Emotions and Decision Making in Trust Games

Francesco Sardo

University of Turin, Italy
fran.sardo@gmail.com

Abstract. This paper presents a computational model that formalizes the reasoning process of a Trust Game player. Experimental results show that standard theories based on the concept of utility fail to describe how humans behave in social games, while on the other hand neuroimaging data confirms the importance of emotions involved in strategic thinking. This model uses a logical structure of emotion elicitation to be able to simulate the influence of emotions in reasoning and gives a formal description that can be executed by a software agent.

Keywords: affective computing, emotions, game theory, neuroeconomics, software agents, trust and reputation.

1 Introduction

The concepts of rationality and rational choice have been deeply investigated in many different fields over the past years. One of the most successful theory that tries to formalize what rationality means belongs to the economic field and it's called the *homo economicus* theory. Many scientists support this theory, especially because of its simplicity and mathematical elegance. This theory states that a rational being will always try to maximize its own perceived utility, where *utility* is a measure of relative satisfaction.

Drawing from this idea, many scientists depicted the perfect rational human like a function that tries to maximize the utility variable. Even the famous psychologist Daniel Gilbert argues that if we apply something like Bernoulli's equation before doing any action we'll always know "how to do exactly the right thing at all possible times" [1]; because following the utility theory "the expected value of any of our actions - that is, the goodness that we can count on getting - is the product of two simple things: the odds that this action will allow us to gain something, and the value of that gain to us [..] if we can estimate and multiply these two things, we will always know precisely how we should behave" [2].

The effectiveness of simple formulas like Bernoulli's influenced researchers to the point of taking the homo economicus concept for true. Thus, when economists evaluated people performance in a series of experiments they found that humans are totally irrational, because they behave very differently from what theory predicts. This is, we have to say, a very strong claim, because researchers themselves defined the concept of rationality a priori, without trying to understand why humans behave the way they do. So, either we take the theory of *homo economicus* for correct and realize that the human being is not a rational being, or we should try to understand more precisely how do we make choices and find a logic behind it, and call this rationality.

In this paper we present a computational framework that gives a description of a software agent playing a Trust Game, an experiment that belongs to the Game Theory. We don't want to analyze a player's behaviour in an utilitarian view: instead, we propose a model inspired by recent findings on how emotions influence our choices. In fact, there is growing evidence through neuroimaging that emotion activation almost predicts player choices in some experiments of the Game Theory. Moreover, researchers started to notice that including emotions in standard models of reasoning helps in understanding how human behave, because emotional processes may have a deeper influence than we ever thought. Some scientists even claim that higher faculties of the human mind are built on low-level processes already present in a primordial brain, like motor skills [3] and emotions regulation [4]. That's why our model will make large use of the emotional component.

2 Trust Game

The Trust Game is a particular experiment that belongs to the Game Theory. It consists of a first player, called trustor, who receives a certain sum of money from the experimenter, and a second player, called trustee. The trustor can give part of his endowment to the trustee, who receives that money multiplied by a constant (usually by 3), then the trustee can give back part of his gain to the trustor.

We chose this experiment as a case study for our emotional framework because there is strong evidence about how emotion affects trust and human affiliation [5] and because experiments with human players firmly contradicts standard economics theories [6]. What the standard theory says is that the trustee has no rational reasons to give some money back to the trustor. So, the optimal strategy chosen by the two players would be to never trust each other, because the other player will rationally defect. This strategy is *Pareto optimal* in the sense that if one player deviates from the strategy and the other sticks with the optimal strategy, the one who deviates will inevitably lose more than he would lose with the optimal strategy.

If we accept this theory, cooperation and trust shouldn't have emerged in a society of rational individuals. Many economists wanted to prevent this result and changed the utility concept to include social welfare among human perceived gains. Adding a social variable into the utility function has the advantage to mathematically restore cooperation as the the rational outcome chosen by the players. However many other economic irrational decisions, like the famous "buying a national lottery ticket", may be perfectly rational if we consider the release of serotonin in the brain caused by the anticipation of possibly winning: buying the ticket provides a good feeling for a small cost, until the drawing indicates you've lost. Utility functions try to design a general law that could tell wether an action is rational or not, while the human brain is a far more complex system and requires that take into account, for instance, internal feedbacks of actions. We believe that is easier to explain decision making in terms of processes rather than through the maximization of an utility function.

This paper tries to highlight the emotional component involved in trust decisions, based on what is known in the psychological literature. Many more steps still must be taken in this direction to better explain how strategic choices are made.

2.1 Emotions and strategic thinking

Experiments with humans playing the Trust Game show that cooperation and trust exist in our society. Even when the game is played one-shot, the trustor most likely entrust a positive sum of money to the trustee, and the trustee usually gives some money in return, even if he has no evident advantages in doing that [6]. Experimenters also noticed that emotions work as an important bias that can influence how much players entrust each other: negative emotions have a negative impact on trust while positive emotions have a positive impact [5].

Negative emotional states observed behaviorally as a result of nonreciprocity have been proposed as a mechanism by wich iniquity is avoided and may have evolved precisely to foster mutual reprocity, to make reputation important, and to encourage punishment of those seeking to take advantage of others [7]. Neuroscientific studies offer the potential to go beyond speculation to examine the correlations between an emotional reaction and subsequent social decision, as well as to investigate whether areas specialized for the processing of basic emotions may be involved for more complex affective reaction. Researchers found brain areas that exibited greater activation as the inequity of the offer increased, and the activation of this area predicted the player's decision to either accept or reject the offert [8].

The evidence that underlines the importance of emotions in strategic choices can be used to give people a better undertanding on how their choices are made and can help them exploit their own affective processes to empower their reasoning. In a related work, authors of [9] illustrate an affective negotiation support system (NSS) that should help negotiators cope with the interplay between affect and the negotiation process. An NSS can help to make someone aware of one's mood and the effects this mood can have on the decision making process. This is an important work because it points out how strategic thinking is not just a matter of cold-blooded decisions but it coexists with emotional activities. Monitoring and exploiting the emotional component is an essential part of the reasoning process, because emotions evaluate the various alternatives and determine the direction of our thought (for example, positive moods favor creative, big-picture integrative thinking, while negative moods favor critical evaluation of details).

For the reasons stated above, it makes sense to propose an alternative scenario to interpret the decisions of Trust Game players. This scenario should make use of the emotional component rather than focusing on the utility theory. This is just a preliminary work and many problems still remain open, but thanks to recent works in computational models of emotions over the last years, we have a handful of formal tools to simulate the role of emotions in strategic decision making.

3 Computational model

In order to build a computational model that incorporates an emotional component we first need to rely on a rigorous formalization of emotional elicitation. Examples of logical formalizations of emotional elicitation are, for instance, Adam et al.'s [10] and Steunebrink's [11]. Both of them are based on the widely accepted OCC model [12] and both were built with software agents programming in mind.

However, Adam et al.'s formalization of OCC's emotion types has been tailored to their BDI-based logical framework, while Steunebrink's is realized trough various stages, where only the last stage commits to BDI. For the purpose of this paper we don't need software agents with full BDI capabilities, thus a more high-level logical description of emotional elicitation is to be prefered. We therefore chose to draw inspiration from the first high-level formalizations described in Steunebrink's work.

Once emotions are triggered throught logical formulas they can be used effectively inside our angent to influence the reasoning process.

3.1 Trustor role

We aim to build a software agent that plays the role of a trustor in a Trust Game with many players, repeated an undefinite number of rounds. Our agent should be able to perform the following tasks: receive the endowment for the current turn, choose a trustee to entrust, select a sum of money to invest, send the money, receive the response, evaluate it and update his model of trust and reputation. To ensure that our agent performs these tasks we divided the reasoning process in four distinct macro-processes, shown in Figure 1.

In the first process, the emotional trigger phase, the agent considers the various investment options to be made for the current round. According to the definitions of [12], pre-conditions for triggerng the prospect based emotions 'hope' and 'fear' are satisfied, as our agent is thinking about a future outcome desirable or undesirable. When thinking about the investment outcome, the agent is emotionally evaluating the investment itself: if the desiderable outcome is thought to be likely (according to the agent's internal model) the investment will be associated with a 'hope' emotion; otherwise the investment will be associated with a 'fear' emotion. To be precise, prospect-based emotions refers to the consequence of an event (in his case, the investment) and it is inaccurate to associate them to the event itself. For a more accurate logical notation, see [11], but for the purpose of this paper we felt that this was an unnecessary complication.

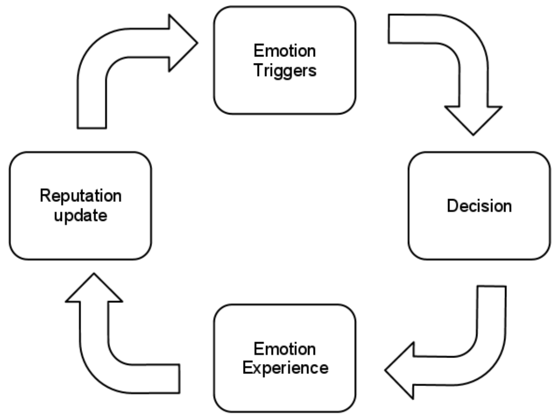

Fig. 1. Trustor reasoning process divided in four main phases.

During the decision phase, emotional information is used to choose the investment itself: valenced feelings help bias a person away from bad outcomes, and toward good ones [13]. As Damasio argued in his researches [14,15], emotions are powerful heuristic used by the human brain to solve problems computationally hard. As already demonstrated in [16], software agents reasoning can benefit from the emotional component to thin the space of options available, thus reducing the combinatorial explosion of some decision problems. Once emotions are cognitively generated, they lead reasoning in different ways depending on their valence. In our case, investment options associated with a positive valenced feeling will be obviously considered first, while investment options associated with a negative valenced feeling will be considered only if no other option is available.

Once the proposal is sent, the emotional experience process evaluates the outcome of the agent's actions. It is important to notice that in this stage we are not talking about emotional triggers, but real emotional experience. Emotional experience is crucial because it allowes for the organism to learn the correlations between perceptions of external world and internal bodily states [17]. As [18] points out, to learn to evaluate external perceptual images and give them meaning, it is necessary to be able to perceive what influence they have on the body. Software agents doesn't have bodies, but experienced emotions can be seen as labels that evaluate one's bodily states. Upon receiving an answer from his parner, the trustor generates automatically evalutative emotions depending on the outcome of the current turn.

On a related topic, is important to point out that emotion experience usually leads to emotion expression [13], and emotion expression itself is one of the key topic in affective computing. We could then think of an hypothetical scenario were emotions experienced by the partecipants of the Trust Game are displayed and perceived by other players, so that they can be used as a useful information inside each reasoning process. In our scenario Trust Game players will only have a partial world observability (meaning that they can have a full history of their own interactions, but they can't see the interactions of other players), but they can use the emotion expression of other player as an information. Thus, emotions generated in response to the actions performed by the trustee are observable by other trustors, and they can in turn generate other emotions accordingly (the so called "fortune-of-others" emotions). This is obviously a scenario far from reality, where deceive is very usual, but it could serve the purpose of better understanding why emotion expression has evolved at all. We think indeed that emotion expression is a powerful tool that provides an interaction feedback to other peers. Other trustors can benefit from the displayed evaluation of an interaction to adjust their own reputation system, that is, how much they entrust each trustee. In this sense, emotion expression can be seen as a public feedback that evaluates the interaction occurred.

Finally, experienced emotions allow the trustor to change his own trust and reputation system. Emotional values should help the trustor to learn which parner to trust for the next interactions, while avoiding untrustful partners. Machine learning and fuzzy logic approaches should be preferred in this stage, but we chose to stick to a simple logical framework. The reason for this choiche is that we lack enough information to tailor a learning system linked to the emotional experience. A good starting point would be Picard's "negative resistances" model [13], but our intention here is to sketch a simple logical framework. In the next section, we will discuss in details a computational model that describes the reasoning process of the trustor.

3.2 A formalization

1. Emotional Triggers phase The OCC model defines 'hope' and 'fear' as "(pleased) about the prospect of a desiderable event" and "(displeased) about the prospect of a undesiderable event". We thus have to define what a desiderable (or undesiderable) event is. Following data gathered by [6], players who play the role of a trustor give positive feedback when they receive two-thirds of the gain of the trustee, while giving a negative feedback for the loss of money. We can consider X as the hypothetical event of receiving a certain sum of money by the trustee: X is considered a desiderable event if the sum received is greater than the two-third of the gain of the trustee, while X is considered a undesiderable event if the sum is lower than the money invested. In the first case we say that the trustee collaborates, while in the other he defects. During the emotion trigger phase our agent will consider, for each trustee, an investment hypothesis, activating either 'hope' or 'fear' emotions, as follows:

$$Pleased_i^T(X) \land Prospect_i(X) \land Likely_i(X) \to Hope_i^T(X)$$
$$Displeased_i^T(X) \land Prospect_i(X) \land Likely_i(X) \to Fear_i^T(X)$$

Where $Hope_i^T(X)$ means emotion 'hope' is triggered for agent i with respect of the prospected event X; same thing applies for 'fear' emotion. So, when a prospected event X is considered desiderable (the amount of money received is greater than the two-thirds of the gain of the trustee) and the event is likely (its likelihood is greater than a certain threshold *alpha*), emotion 'hope' is activated with respect of that event. The estimated probability that a parner will cooperate is derived from the trust and reputation system maintained by the trustor. Note that the trustor does not feel a real fear (or hope) until the actual investment has not been made: the emotions at this level are just conditional on his endowing the trustee with money. Nevertheless, emotional triggers have a precise value which can be used as an information to discard some options during the decision making process.

2. Decision phase Emotional information activated in the previous stage is used in the decision process to select the candidate for the current investiment. Investment options labeled with a positive valenced emotion ('hope') will be considered first, as follows:

```
for each hypothesis that triggered hope':
        select one hypothesis based on the weighted
           probability of success;
        exit;
for each hypothesis that triggered fear' with a
   likelihood lower than a threshold beta:
        select one hypothesis based on the weighted
           probability of insuccess;
        exit;
do nothing this turn;
exit;
```

In other words, an agent first considers the options labeled with 'hope', and choose an investment among them weighting the probability of its success. A greedy strategy will only consider the investment with the highest probability of success, but we wanted to give our agent the ability to select some exploratory choices. A random selection biased by the weighted probability of success is the best way we can think of to prevent a local optimum strategy and to simulate the innate propensity to exploratory searches.

Only if no option is labeled with 'hope', the agent starts considering dangerous options labeled with 'fear'. In fact, we should remind that emotions do not determine uniquely the behavior, but instead provide an action tendency that other reasoning mechanisms may decide to ignore. A famous example is that of a rat overcoming its fear to get a piece of cheese guarded by a cat. Our agent has the target of maximizing his own profit at the end of the game, so he might consider a risky investment. In this case, undesiderable options labeled with 'fear' that have a probability of success lower than a threshold *beta* (that is, the investment has a relative high probability of success) will be considered in this stage. If no option meets this requirement, the trustor will skip the investment stage this turn, and he will only generate fortune-of-others emotions in the next stage.

On a side note, is important to note that although in our model emotions depend on valence and likelihood, and decisions depend on emotions, is not correct to state that emotions are redundant and to think that we can completely to do away with emotions. Not considering emotions would mean building an ad-hoc solution tailored to the given problem. Instead, integrating the emotional component inside computational models is a useful step forward in the simulation of human mind behaviour. For instance, when data about bodily states' influence on emotion will be collected in detail, it would be possible to simulate the interconnected feedback between cognitive generated emotions and bodily generated emotions. If we have decision models that make use of the emotional component, the integration should be seamless.

We emphasize that both *alpha* and *beta* thresholds can vary accordingly to the agent current mood. For example, a streak of positive outcomes can lower the *alpha* threshold while raising the *beta* threshold. This can well simulate the influence that mood plays on decision making: according to [19] a person in a good mood is more cooperative than a person in a bad mood. In our framework a good mood will broaden the options available, increasing the chances of cooperation, while a bad mood will narrow the number of selected candidates.

Once a partner has been selected, a proposal is sent. The amount of money sent is proportional to how much the other player is trusted [6], so this sum is basically determinated by the agent reputation model.

3. Emotional Experience phase In this phase event Y represent the actual answer given by the trustee that can be evaluated through 'joy' and 'distress' emotions, as follows:

$$Pleased_i^T(Y) \wedge Actual_i(Y) \rightarrow Joy_i^T(Y)$$
$$Displeased_i^T(Y) \wedge Actual_i(Y) \rightarrow Distress_i^T(Y)$$

where $Joy_i^T(Y)$ means emotion 'joy' is triggered for agent i with respect of the actual event Y. The triggering conditions for the derived emotions are then met, as follows:

$$Joy_i^T(Y) \wedge Past\,Hope_i^T(X) \wedge Confirms_i(Y,X) \rightarrow Satisfaction_i^E(Y)$$
$$Joy_i^T(Y) \wedge Past\,Fear_i^T(X) \wedge Disconfirms_i(Y,X) \rightarrow Relief_i^E(Y)$$
$$Distress_i^T(Y) \wedge Past\,Hope_i^T(X) \wedge Disconfirms_i(Y,X) \rightarrow Disapp._i^E(Y)$$
$$Distress_i^T(Y) \wedge Past\,Fear_i^T(X) \wedge Confirms_i(Y,X) \rightarrow Fears\text{-}confirmed_i^E(Y)$$

Following Steunebrink's original logical structure, derived emotions produced as logical consequences of the previous formulas should have been simple triggered emotions. But as discussed previously, it is important that we consider these emotions as experienced emotions, instead of simple triggers. In this paper we omit the complexity of how to model the emotional experience, focusing on its effect. As described in [20], computational models of emotional experience are at an early development, but they can bring many benefits to the study of emotional mechanisms. Here we don't want to give details on how to model emotion experience. Instead, we simply note that in this stage emotion experience is a key requirement to evaluate the partner answer: simple emotional triggers don't activate the complex bodily and cognitive reactions in response to a given situation.

To have players with emotion experience means also being able to reason about emotional expressions and considering the use other players can make of this information. If we place ourselves in an ideal setting in which emotional expression is visible to other participants without deceit and distortions, we find a possible explanation for the emergence of emotional expressions as an interaction feedback. A player can then activate fortune-of-others emotions, as follows:

$$Satisfaction_j^E(Y) \vee Relief_j^E(Y) \rightarrow Presume_i\, Des_j(Y)$$
$$Fears\text{-}confirmed_j^E(Y) \vee Disapp._j^E(Y) \rightarrow Presume_i\, Undes_j(Y)$$

$$isSelfish_i \wedge Presume_i\, Undes_j(Y) \rightarrow Gloating_i^E(Y,j)$$
$$isSelfish_i \wedge Presume_i\, Des_j(Y) \rightarrow Resentment_i^E(Y,j)$$
$$\neg isSelfish_i \wedge Presume_i\, Undes_j(Y) \rightarrow Pity_i^E(Y,j)$$
$$\neg isSelfish_i \wedge Presume_i\, Des_j(Y) \rightarrow Happy\text{-}for_i^E(Y,j)$$

In these formulas we have introduced a distinction between two types of agent: a selfish and an unselfish agent. A selfish agent will be focused only on his personal gains, while the other will care more about other trustor gains and losses and will take other players' emotional feedbacks in serious consideration when updating his own trust and reputation system. The prediction is that the lower is the number of selfish agents the greater will be the total gain realized at the end of the game by the trustors.

The emotions arised during this phase will be used inside the next phase to change the trust values of his possible partners.

4. Reputation update phase In this phase, generated emotions are used as a sort of learning mechanism that allows our agent to change the reputation of each trustee involved. In fact, behavioural evidence suggest that the ability to have sensations is strongly connected to basic mechanisms of learning and decision making [21,22]. These studies directly relate emotion to reinforcement learning. Another artificial intelligence work [23] relates artificial affect with information processing, showing a computational explanation compatible with the psychological literature on affect and learning.

Our approach is purely logical, so that the reputation update phase is relized through the following formulas:

$$Happy\text{-}for_i^E(Y,j) \vee Gloating_i^E(Y,j) \rightarrow Admiration_i^T(Y.Source, Y)$$
$$Resentment_i^E(Y,j) \vee Pity_i^E(Y,j) \rightarrow Reproach_i^T(Y.Source, Y)$$

$$Satisfaction_i^E(Y,X) \vee Relief_i^E(Y,X) \rightarrow Admiration_i^T(Y.Source, Y)$$
$$Disapp._i^E(Y,X) \vee Fears\text{-}confiermed_i^E(Y,j) \rightarrow Reproach_i^T(Y.Source, Y)$$

$$Admiration_i^E(Y.Source, Y) \rightarrow Incr\, Reputation_i(Y.Source)$$
$$Reproach_i^E(Y.Source, Y) \rightarrow Decr\, Reputation_i(Y.Source)$$

We can imagine our reputation model like a set of crisp reputation values like <*very low, low, medium, high, very high*> altough a fuzzy logic approach would have been preferable. Unfortunately, we don't have enough information to give emotions an intensity variable useful to a machine learning or fuzzy logic approach, so that the *increase_reputation* and *decrease_reputation* predicates are basically left or right shift in the reputation values set.

In the forumulas above emotions produced by derived emotions and fortune-of-others emotions are written as emotional triggers. This is because there is not a direct relation between, for instance, the experience of a 'satisfaction' emotion and the experience of an 'admiration' emotion directed to the agent that caused the event that brought satisfaction. Instead, the experience of a 'satisfaction' emotion simply cause the trigger for 'admiration' to be activated: that is, the necessary but not sufficient conditions for 'admiration' to be experienced are met. In fact, a trustor can activate simultaneously both 'admiration' trigger (for being satisfied of the amount of money received) and 'reproach' trigger (for blaming trustee's behaviour towards another trustor) toward the same trustee. In this case, what will be the resulting experienced emotion? Our answer is that an agent should always give priority to the emotional experience relative to his own interactions, but when this is not possibile (that is, when he is evaluating the actions of a trustee he has not interacted with), he should rely on a majority vote; if there is still a tie, the reputation of the trustee will not change. Once is know which is the experienced emotion relative to a certain trustee, his reputation is changed.

4 Conclusions and related works

This project aims at describing computationally the strategical decision making within the Trust Game experiment. The model presented is in accordance with the latest psychological theories on decision making, particularly the Somatic markers hypotesis [14]. This theory points out how modern economic theories ignores the influence of emotions on decision-making, while neuroscience evidence suggests that sound and rational decisions, in fact, depend on prior accurate emotional processing [25].

The dominant approach to the strategic thinking problem, until now, was similar to the physical explanation of reality, that is looking for a small set of general laws from which everything can be explained. This approach will not work for the human mind, because the brain is a layered system built in the course of evolution through small incremental additions [26], thus resulting in a very complicated architecture that can not be explained by a small set of laws [4].

The emotional component, in particular, plays a fundamental role that cannot be ignored if we want to build models that aims to simulate and understand the rationality behind human choices. This paper presents a software agent scheme of reasoning with an emotional component. A related work [16] uses a logical emotion framework to help a software agent cope with the non-determinancy of certain choices. We, in turn, emphasize social interactions placing our agent in a multiagent context, where he can activate emotions regarding the outcome of the choices of other players (fortune-of-others emotions).

Other works about emotional mechanism may also suggest interesting ideas to the trust and reputation area of study [27]. In this paper we underline the importance of emotional expression as an interaction feedback between people. Since most of the facial expressions are culturally independent and are recognized by most people [28], emotions provide useful information to know how situations are evaluated by other people to act accordingly. A better understanding of how emotion mechanism affects trust and human affiliation is needed (for instance, cfr. [29] for a distinction between cognitive and affective dimensions on trust).

This computational framework has the advantage of being easily implementable inside a software agent. This way it is possibile to test the performance of the software agent against normal human players. Our aim is not to build an agent that achives the best overall performance within the game, but to build a belivable Trust Game player. To this end, it is possible to simulate a Trust Game run where each player is separate from the others and has a monitor as his only interface. Each turn the players can select an amount of money to send to a chosen trustee, and they can in turn rate the interaction via four standard feebacks ('satisfaction', 'disappointment', 'relief' and 'fears-confirmed'), so that the software agent can retrieve the emotional reaction of other players. If the framework is well built, partecipants won't be able to tell who is the non-human player (as a sort of *Turing test*).

Finally, another interesting experiment is a game played only by non-human trustors. Thanks to data collected by [6] we know in detail how human players behave in Trust Game experiments. While keeping constant the original experiment parameters, it is possible to obtain statistics about the performance of software agents and compare them to the original data. Playing only with non-human trustors is also a good way to test the difference between selfish agents and unselfish agents. Our expectation is that, on the short run, selfish trustors may have a personal advantage respect to other trustors [30] in terms of money gained from the experiment, but we expect the total sum of money gained by the trustors at the end of the game to be lower if there is a high percentage of selfish agents.

References

1. Gilbert, D. (2006). Stumbling on Happiness. *Knopf.*

2. Gilbert, D. (2005). Our Mistaken Expectations. *TEDGlobal 2005.*
3. Lakoff, G.; Núñez, R. E. (2000). Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being. *Basic Books.*
4. Minsky, M. L. (2007). The Emotion Machine. *Simon and Schuster.*
5. Myers, D.; Tingley, D. (2011). The Influence of Emotion on Trust. *PhD thesis.*
6. Keser, C. (2007). Game-theory based trust experiment relating reputation management with trust and trustworthiness.
7. Nowak, M. A.; Page K. M.; Sigmund, K. (2000). Fairness versus reason in the ultimatum game. *Science 289, 1773.*
8. Sanfrey, A. G.; Rilling J. K.; Aronson J. A.; Nystrom L. E.; Cohen J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science 300, 1755.*
9. Brookens, J.; Jonker, C. M.; Meyer, J. C. (2010). Affective Negotiation Support Systems. *Journal of Ambient Intelligence and Smart Environments, 2(2),* 121-144.
10. Adam, C.; Herzig A.; Longin D. (2009). A Logical Formalization of the Occ Theory of Emotions. *Synthese 168 (2):201 - 248.*
11. Steunebrink, B. R. (2010). The Logical Structure of Emotions. *PhD thesis.*
12. Ortony, A.; Clore, G. L.; Collins A. (1988). The Cognitive Structure of Emotions. *Cambridge University Press.*
13. Picard, R. W. (1997). Affective Computing. *MIT Press.*
14. Damasio, A. R. (1994). Descartes' Error: Emotion, Reason and the Human Brain. *Grosset/Putnam.*
15. Bechara, A.; Damasio, H.; Damasio, A. R. (2000). Emotion, Decision Making and the Orbitofrontal Cortex. *Oxford Univ Press.*
16. Steunebrink, B. R.; Dastani, M.; Meyer, J. C. (2010). Emotions to control agent deliberation. *AAMAS 2010.*
17. LeDoux, J. E. (1992). Brain mechanisms of emotion and emotional learning. *Current opinion in neurobiology.*
18. Velik, R. (2010). Minds and Machines. *Springer.*
19. Hertel, G.; Neuhof, J.; Theuer, T.; Kerr N. L. (2000). Mood effects on cooperation in small groups: Does positive mood simply lead to more cooperation? *Taylor & Francis.*
20. Broekens, J. (2010). Modeling the Experience of Emotion. I*nternational Journal on Synthetic Emotions, 1(1), 1-17.*
21. Berridge, K. C. (2003). Pleasures of the brain. *Brain and Cognition, 52, 106–128.*
22. Cohen, J. D.; Blum, K. I. (2002). Reward and decision. *Neuron, 36, 193–198.*
23. Broekens, J.; Kosters, W. A.; Verbeek, F. J. (2007). Affect, Anticipation and Adaptation: Affect-Controlled Selection of Anticipatory Simulation in Artificial Adaptive Agents. *Adaptive Behavior, 15(4), 397-422.*
24. Damasio, A. R. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex [and discussion]. *Philosophical Transactions: Biological Sciences.*
25. Bechara, A.; Damasio, A. R. (2005). The Somatic Marker Hypothesis: A Neural Theory of Economic Decision. *Games and Economic Behavior.*
26. Bara, B. G. (1999). Pragmatica cognitiva: i processi mentali della comunicazione. *Bollati Boringhieri.*
27. Dunn, J. R.; Schweitzer, M. E. (2005). Feeling and believing: the influence of emotion on trust. *Journal of Personality and Social Psychology, 88.*
28. Ekman, P. (1982). Emotion in the Human Face.
29. Johnson, D.; Grayso, K. (2003). Cognitive and affective trust in service relationships. *Journal of business research.*

30. Sen, S.; Biswas, A.; Debnath, S. (2000). Believing others: pro and cons. *MultiAgent Systems Proceedings.*