viewed as corresponding to groups of connected neurons in GAGE, including ones in the prefrontal cortex whose joint spiking activity corresponds to the unit's activation, and ones in the amygdala and nucleus accumbens whose joint spiking activity corresponds to the unit's valence. HOTCO uses activations and valences of units to integrate cognition and emotion, and GAGE uses firing behavior of groups of spiking neurons in different brain areas to accomplish the same task in a more neurologically realistic way. However, GAGE is much more complex to program and run simulations, so for the purpose of modeling group decisions we will use HOTCO instead, viewing it as a useful approximation to GAGE's closer approximation to brain activity.

In section 4, we will describe HOTCO 3, which expands HOTCO 2 into a multiagent system with emotional communication. HOTCO 2 only models a single agent, operating alone by virtue of a neural network that generates emotional decisions. In contrast, HOTCO 3 models any number of agents, each one of which is an agent running a HOTCO 2 simulation. The multiple agents in HOTCO 3 communicate and shape each other's decisions by transmitting emotional information, a social process that we will now characterize.

### **3. Social Mechanisms of Emotional Transmission**

Group decisions that are not based on autocratic pronouncement or simple voting require achievement of some sort of consensus. In a purely cognitive model of consensus concerning what to believe, group agreement can be reached simply by exchanging information about relevant hypotheses, evidence, and explanatory relations between them (Thagard, 2000, ch. 7). However, if decision making by individuals is



inherently emotional, then group decision making will require attainment of a kind of emotional consensus in which members of the group share similar positive and negative feelings about different actions and goals. Cognitive consensus can come about by verbally exchanging propositional information, but emotional communication is much more complicated. In this section we will describe what seem to be the most important mechanisms of emotional transmission leading toward consensus, and in the following describe how some of them can be implemented computationally.

What is a social mechanism? In general, a mechanism consists of “entities and activities organized such that they are productive of regular changes” (Machamer, Darden, and Craver, 2000, p. 3). Familiar examples are machines such as bicycles that have parts interacting in ways that produce desired results, and biological systems such as organs that have parts such as cells interacting in ways that perform functions useful to the organism involved. In social mechanisms, the parts are human agents, the activities are communication and other interactions between agents, and the regular changes involve groups of humans.

Suppose now that we have a university department whose members need to make a joint decision about whom to hire, and that each member has made a preliminary choice based on emotional coherence. If there is no general agreement, then the people need to talk to each other in order to try to convince their colleagues to hire their preferred candidate. Communication of emotional valences can take place by very different kinds of social mechanisms, including means-ends and analogical arguments, emotional contagion, altruism, and empathy. Each of these mechanisms can be characterized by four parts: a *trigger* that initiates the mechanism by which a sender influences the

emotional decision making of a receiver, *inputs* that the sender gives to the receiver, *changes* in the mental state of the receiver, and *effects* on the receiver's decision making. These mechanisms are social in that they involve interactions of 2 or more agents, but they do not require that the sender intentionally produce the communication. For example, with emotional contagion senders may not even be aware that they are communicating their enthusiasm or other reaction to receivers, and with altruism all that matters is that receivers care enough about the senders to pick up some of their values. The social mechanisms of emotional communication are processes that involve interactions between agents, and are only sometimes intentional strategies that a sender uses to convince a receiver.

The most familiar kind of communication in group decisions is verbal argument, which is triggered by the desire of one person, the sender, to influence the decision of another person, the receiver. The most straightforward kind of argument is means-ends, in which the sender tries to convince the receiver of an action by showing the receiver that the action contributes to one or more of the receiver's goals. For example, in hiring, the sender may argue that the department should hire Chris because of fine teaching, on the assumption that the receiver has the goal of having a department with good teachers. Here the input to the receiver is the verbal claim that Chris is a good teacher, and the desired change is that the receiver should acquire a stronger association between the action of hiring Chris and the goal of teaching. If this change is made, then it could have the desired effect on the receiver's decision making, producing agreement between the receiver and sender and moving the department toward consensus.

Analogical argument operates more indirectly by comparing prospective actions to ones that have worked well or badly in the past. In a hiring discussion, someone might argue for Pat as “the next Herbert Simon,” drawing an analogy between the intellectual abilities and accomplishments of Pat and those of a distinguished researcher. Alternatively, if department members are familiar with a previous disastrous hire, Schmerdly, the sender may use a negative analogy and say that they should not hire Jordan who is another Schmerdly. In both these cases, the inputs are the verbal comparison that the sender gives to the receiver, and the desired changes are rearrangements of the receiver’s cognitive/emotional network to make a connection between the source analog (Simon or Schmerdly) and the target analog, Pat. This connection makes possible the spread of the positive or negative valence of the source to the target, increasing or decreasing its valence, which may have the desired effect of making the target action more or less emotionally attractive. See Thagard and Shelley (2001) for an extended discussion of emotional analogies, and Blanchette and Dunbar (2001) for empirical evidence that analogies are often used in persuasion.

Means-ends and analogical argument are verbal means by which senders attempt to modify the emotional reactions of receivers, but non-verbal means may be more effective. Hatfield, Cacioppo, and Rapson (1994) provided an extensive discussion of emotional contagion, in which one person “catches” another’s emotions. Their theory is summarized in the following propositions (pp. 10-11):

Proposition 1. In conversation, people tend automatically and continuously to mimic and synchronize their movements with the facial

expressions, voices, postures, movements, and instrumental behavior of others.

Proposition 2. Subjective emotional experiences are affected, moment to moment, by the activation and/or feedback from such mimicry.

Proposition 3. Given propositions 1 and 2, people tend to “catch” others’ emotions, moment to moment.

The trigger for the mechanism of emotional contagion is face-to-face meetings between individuals. The inputs that the sender provides to the receiver are non-verbal, consisting of bodily states that the receiver unconsciously imitates. Then the receiver undergoes physiological changes that change the psychological valences associated with different aspects of the situation, with the possible effect that the overall decision of the receiver shifts. Unlike means-ends and analogical arguments, contagion is a social process that is usually not an intentional strategy used by a sender to convince a receiver.

In hiring situations, emotional contagion can work with both positive and negative emotions. If there is a discussion of a candidate, and one person is displaying great enthusiasm by smiling and displaying positive body language, then there will be a tendency for others to catch some of the enthusiasm. On the other hand, if people exhibit skepticism about a candidate by scowling, sneering, or turning away, then others may pick up some of the negative attitude toward the candidate. Barsade (2002) provides an insightful discussion of the contribution of emotional contagion to group processes.

Other social mechanisms for emotional transmission - altruism, compassion, and empathy - may combine verbal and non-verbal information. Altruism is the unselfish

regard for the welfare of others. If you care about a person, your altruism may lead you to acquire an emotional attitude toward something that is valued by that person. Altruistic transmission is triggered in situations where the receiver cares for the sender, and receives verbal or non-verbal input concerning what the sender's goals are. Then the altruistic receiver may change by acquiring at least a weak version of the sender's goals, changing the receiver's emotional constraint network in a way that may have the effect of generating a different decision. For example, if you have a colleague you care about who desperately wants to hire a candidate who would be a useful collaborator, then you may adopt such a goal as a way of helping out your colleague. You then come to the same decision as your colleague in part because you have adopted his or her goals as your own.

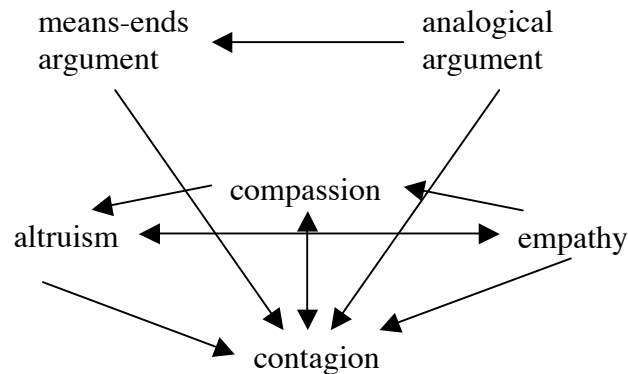
Compassion is triggered when you notice that someone is distressed, and might be viewed as a special kind of altruism. It is probably not a factor in faculty hiring, but could easily affect personal and political decisions. Whereas altruism assumes that you care enough generally for people to adopt their goals, compassion requires only that you do not want to see them suffer. The input to the compassion mechanism is verbal or nonverbal information conveyed from the sender to the receiver concerning the sender's distress. This information then augments the receiver's emotional constraint network to include the goal of relieving the sender's distress. The result may be a shift in the receiver's decision making toward an action that is at least partially in the interest of the sender.

The most complicated mechanism of emotional transfer is empathy, which we will treat as a special form of analogy (Barnes and Thagard, 1997; Thagard and Shelley,

2001). Empathy involves putting yourself in other people's shoes so that you understand their emotions by virtue of having an experience that is an image of theirs. The trigger for this mechanism is information about the sender's situation that reminds the receiver of a similar situation that the receiver had personally experienced emotionally. Then the receiver can infer from the similarities between the two people and their situations that the sender is experiencing something like what the receiver did. In empathy, this inference is not verbal, but rather involves projecting the receiver's imagined experience back on to the sender. Empathy does not by itself change the emotional constraint network of the receiver, but may inspire altruistic feelings that do. For example, suppose you and I disagree about whom to hire, and you say that you feel intellectually isolated in the department. If I remember another department in which I felt distressed about being intellectually isolated, then I may come to appreciate your distress, and then through compassion and altruism come to adopt the goal of relieving your intellectual isolation. The new emotional constraint network that incorporates the goal of relieving your distress that empathy has enabled me to perceive may then lead me to prefer a decision similar to the one you prefer.

In any realistic social situation, the mechanisms of means-ends argument, analogical argument, emotional contagion, altruism, compassion, and empathy may interact with each other. For example, a means-ends argument may inspire enthusiasm in one person whose excitement spreads by emotional contagion to another. Or, as we have just suggested, empathy may generate compassion and altruism. Figure 2 summarizes some of the ways in which the mechanisms of emotional transmission may interact. Even though it shows that all other mechanisms may influence contagion, this

should not be taken to mean that contagion is somehow the fundamental process of emotional communication. Each of the other mechanisms can be used to influence the emotional state of a receiver independently of whether contagion is also operating as well. How frequently the different mechanisms are used in group decisions, and how frequently they interact, are empirical questions, currently unanswered. In order to make discussion of these mechanisms more precise, we now describe a computational model of emotional group decision making.



**Figure 2.** Interactions among mechanisms of emotional transmission.

#### 4. A Computational Model of Emotional Consensus

In order to get a deeper understanding of the social mechanisms that can lead to emotional consensus, we have constructed a computational model, HOTCO 3, which builds on earlier models of emotional coherence and cognitive consensus (Thagard, 2000, 2003). This section describes the data structures and algorithms of HOTCO 3, and their application to simulations of a movie decision and a hiring decision. HOTCO 3 is programmed in Common LISP.

The basic structure in HOTCO 3 is a person, whose properties include a group that the person belongs to, a list of people that the person cares for, and a set of inputs

concerning actions, goals, facilitation relations, and emotional valences. For example, the facilitation relations for a moviegoer might include (FACILITATE 'ACTION-MOVIE 'EXCITEMENT) which represents the proposition that going to an action movie can contribute to this person's goal of being excited. If the person's emotional valences include EXCITEMENT, then this input will encourage the spread of positive valence to ACTION-MOVIE. For modeling each person's decision making, the basic data structures are units, which are highly artificial neurons that possess activations and valences, which are real numbers between 1 and -1. Thus there will be a unit to stand for ACTION-MOVIE, EXCITEMENT, and other actions and goals. These units are connected by links, including excitatory links created by facilitation relations such as the one between ACTION-MOVIE and EXCITEMENT, and inhibitory links created by incompatibility relations, such as the one between ACTION-MOVIE and ROMANTIC-COMEDY.

The algorithms for individual decision making are mostly those for previous versions of HOTCO. Units for the person's fundamental goals receive activations and valences from an external source, and these spread to other connected units. Parallel updating leads after a reasonable number of iterations (typically 50-100) to stability of activations and valences that represent the overall decision of the individual. Units with high activations and valences are interpreted as indicating which action or actions the individual choose to perform, while units with low activations and valences stand for rejected choices.

What is novel about HOTCO 3 is that it simulates decision making by a group of interacting emotional agents. The main procedure works as follows:



1. Perform individual decision making for each member of the group.
2. If all members agree about what to do, then stop: consensus has been reached.
3. Otherwise, simulate a meeting to exchange factual and emotional information between two randomly selected members of the group.
4. Reevaluate individual decision based on emotional coherence for both newly informed members.
5. Repeat from step 2.

The key step in this process is the third, information exchange, which should include all of the mechanisms described in section 3, although currently only means-ends, contagion, and altruism are implemented in HOTCO 3. Contagion works by identifying actions on which a group member meeting with another has a strong positive or negative emotional valence. A member with a strong valence becomes a sender whose emotional reactions are received in weakened form by the receiver. Obviously, HOTCO 3 lacks the physiological processes by which contagion takes place in people, but it simulates the result of having some part of the strong emotional reactions of a sender transferred to a receiver. Contagion in HOTCO 3 works by providing a unit representing the action for the receiver an emotional input that is a reduced version of the strong valence that the sender attaches to the action. Thus emotional valence spreads by contagion from the sender to the receiver. For example, if the sender attaches strong positive valence to that unit for ACTION-MOVIE, then a new link will encourage spread of positive valence to the receiver's unit for ACTION-MOVIE. Thus the receiver acquires by contagion a weakened version of the sender's attitudes towards various actions.

The altruism mechanism in HOTCO 3 affects goals rather than actions. A participant in a meeting who cares for the other member becomes a receiver in the altruism process. The other member becomes a sender, with the effect that part of the emotional valence attached to the sender's goals is transferred to the receiver's goals. This is accomplished by creating a new link between the unit that represents a goal of the receiver and the unit for emotional input. For example, if the sender has the valenced goal EXCITEMENT, then the receiver will acquire some emotional input to the unit that represents this goal. Thus, because the receiver cares for the sender, the sender appreciates, at least weakly, the sender's goals. The altruism procedure also transfers from the sender to the receiver information about what actions facilitate the receiver's newly acquired goals, making it more likely that the receiver will reach the same decision as the sender. The compassion procedure is not yet implemented in HOTCO 3, but would work much like the altruism procedure except that it would be triggered by perception of a sender's distress.

HOTCO 3 models a kind of means-ends argument by having a sender notice that the receiver has a goal that can be accomplished by one of the sender's preferred actions. For example, if the receiver has the goal EXCITEMENT, but has not noticed that ACTION-MOVIE facilitates it, then the sender can transfer to the receiver the information that ACTION-MOVIE facilitates EXCITEMENT. Then the LISP expression (FACILITATES 'ACTION-MOVE 'EXCITEMENT) becomes part of the input that determines the decision made by the receiver. This is not a full-blown argument, since HOTCO 3 does not do natural language processing, but it has the same effect in that the sender is trying to convince the receiver to do what the sender wants by

pointing out to the receiver that the action favored by the sender actually contributes to the goals of the receiver. Means-ends argument is a different mechanism from altruism, in which the receiver adopts the senders goals in addition to input about facilitation.

For the three mechanisms currently implemented in HOTCO 3, Table 1 summarizes the triggering conditions, inputs that the sender transmits to the receiver, and changes and effects experienced by the receiver. Only in the case of means-ends argument is emotional transmission triggered by the conscious intention of the sender to convince the receiver.

	<i>trigger</i>	<i>inputs from sender to receiver</i>	<i>changes in and effects on receiver</i>
<i>contagion</i>	receiver perceives sender	sender's valence of action transferred to receiver	receiver acquires some of sender's attitude to action
<i>altruism</i>	receiver cares about sender	sender's valence of goal transferred to receiver	receiver acquires some of sender's attitude to goal
<i>means-ends</i>	sender tries to convince receiver	information that an action facilitates a goal	receiver sees greater attractiveness of goal

**Table 1.** Summary of mechanisms of emotional transfer.

We have not yet implemented analogical argument and empathy, because of technical complications in adapting the analogical mapping program ACME (Holyoak and Thagard, 1989) to use the kind of information available to HOTCO 3. But it is easy to see how they would work in principle. In analogical argument, a sender would try to convince a receiver to change decisions by communicating to the receiver a description of a situation similar to the one at issue. Then the receiver would make new analogically derived connections that would shift the emotional inputs in a way that could result in a different decision. In empathy, the receiver would end up modifying a decision because

of noticing an analogical connection between the sender's situation and some previous emotional experience of the sender. In both analogical argument and empathy, the receiver would end up with a modified emotional constraint network that could increase the likelihood of consensus.

There are undoubtedly other social mechanisms of emotional transmission that we have not considered. For example, senders who want to increase the likelihood that receivers will acquire their emotional attitudes can employ social strategies such as providing compliments, gifts, or sexual favors. Such behaviors are beyond the scope of HOTCO 3, although certainly prevalent in human activity. In the conclusion, we briefly discuss coercive social processes that can also influence group consensus.

## 5. Simulations

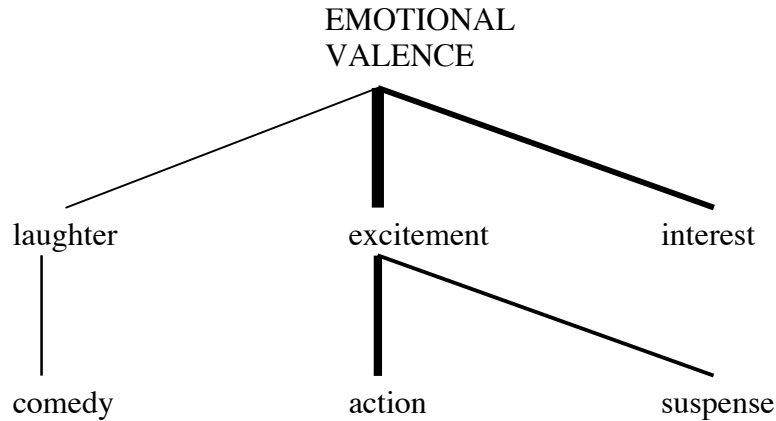
To test the HOTCO 3 implementation of emotional contagion, altruism, and means-ends reasoning, we have used it to simulate simple cases of group decision making. The first concerns a male-female couple making a joint decision about what movie to attend. In accord with their stereotypical preferences, we will call them Mars and Venus; they need to choose between going to an action, comedy, or suspense movie. They are in conflict because Mars mostly wants excitement and so prefers an action movie, whereas Venus mostly wants laughter and so prefers a comedy. The input to HOTCO 3 gives Mars the following emotional valences for various goals: 1 for excitement, .8 for interest, and .5 for laughter. His beliefs about how to accomplish these goals are represented by the following facilitation relations:

```
(facilitate 'go-action 'M-excitement 1)
```

```
(facilitate 'go-suspense 'M-excitement .6)
```

(facilitate 'go-comedy 'M-laughter .5)

The first of these indicates that Mars believes that going to an action movie would maximally contribute to his goal of excitement. The structure of the emotional coherence network for Mars is displayed in figure 3.



**Figure 3.** Emotional constraint network for Mars. The thickness of the lines corresponds to the strength of the connections. Not shown are inhibitory connections among the three actions.

The structure of the coherence network for Venus is similar, but she attaches different valences to the goals: 1 for laughter, .8 for interest, and .3 for excitement. Here beliefs about how different options would contributed to these goals are represented by the following facilitation relations:

(facilitate 'go-action 'V-excitement .2)

(facilitate 'go-suspense 'V-excitement .5)

(facilitate 'go-comedy 'V-laughter 1)

The inputs for the Mars and Venus decisions distinguish M-excitement, which is Mars' excitement, from V-excitement, which is Venus'; this assumes that Mars and Venus are

initially concerned with their own goal satisfaction rather than the other's. It is obvious from the above inputs that Mars will initially want to go to an action movie, whereas Venus will want to go to a comedy. How can the conflict be resolved?

We have run this example using both contagion and altruism, and the results surprised us. Mars and Venus had repeated meetings to exchange emotional information. With contagion alone, Mars and Venus acquired some of the positive and negative valences that the other attached to the various options. For example, Mars attached a new positive valence to going to a comedy. But contagion did not change their decisions: Mars still preferred the action movie, and Venus the comedy. Similarly, altruism did not produce a consensus, even though Mars adopted some of Venus's goals such as V-excitement, i.e. excitement for Venus as well as himself, and Venus adopted some of Mars' goals. Curiously, however, when emotional transmission included both contagion and altruism, Mars changed his decision and agreed to go to the comedy. He was influenced both by contagion of Venus's desire to go to the comedy, and by his altruistic adoption of her goals. Figure 4 graphs how meetings lead to consensus when both contagion and altruism influence Mars.

An extension to the Mars-Venus example shows that the phenomenon scales to a decision with three agents. A third agent, "Neptune" (of indeterminate gender), was added to the original simulation. Neptune's valences were 1 for interest, .6 for excitement and .4 for laughter. The facilitation relations were:

```
(facilitate 'go-action 'N-excitement .6)
```

```
(facilitate 'go-suspense 'N-interest 1)
```

```
(facilitate 'go-comedy 'N-laughter .5)
```

As a result, Neptune's first choice is to see a suspense film. As with the two agent Mars-Venus example, when contagion was the only form of emotional communication, the group fails to reach consensus. Restricting emotional communication to altruism gives similar results. However, when the two are combined the group reaches consensus in just a few meetings. The exact number varies, since, unlike a group with only two members, the order of meetings may be different in different simulation runs. This possibility adds additional complexity to larger groups, which we investigated in another example concerning hiring.

This second case is a highly simplified approximation to a university department making a faculty hiring decision. The department has only three members, Professors A, B, and C, and they are trying to choose collectively between two candidates, Pat and Chris. The professors differ in their emotional attachments to different hiring criteria: A cares most about research, B about teaching, and C about administration. They have different beliefs about how the candidates would contribute to these goals. For example, Professor A attaches full 1.0 emotional valence to research, in contrast to .8 for teaching and .5 for administration. The comparable valences for Professor B, respectively, are .5, 1, and .5, and for Professor C, .5, .5, and .8. Here are the facilitation relations that represent their beliefs about the candidates:

Prof A:

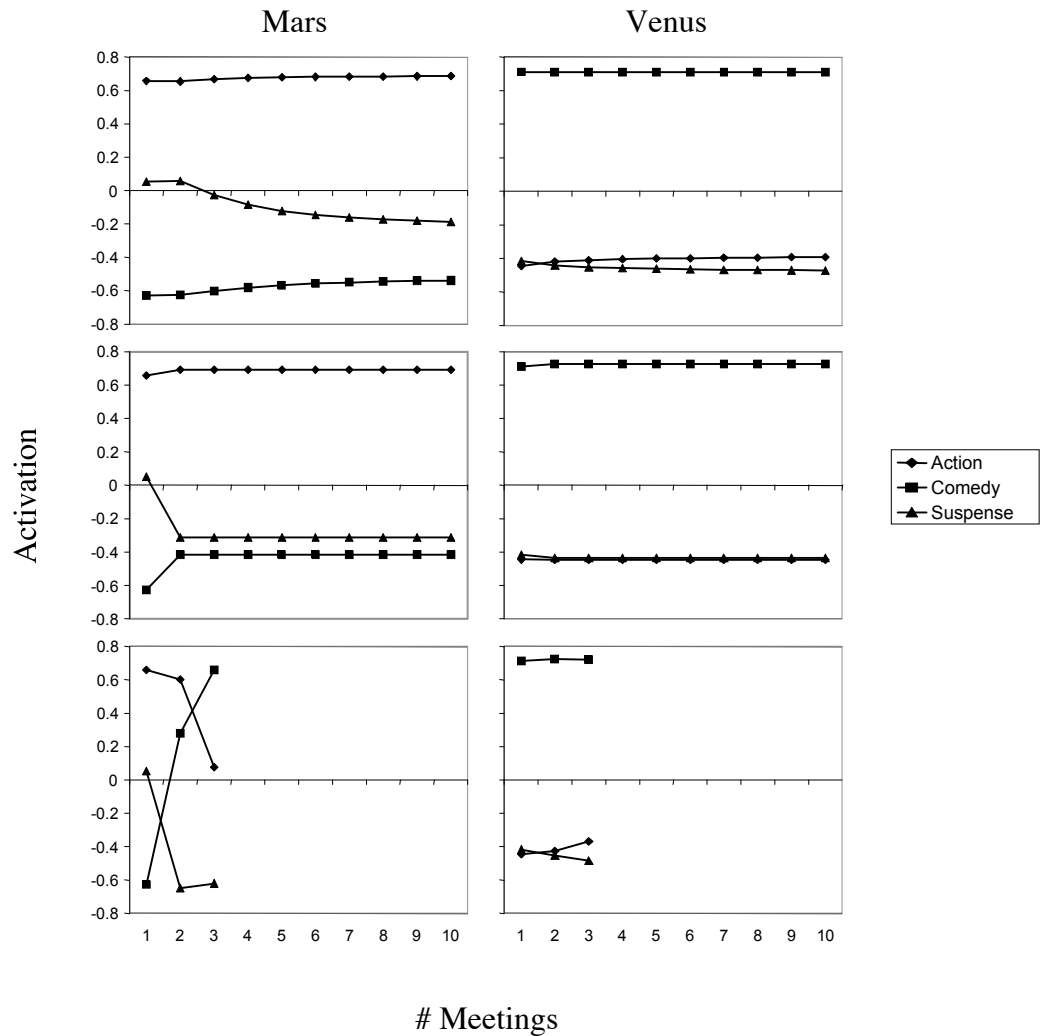
```
(facilitate 'hire-Chris 'research .8)
(facilitate 'hire-Pat 'research .6)
(facilitate 'hire-Pat 'teaching .3)
```

Prof B:

```
(facilitate 'hire-Pat 'research .7)
(facilitate 'hire-Chris 'research .8)
(facilitate 'hire-Chris 'teaching .7)
```

Prof C:

```
(facilitate 'hire-Pat 'research .6)
(facilitate 'hire-Chris 'teaching .7)
(facilitate 'hire-Pat 'administration 1)
```



**Figure 4.** Graphs of the activation of each of the three possible decisions for Mars and Venus. The top row uses contagion as the only emotional communication, the middle only altruism, and the bottom both contagion and altruism. Note that consensus is reached almost immediately in the third case, so the simulation stops after 3 meetings.



While it may not be immediately obvious, Professors A and C initially prefer Pat, and Professor B prefers Chris. Note that the simulation assumes that all three professors have the same three goals, but have an incomplete set of beliefs about how different candidates contribute to them. This makes room for the use of means ends arguments to convince a decision maker of the need to augment their network with additional facilitation relations. For example, Professor C can point out to Professor B that Pat can make a good contribution to administration. The emotional constraint networks are thus not as complete as the one depicted in figure 1.

What happens when the professors meet to discuss the candidates and exchange emotional reactions about them? The interesting (because we did not expect it) answer is that it depends very much on the order of the meetings. HOTCO 3 randomly chooses two members of a group to meet and exchange information by the various social mechanisms. Altruism is not relevant here, since all three professors have the same goals. We had expected that Professors A and C would eventually convince B, and this usually happens.

But surprisingly, sometimes consensus is not reached, if B first meets with A and convinces A by adding new facilitation relations about Chris. Then A switches to Chris, but this does not produce consensus, because C still prefers Pat, and if B meets with C then contagion and means-ends transfers lead B to come to prefer Pat too. The odd result then is that consensus is not reached because both Professors A and B have changed their minds in opposite directions! The next experiment to be done is to see

what happens if A, B, and C have a series of joint meetings in which everyone transmits to everyone else.

Why is it that altruism has such an impact in the Mars-Venus-Neptune example, and yet, though altruism is not active in this simulation, consensus is still achieved? This is a side effect of the fact that all of the agents in this simulation have identical goals initially. As a result, altruism is never used, since no agent can adopt a new goal altruistically. This fact allows the direct use of means-ends communication to fill the gap. To further see the impact of altruism on this example, we implemented a quantitative version of altruism in which agents adopt goals as a kind of weighted sum of the goals of themselves and others. This, as expected, reduced the time needed to reach consensus, but otherwise made little difference.

Additional computational experiments with larger groups show that the number of meetings required to reach consensus increase with the size of the group. Additionally, increasing the number of possible choices also increases the average number of meetings required for consensus. Furthermore, these experiments support the observation that meeting order strongly affects the time needed to reach consensus, as well as whether consensus is reached at all. Some meeting orders produce a form of group polarization. If the first meetings are between agents with similar initial choices, this leads to a strengthening of those agents' feelings. A small group of like-minded agents meeting repeatedly may generate a contagion feedback-loop, with each agent strengthening the other's feelings. If two or more such clusters form, neither has the ability to budge the other's now entrenched position, thus blocking any consensus from forming.

Obviously, both the movie and hiring simulations are highly artificial in their simplifications and choice of inputs. But they still show that even small simulated groups can exhibit interesting decision making behavior as the result of emotional transmission. Future work may implement additional mechanisms such as empathy and apply them to larger groups of more complex individuals.

## **6. Related Work**

In recent years, there has been rapid growth in the computer simulation of organizations as multiagent systems (e.g. Wooldrige, 2002). But little research has been done on the development of consensus in such systems, and we have not been able to find any previous work on emotional consensus.

Johnson and Feinberg (1989) describe computer simulation of consensus as it occurs in crowds. Their model differs from ours in three important ways. First, for the crowd to achieve consensus, not every member had to agree. Rather, consensus was said to be achieved when the amount of variability fell below a certain amount. Second, each agent was highly simplified, varying only in their initial predisposition to a certain plan of action. Furthermore, the agents did not have any reasoning ability whatsoever. Instead, all agents had a fixed likelihood that they would change their decision. Third, the model included information on the spatial distribution of agents. Individual communications were dependent on this distribution. Finally, two forms of communication were modeled. First, there was communication within small subgroups, and then communication between such groups.

Hutchins (1995, ch. 5) discusses communication among members of navigation teams on ships. He describes a computer simulation of a system of members, each of which is a constraint satisfaction system much like those in HOTCO 3, although their function is to interpret events rather than to make decisions. Hutchins considers four parameters that characterize communication: the pattern of interconnections among members of the community; the pattern of interconnectivity of units in the constraint networks; the strength of connections between communicating networks, that is, degree of persuasiveness; and the time course of communication. Hutchins provides a highly illuminating discussion of the advantages and disadvantages of different modes of decision making including consensus and voting, but does not discuss the role of emotion in decision making.

Ephrati and Rosenschein (1996) computationally model an economic decision process that can be used to derive multiagent consensus. Consensus is reached by a process of voting with financial bids: each agent expresses its preferences and a group choice mechanism is used to select the result. The mechanism they advocate involves a tax that charges voters based on the extent to which their bids affect final outcomes. The procedure is clever but complex, and seems more suited to automated systems than to modeling decision making in groups of humans.

Another model (Moss, 2002) attempts to simulate multilateral negotiation. Like our simulation, Moss is interested in multiple agents coming to a consensus. However, the reasoning ability of his agents is limited. Agent goals are represented by two strings, one giving the desired outcomes, and one establishing the importance of each goal to the agent. In addition, each agent has a representation of other agents consisting of values on

a variety of scales, and a ranking of the importance of each of these scales. Both the values and the ranking change with each negotiation. Negotiation between two agents consists of making tradeoffs, in order to achieve consensus. This model works well for situations of bilateral negotiation, but never reaches consensus in cases of multi-lateral negotiation, since reaching an agreement with one agent may cause a new disagreement with another. In contrast to Moss, our model did generally achieve consensus. Like the other research described in this section, Moss ignores emotional aspects of consensus formation.

## **7. Conclusion**

This paper has presented a novel theory of the contribution of emotions to group decision making. We have described social and psychological mechanisms that can encourage the emergence of emotional consensus within a group of decision makers, each of whom makes decisions based on a combination of cognitive and emotional factors. Mechanisms such as emotional contagion, altruism, means-ends argument, analogy, and empathy can transfer emotional attitudes across individuals and lead to the resolution of conflicts. Some of these are modeled computationally by HOTCO 3, a program that shows how group interactions of emotion-driven decision makers can lead to interesting developments, including consensus.

There are undoubtedly additional mechanisms of social transmission of emotions that we have not addressed. On the dark side of group decision making, there are coercive social processes that operate much less benignly than contagion, argument, and empathy. For example, intimidation can lead to group agreement because powerful

members convince others to go along with them out of fear of being harmed. Propaganda uses non-argumentative rhetorical techniques such as appeals to xenophobia to manipulate people's emotional attitudes. We have chosen to concentrate on the more positive contributions of emotions to group decision making, but must acknowledge that emotion can also have less savory contributions. Additional mechanisms for emotional transmission might also include the opposite of contagion, where a history of conflicts leads one member to adopt emotional attitudes opposite to another on the assumption that the other is always wrong. Other possible sources of emotional influence are respect for authority and the desire to conform with group opinions and behavior.

Real-life decision making is much more emotionally complex than the account we have presented here. There are many different emotions that involve a much broader range of feelings than is captured by our discussion of positive and negative valences. On the negative side, anger is different from fear; on the positive side, enthusiasm is different from desperate craving. We have also not discussed complex social emotions, such as those that arise in social situations when people strive to present themselves positively to other group members, to avoid embarrassment, and to maintain group solidarity. For the individual, the relation between actions and goals is often not so narrowly verbal as our facilitation relations suggest, but could reflect visual and other non-verbalized experiences, just as the emotional input and outputs to decision are more a matter of experience than verbal representation.

Despite these limitations, this paper has contributed to the theory of group decision making by highlighting the role of emotions at both the individual and group levels. It has also shown how computational modeling can help to specify the kinds of

psychological and social mechanisms that underlie group conflict and movement toward agreement. Simulations with HOTCO 3 also reveal some surprising results of those mechanisms, even in simple examples.

The focus in this paper has primarily been descriptive, aiming to describe how emotions operate in group decisions. But future work should also address normative issues, assessing how group decision making can be improved by taking into account the ineliminable role of emotions in important choices. We can ask not only the empirical question of which mechanisms of emotional transmission are most effective in bringing a receiver around to a sender's point of view, but also the normative question of which mechanisms are most effective in securing a consensus that best reflects the interests of the whole group. Presumably coercive mechanisms such as threats and propaganda are inferior in this regard to logical and ethical mechanisms such as means-ends argument and altruism. Perhaps future computer simulations will provide a tool for evaluating the desirability of different ways in which transmission of emotional information can lead members of a group toward consensus.

**Acknowledgement:** This research was supported by a grant from the Natural Sciences and Engineering Research Council of Canada. We are grateful to an anonymous referee for very helpful comments.

### References

Barnes, A., & Thagard, P. (1997). Empathy and analogy. *Dialogue: Canadian Philosophical Review*, 36, 705-720.

- Barsade, S. (2002). The ripple effect: Emotional contagion in group behavior. *Administrative Science Quarterly*, 47, 644-675.
- Blanchette, I., & Dunbar, K. (2001). Analogy use in naturalistic settings: The influence of audience, emotion, and goals. *Memory & Cognition*, 29, 730-735.
- Damasio, A. R. (1994). *Descartes' error*. New York: G. P. Putnam's Sons.
- Ephrati, E., & Rosenschein, J. S. (1996). Deriving consensus in multiagent systems. *Artificial Intelligence*, 87, 21-74.
- Finucane, M. L., Alhakami, A. S., Slovic, P., & Johnson, S. M. (2000). The affect heuristic in judgements of risks and benefit. *Behavioral Decision Making*, 13, 1-17.
- Hatfield, E., Cacioppo, J. T., & Rapson, R. L. (1994). *Emotional contagion*. Cambridge: Cambridge University Press.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Johnson, N. R., & Feinberg, W. E. (1989). Crowd structure and process: Theoretical framework and computer simulation model. In E. Lawler & B. Markovsky (Eds.), *Advances in group processes* (Vol. 6, pp. 49-86). Greenwich, CT: JAI Press.
- Kunda, Z., & Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review*, 103, 284-308.



- Lerner, J. S., Small, D. A., & Loewenstein, G. F. (2004). Heart strings and purse strings: Carryover effects of emotions on economic decisions. *Psychological Science, 15*, 337-341.
- Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin, 127*, 267-286.
- Millgram, E., & Thagard, P. (1996). Deliberative coherence. *Synthese, 108*, 63-88.
- Moss, S. (2002). Challenges for agent-based social simulation of multilateral negotiation. In K. Dautenbaum, A. Bond, L. Canamero & B. Edmonds (Eds.), *Socially intelligent agents: Creating relationships with computers and robots*. Norwell, MA: Kluwer.
- Thagard, P. (2000). *Coherence in thought and action*. Cambridge, MA: MIT Press.
- Thagard, P. (2003). Why wasn't O. J. convicted? Emotional coherence in legal inference. *Cognition and Emotion, 17*, 361-383.
- Thagard, P. (forthcoming). The emotional coherence of religion. *Journal of Cognition and Culture*.
- Thagard, P., & Millgram, E. (1995). Inference to the best plan: A coherence theory of decision. In A. Ram & D. B. Leake (Eds.), *Goal-driven learning*: (pp. 439-454). Cambridge, MA: MIT Press.
- Thagard, P., & Shelley, C. P. (2001). Emotional analogies and analogical inference. In D. Gentner, K. H. Holyoak & B. K. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 335-362). Cambridge, MA: MIT Press.

Wagar, B. M., & Thagard, P. (2004). Spiking Phineas Gage: A neurocomputational theory of cognitive-affective integration in decision making. *Psychological Review*, *111*, 67-79.

Wooldridge, M. (2002). *An introduction to multiagent systems*. Chichester: John Wiley & Sons.