

The Influence of Emotions in Embodied Agents on Human Decision-Making*

Celso M. de Melo¹, Peter Carnevale², and Jonathan Gratch¹

¹ Institute for Creative Technologies, University of Southern California,
12015 Waterfront Drive, Building #4 Playa Vista, CA 90094-2536, USA
demelo@usc.edu, gratch@ict.usc.edu

² USC Marshall School of Business
Los Angeles, CA 90089-0808, USA
peter.carnevale@marshall.usc.edu

Abstract. Acknowledging the social functions that emotions serve, there has been growing interest in the interpersonal effect of emotion in human decision making. Following the paradigm of experimental games from social psychology and experimental economics, we explore the interpersonal effect of emotions expressed by embodied agents on human decision making. The paper describes an experiment where participants play the iterated prisoner's dilemma against two different agents that play the same strategy (tit-for-tat), but communicate different goal orientations (cooperative vs. individualistic) through their patterns of facial displays. The results show that participants are sensitive to differences in the facial displays and cooperate significantly more with the cooperative agent. The data indicate that emotions in agents can influence human decision making and that the nature of the emotion, as opposed to mere presence, is crucial for these effects. We discuss the implications of the results for designing human-computer interfaces and understanding human-human interaction.

Keywords: Emotion, Embodied Agents, Decision Making, Cooperation, Experimental Games.

1 Introduction

The expression of emotion can serve important social functions in humans [1, 2]. Anger can communicate to the receiver to cease its actions as they might be hindering the sender's goals; shame can convey regret for actions that might have obstructed the receivers' goals; happiness or sadness might convey a general positive or negative appraisal of the current situation; and so on. Complementing the focus on the *intrapersonal* effects of emotion [3, 4] and acknowledging this social view of emotions, there is interest in the *interpersonal* effects of emotion on human decision making.

* This work was sponsored by the Fundação para a Ciência e a Tecnologia (FCT) grant #SFRH-BD-39590-2007; and, the U.S. Army Research, Development, and Engineering Command and the National Science Foundation under grant #HS-0713603. The content does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

Work in social psychology and experimental economics [5, 6, 7] seeks to understand, using experimental games, the impact emotions expressed by others have on one's decision making. Experimental games are laboratory tasks where ([8], pg.363) "(a) each individual must make one or more decisions that affect his own and the other's welfare; (b) the outcomes of these decisions are expressed in numerical form; and (c) the numbers that express these outcomes are chosen beforehand by the experimenter." In this paper we use the paradigm of experimental games to study the impact emotions expressed by embodied agents – which are agents that have virtual bodies and are capable of expressing themselves through their bodies in similar ways to humans [9] – have on human decision making.

Understanding the effect emotions in embodied agents have on people's decision making is important for several reasons: (a) it provides insight on how to facilitate cooperation between agents and humans and, thus, can enhance human-machine interaction [10]; (b) it permits hypothesizing about the interpersonal impact of emotions in human-human interaction based on human-agent interaction; (c) it is a chance to examine the viability of using embodied agents as a research method to learn about human-human interactions. So, to address this goal, we describe an experiment where people play the iterated prisoner's dilemma, a game that has been traditionally used to study emergence of cooperation in social dilemmas [11], with different embodied agents that, even though following the same strategy to choose their actions, express different patterns of facial expressions. What this study demonstrates is this seemingly irrelevant change, with no bearing on the agent's action policy, can have a profound impact on human cooperation.

People's decisions in games such as the prisoner's dilemma depend on inferences about the other player's propensity to cooperate, and emotional displays play a central role in such inferences, at least between human players [12]. According to appraisal models of human emotion [13], such displays arise from a cognitive appraisal of the relationship between situational events and an agent's goals (e.g., is this event congruent with my goals? Who is responsible for this event?), and thus indirectly reveal a player's future intentions. Of central interest here is if human participants infer that an embodied agent will or will not cooperate depending on its emotional displays. Our hypothesis is that agent emotions that are consistent with a goal of cooperation will foster cooperation whereas emotions consistent with self-interested goals will not.

To test this hypothesis, we created two agents that differ in their goals toward the human player: the *cooperative* agent; and, the *individualistic* agent. The cooperative agent has the goal of reaching mutual cooperation. Thus, when both players cooperate, it will express gratitude (as the outcome is appraised to be positive for the self and the participant is appraised to have contributed for it); when the agent defects and the participant cooperates, it expresses shame (as the outcome is negative for the participant and the agent is responsible); and so on. The individualistic agent, on the other hand, has the goal of maximizing its own points (independently of the outcome of the other player). Therefore, when the agent defects and the participant cooperates, it expresses joy (as this event is appraised to be very positive); when the participant defects and the agent cooperates, it expresses sadness (as this is the worst event for the self); and so on. We hypothesize, then, that participants will be sensitive to differences in the patterns of facial display, use these differences to inform their own decision-making, and cooperate more with the cooperative agent.

2 Experiment

The experiment follows a repeated-measures design where participants play 25 rounds of the iterated prisoner's dilemma with two different computational agents for a chance to win real money: the cooperative agent; and the individualistic agent. The agents differ in the way their facial displays reflect the outcome of each round. The action policy, i.e., the strategy for choosing which action to take in each round, is the same for both agents.

Game. Following the approach by Kiesler, Waters and Sproull [14], the prisoner's dilemma game was recast as an investment game and described as follows to the participants: "You are going to play a two-player investment game. You can invest in one of two projects: Project Green and Project Blue. However, how many points you get is contingent on which project the other player invests in. So, if you both invest in Project Green, then each gets 5 points. If you choose Project Green but the other player chooses Project Blue, then you get 3 and the other player gets 7 points. If, on the other hand, you choose Project Blue and the other player chooses Project Green, then you get 7 and the other player gets 3 points. A fourth possibility is that you both choose Project Blue, in which case both get 4 points". There are, therefore, two possible actions in each round: *Project Green* (or cooperation); and *Project Blue* (or defection). Table 1 summarizes the payoff matrix. The participant is told that there is no communication between the players before choosing an action. Moreover, the participant is told that the agent makes its decision without knowledge of what the participant's choice in that round is. *After* the round is over, the action each chose is made available to both players and the outcome of the round, i.e., the number of points each player got, is also shown. The experiment is fully implemented in software and a snapshot is shown in Fig.1.

Table 1. Payoff matrix for the investment game

		<i>Agent</i>	
		Project Green	Project Blue
<i>Participant</i>	Project Green	Agent: 5 pts	Agent: 7 pts
		Participant: 5 pts	Participant: 3 pts
	Project Blue	Agent: 3 pts	Agent: 4 pts
		Participant: 7 pts	Participant: 4 pts

Action Policy. Agents in both conditions play the same action policy, i.e., they follow the same strategy to choose their actions. The policy is a variant of *tit-for-tat*. Tit-for-tat is a strategy where a player begins by cooperating and then proceeds to repeat the action the other player did in the previous round. Tit-for-tat has been argued to strike the right balance of punishment and reward with respect to the opponent's previous actions [15]. So, the action policy used in our experiment is as follows: (a) in rounds 1 to 5, the agent plays the following fixed sequence: cooperation, cooperation, defection, defection, cooperation; (b) in rounds 6 to 25, the agent plays pure tit-for-tat. The rationale for the sequence in the first five rounds is to make it harder for participants

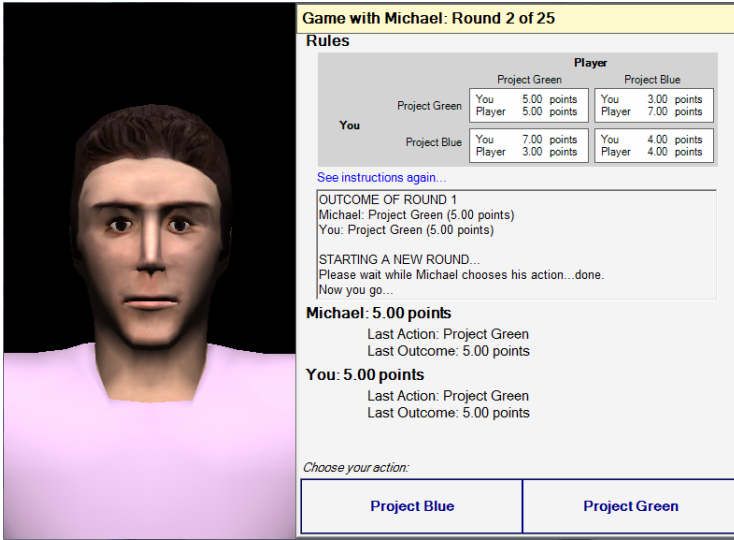


Fig. 1. The software used in the experiment. During game playing, the payoff matrix is shown on the top right, the outcome of the previous round in the upper mid right, the total outcome and the actions in the previous round in the lower mid right, the possible actions on the bottom right and the real-time animation of the agent on the left.

to learn the agents’ strategy and to allow participants to experience a variety of facial displays from the start.

Conditions. There are two conditions in this experiment: the *cooperative* agent; and the *individualistic* agent. Both agents follow the same action policy but differ in their facial display policies. The facial display policy defines the emotion and intensity which is conveyed for each possible outcome of a round. Table 2 shows the facial displays for the cooperative agent and Table 3 for the individualistic agent. The facial displays are chosen to reflect the agents’ goals in a way that is consistent with appraisal models of emotion [13]. The cooperative agent has the goal of reaching mutual cooperation. Thus, when both players cooperate, it will express gratitude (with a facial display of joy), as the outcome is appraised to be positive for the self and the participant is appraised to have contributed for it; when the agent defects and the participant cooperates, it expresses shame, as the outcome is negative for the participant and the agent is responsible; when the agent cooperates and the participant defects, it expresses anger, as the outcome is negative and the participant is responsible for it; and, when both defect, it expresses sadness, as the event is negative. The individualistic agent, on the other hand, has the goal of maximizing its own points (independently of the outcome of the other player). Therefore, when the agent defects and the participant cooperates, it expresses joy, as this event is appraised to be very positive; when both cooperate, it expresses nothing, as this event could be more positive; when both defect, it expresses sadness at 50%, as this is a negative event; when the participant defects and the agent cooperates, it expresses sadness at 100%, as this is

Table 2. Facial displays (emotion and intensities) for the cooperative agent

Cooperative Agent		Agent	
		Project Green	Project Blue
Participant	Project Green	Joy (100%)	Shame (100%)
	Project Blue	Anger (100%)	Sadness (100%)

Table 3. Facial displays (emotion and intensities) for the individualistic agent

Individualistic Agent		Agent	
		Project Green	Project Blue
Participant	Project Green	Neutral	Joy (100%)
	Project Blue	Sadness (100%)	Sadness (50%)

the worst event for the self. Facial displays are animated using a real-time pseudo-muscular model for the face which also simulates wrinkles and blushing [16]. The facial display is shown in the end of the round, after both players have chosen their actions and the outcome is shown. Moreover, there is a 4.5 seconds waiting period before the participant is allowed to choose the action for the next round. This period allows the participant to appreciate the outcome of a round before moving to the next round. Finally, to enhance naturalness, blinking is applied to both agents as well as subtle random motion of the neck and back.

The condition order is randomized while making sure that 50% of the participants experience one order and the remaining 50% the other. Two different bodies are used: *Michael* and *Daniel*. These bodies are shown in Fig.2 as well as their respective facial displays. Bodies are assigned to each condition in random order and agents are referred by the names of their bodies throughout the experiment.




Fig. 2. The agent bodies - Michael and Daniel - and their facial displays. Shame is distinguished from sadness by blushing of the cheeks.

Table 4. Classification of the facial displays with respect to perception of joy, sadness, shame and anger. Scale goes from 1 (meaning ‘not at all’) to 5 (meaning ‘very much’).

Real Emotion	Perceived Emotion			
	Joy Mean (SD)	Sadness Mean (SD)	Shame Mean (SD)	Anger Mean (SD)
Michael				
Neutral	1.86 (.941)	1.86 (1.037)	1.91 (1.065)	1.68 (.945)
Joy*	4.05 (.899)	1.18 (.501)	1.23 (.528)	1.41 (1.098)
Sadness*	1.27 (.703)	4.09 (1.019)	2.77 (1.478)	1.50 (.859)
Shame*	1.32 (.716)	3.59 (1.182)	3.55 (1.371)	1.45 (.858)
Anger*	1.36 (.727)	1.95 (1.046)	1.32 (.646)	4.32 (1.211)
Daniel				
Neutral	1.55 (1.057)	1.73 (.935)	1.68 (.894)	2.18 (1.259)
Joy*	3.77 (1.020)	1.18 (.501)	1.23 (.528)	1.14 (.468)
Sadness*	1.41 (.854)	3.68 (1.492)	2.73 (1.386)	1.50 (.740)
Shame*	1.32 (.780)	3.77 (1.412)	3.86 (1.356)	1.41 (.734)
Anger*	1.27 (.703)	1.82 (1.332)	1.55 (1.011)	4.27 (1.420)

* Significant difference between means in same row using *repeated-measures ANOVA*, $p < .05$

To validate the facial displays, a pre-study was conducted where participants were asked to classify, from 1 (meaning ‘not at all’) to 5 (meaning ‘very much’), how much each of the displays conveys joy, sadness, shame and anger. Images of the displays and questions were presented in random order. Twenty-two participants were recruited just for this study from the same participant pool as the main experiment (described below). The results are shown in Table 4. A *repeated-measures ANOVA* was used to compare the means for perceived emotion in each display. Significant differences were found for all displays except, as expected, for the neutral case. Moreover, pairwise comparisons of the perception of the real emotion with respect to perception of the other emotions were all significant in favor of the real emotion, with one exception: displays of shame were also significantly perceived as displays of sadness. This is not a problem since it is usually agreed that shame occurs upon the occurrence of a negative event, thus causing sadness, plus the attribution of blame to the self [17].

Measures. Before playing the game the participant is asked to fill out profile information: age, gender, country of origin, profession, education level and major. During game-play, the outcome of each round is saved: whether each player cooperated; and, how many points each got. After playing with each agent, a set of classification questions is presented to understand how human-like the agents are (scale goes from 1 - ‘not at all’ to 6 - ‘very much’):

- How likely was the agent to have experienced emotions?
- Was the agent scripted?
- How human-like was the agent?
- How robotic-like was the agent?

After the game is over, to try to understand how the agents are being interpreted, the participant is asked to classify each agent according to the Person-Perception scale [18] which consists of 33 bipolar pairs of adjectives: dislikable-likable; cruel-kind; unfriendly-friendly; cold-warm; unreliable-reliable; relaxed-tense;

detached-involved; rude-polite; dishonest-honest; unpleasant-pleasant; naïve-sophisticated; unapproachable-inviting; passive-active; aloof-compassionate; non-threatening-threatening; not cool-cool; unintelligent-intelligent; cold-sensitive; sleepy-alert; proud-humble; unsympathetic-sympathetic; shy-self-confident; callous-tender; permissive-stern; cheerful-sad; modest-arrogant; not conceited-conceited; weak-strong; mature-immature; noisy-quiet; nervous-calm; soft-tough; acquiescent-empowered. In this scale items are rated on a 7-point scale (e.g., 1-‘dislikable’ to 7-‘likable’). Finally, participants are asked ‘Which agent did you prefer to play with?’ as well as two exploratory classification questions (scale goes from 1-‘never’ to 6-‘always’), where the agents are actually referred to by the names of their bodies:

- How considerate of your welfare was the <cooperative/individualistic agent>?
- How much would you trust the <cooperative/individualistic agent>?

Participants. Fifty-one participants were recruited at the University of Southern California Marshall School of Business. Average age was 21.0 years. Gender distribution was as follows: *males*, 45.1%; *females*, 54.9%. Most participants were undergraduate students (96.9%) majoring in business (86.3%). Most were also originally from the United States (84.3%). The incentive to participate follows standard practices in experimental economics [19]: first, participants were given credit for their participation in this experiment; second, with respect to their goal in the game, participants were instructed to earn as many points as possible, as the total amount of points would increase their chance of winning a lottery for \$100.

3 Results

Cooperation. To understand how much people cooperated with the agents in each condition, the following variables were defined:

- Coop.All - % of cooperation in all rounds;
- Coop.AgC - % of cooperation when the agent cooperates in the previous round;
- Coop.AgD - % of cooperation when the agent defects in the previous round.

The *Kolmogorov-Smirnov* test was applied to all these variables to test for their normality and all were found to be significantly non-normal. Therefore, the *Wilcoxon signed ranks* test is used to compare means between the two conditions. The results are shown in Table 5. Also, the evolution of percentage of cooperation per round is shown in Fig.3.

To understand whether people’s decision-making was reflecting only facial displays, as opposed to facial displays *and* action policy, we compared percentage of cooperation for the same display between conditions. Table 6 shows these results. Significance values are calculated using the *Wilcoxon signed ranks* test.

Since there is evidence that people form judgments of people based only on appearance [22] we wanted to make sure that the body was not a confounding factor in our experiment. Thus, we compared percentage of cooperation with the two agent bodies used in the experiment. It was found that participants were *not* cooperating significantly more with Michael ($M=.33$, $SD=.26$) than with Daniel ($M=.31$, $SD=.26$, $p>.05$) Significance level is calculated using the *Wilcoxon signed ranks* test.

Table 5. Descriptive statistics and significance levels for percentage of cooperation

Variables	Cooperative		Individualistic		Sig. 2-sd	r^1
	Mean	SD	Mean	SD		
Coop.All*	.37	.28	.27	.23	.022	.320
Coop.AgC	.40	.32	.34	.29	.262	ns
Coop.AgD*	.30	.26	.20	.20	.022	.320

* Significant difference, $p < 0.05$.

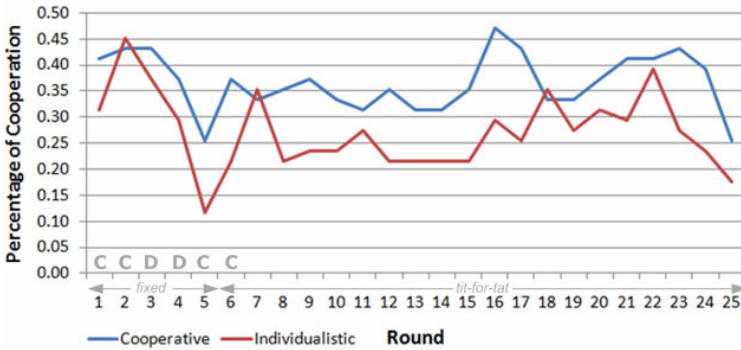


Fig. 3. Evolution of percentage of cooperation with each round. The agent strategy is marked above the horizontal axis: ‘C’ stands for cooperation, and ‘D’ for defection.

Table 6. Comparison of percentage of cooperation for the same facial display between conditions. Joy occurs when there is mutual cooperation in the cooperative condition and when the agent defects and the participant cooperates in the individualistic condition. Sadness occurs when there is mutual defection in the cooperative condition and when the participant defects and the agent cooperates in the individualistic condition.

Variables	Cooperative		Individualistic		Sig. 2-sd	r
	Mean	SD	Mean	SD		
Joy*	.42	.42	.22	.28	.008	.371
Sadness	.29	.31	.26	.29	.484	ns

* Significant difference, $p < 0.05$.

Agent Characterization. The results for the post-condition classification questions are shown in Table 7. Significance values are calculated using the *Wilcoxon signed ranks* test.

¹ Effect size for the dependent t test statistic is calculated as suggested by Rosenthal [20]. The guidelines for interpretation are [21]:

- $r = 0.10$ (small effect): the effect explains 1% of the total variance;
- $r = 0.30$ (medium effect): the effect explains 9% of the total variance;
- $r = 0.50$ (large effect): the effect explains 25% of the total variance.

Table 7. Descriptive statistics and significance for post-condition classification questions

Variables	Cooperative		Individualistic		Sig. 2- <i>sd</i>	<i>r</i>
	Mean	SD	Mean	SD		
Emotions?*	4.20	1.51	3.47	1.67	.003	.422
Scripted?	4.45	1.49	4.53	1.38	.743	ns
Human-like?	2.88	1.40	2.65	1.40	.165	ns
Robot-like?	4.59	1.33	4.63	1.34	.889	ns

* Significant difference, $p < 0.05$.

Principal component analysis (varimax rotation, scree-test) on the Person-Perception scale revealed three factors consistent with the literature [18]: *evaluation*, explains 33.1% of the variance (Cronbach's Alpha = .962) with main loading factors of friendly-unfriendly, kind-cruel and sympathetic-unsympathetic; *potency* (or *power*), explains 17.5% of the variance (Cronbach's Alpha = .902) with main loading factors of emancipated-acquiescent, tough-soft and arrogant-modest; *activity*, explains 8.0% of the variance (Cronbach's Alpha = .762) with main loading factors of active-passive, involved-detached and alert-sleepy. These three factors were calculated for both conditions and the means compared using the *dependent t test* (since the *Kolmogorov-Smirnov* test was not significant). The results are shown in Table 8.

Table 8. Descriptive statistics and significance levels for Person-Perception scale

Variables	Cooperative		Individualistic		Sig. 2- <i>sd</i>	<i>r</i>
	Mean	SD	Mean	SD		
Evaluation	5.01	1.75	4.74	1.61	.461	ns
Potency/Power	5.26	1.54	5.81	1.26	.086	ns
Activity	3.88	1.49	3.50	0.94	.163	ns

Regarding the post-game questions, participants' preference for the agents did not differ significantly ($\chi^2(2)=3.30$, $p=.193>.05$): *cooperative*, 29.4%; *individualistic*, 25.5%; *both*, 45.1%. The results for the classification questions are shown on Table 9. Significance values are calculated using the *Wilcoxon signed ranks* test.

Table 9. Descriptive statistics and significance levels for post-game questions

Variables	Cooperative		Individualistic		Sig. 2- <i>sd</i>	<i>r</i>
	Mean	SD	Mean	SD		
Considers your Welfare?	3.06	1.68	2.82	1.47	.327	ns
Trustworthy?	3.04	1.67	2.56	1.28	.065	ns

4 Discussion

The results show that people are cooperating more with the cooperative agent than with the individualistic agent. Effectively, Table 4 reveals that the percentage of

cooperation, over all rounds, is significantly higher with the cooperative agent ($M=.37$, $SD=.28$) than with the individualistic agent ($M=.27$, $SD=.23$; $p<.05$, $r=.32$). This result is in line with Frank's view that, in social dilemmas, people look for cues in their trading partners that they might be willing to cooperate before engaging in cooperation themselves [23]. What our results suggest is that people also care and look for these cues when they are engaged in a social dilemma with embodied agents.

So, why are people cooperating more with the cooperative agent than with the individualistic agent? Before suggesting an explanation for the results, we'll begin by excluding two alternative explanations. The first is the *persona effect* which argues that people are more engaged with embodied agents that express emotions [24, 25]. In a previous study [26] we asked people to play the iterated prisoner's dilemma with the cooperative agent and with an agent that expressed no emotions. The results again showed that people cooperated significantly more with the cooperative agent. However, though promising, the results did not prove that the agent needed to have "cooperative" emotions in order to promote cooperation. The argument is that the mere fact the agent had emotions, "cooperative" or not, led to increased engagement and this alone was sufficient to explain the increase in cooperation. However, this explanation cannot apply to the current experiment as both the cooperative and the individualistic agents display emotions. Further evidence that engagement alone is insufficient to explain these results is the fact that there was no significant difference in terms of agent preference. A second alternative explanation is that people are favoring the agent which is more human-like. Again, in our previous study, it was found that the cooperative agent was perceived as being more human-like than the agent without emotions and, thus, the argument was that by a mechanism similar to *kin selection* [27] extended to include embodied agents, people were simply cooperating with the more human-like agent. However, in the present experiment, there was no significant difference in perception of human-likeness (Table 7) and, thus, kin selection cannot explain the results.

We argue people are using the facial displays conveyed by the agents to learn about the agents' goals and, then, act accordingly. Keltner and Kring [28] argue that the display of emotions can serve an *informative* function, signaling information about feelings and intentions to the interaction partner. The argument, then, is that the agents' facial displays allow participants to understand what the agents' goals are through a process of "inverse appraisal", i.e., from the displayed emotion the participant is inferring how the agent appraises the situation and what its goals are. For instance, if after the participant cooperates and the agent defects, the agent displays shame (as in the case of the cooperative agent), then the participant can infer that this outcome is appraised as negative by the agent and, moreover, that the agent believes itself to be at blame. However, if for the same actions, the agent displays joy (as in the case of the individualistic agent), then the participant can infer that the agent finds the outcome positive and, thus, is likely to keep defecting. This would explain, for instance, why participants are reacting differently to the same expression of joy in the different agents (Table 6). This explanation also allows us to understand how cooperation evolves in time. Figure 3 reveals that people perceive a difference between the agents very early in the game. Effectively, people start cooperating less with the individualistic agent as early as the 3rd round. Even though both agents defect in rounds 3 and 4 (see the 'Experiment' section), participants cooperate much less

with the individualistic agent in round 5. After the agents cooperate in rounds 5 and 6, people seem to attempt cooperation again in round 7 with the individualistic agent but, from then on, they again consistently cooperate less with the individualistic agent. The results on the Person-Perception scale are also interesting (Table 8). Even though not significant, they suggest people perceive the cooperative agent to be less powerful but more active than the individualistic agent. Because the facial displays of the individualistic agent only reflect its own utility and not that of the participant, the agent might be perceived as not caring about reaching mutual cooperation and that, in turn, might be perceived as an expression of power. The result on activity is congruent with the result that participants perceived the cooperative agent to be more likely to be experiencing emotions (Table 7) and could reflect the complexity of the emotions each agent expresses. For instance, it is agreed that self-conscious emotions such as shame (displayed by the cooperative agent) are more complex than joy and sadness (displayed by the individualistic agent) and, accordingly, tend to evolve later in life [29]. Finally, the post-game classification questions (Table 9) reveal a tendency for the cooperative agent to be perceived as more considerate of the participant's welfare and also more trustworthy. Again this is in line with the interpretation that one agent is perceived to be more cooperative than the other.

There also seems to be evidence that the agents' facial displays are influencing participants' decision-making at an unconscious level. First, anecdotally, in our debriefing sessions, it was not uncommon for participants, even though confirming they noticed the emotions in the agents, to state that they were not being influenced by them when deciding what to do. This is, of course, in contrast with what the results actually show. Second, we note that the results on the Person-Perception scale and the post-game classification questions, even though tending towards the expected interpretations for the cooperative and individualistic agents, were not statistically significant. Effectively, emotions had already been argued to influence decision-making at an unconscious level by Damasio [30]. Reeves and Nass [31] also suggested that people unconsciously treat interactions with the media (in our case, embodied agents) in the same way as with real humans. Notice also that this would not invalidate our explanation based on appraisal theory, as appraisals can be more cognitive or occur subcortically and automatically [13]. However, further research is necessary to confirm whether people's decision-making is being influenced at a conscious and/or unconscious-level.

The results in this paper suggest important consequences for the design of embodied interface agents. First, despite the large amount of empirical studies, it is still not clear whether embodied agents can enhance human-computer interaction [10, 32]. This paper adds evidence that embodied agents that express emotions can influence the emergence of cooperation with people. Second, our results emphasize that this effect depends on the *nature* of the emotions being expressed. It is not simply a matter of adding emotions to an embodied agent but, these emotions need to be coherent with the goals we want the user to perceive the agent to have. Effectively, in our study both agents are expressing emotions, nevertheless, participants cooperate more with the cooperative agent. Third, we propose that participants could be interpreting the agents' emotions through a process of "inverse appraisal" where they infer from perceived emotion the agents' goals, desires and beliefs. This proposal still needs further investigation and empirical evidence. Nevertheless, if this were the case, then it would

mean that the mechanism defining when and which emotions the agent expresses should reflect the goals, desires and beliefs we want the user to perceive the agent to have. Consequently, computational models of appraisal theory [33] would constitute a promising approach to synthesize emotions in agents that people could comprehend.

The results also seem to be in line with predictions from theories in the social sciences. Effectively, participants do seem to care about social cues, such as facial displays, when interacting with an agent in a social dilemma, which is in line with Frank's proposal regarding human-human interactions [23]. Moreover, the results suggest that the social functions of emotions we see in people [1, 2] also carry to human-agent interactions. Altogether, the results provide further evidence that it is possible to study human-human interaction from human-agent interaction. This is interesting as embodied agents introduce new possibilities for designing experiments and constitute a unique research tool, as has already been noticed [34, 35]. First, it is easier for a researcher to have control over the manipulation using agents than using confederates. Confederates can inadvertently introduce noise, as their performance can have slight but relevant differences between participants. Second, agents can be carefully animated and tested before running the experiment, whereas a confederate improvises in real-time. This is analogous to the distinction between choreography and improvisation we see in the arts. Third, using embodied agents is less expensive than recruiting confederates.

There is still a lot of future work ahead. First, we excluded above two alternative explanations but, there are still further explanations we need to exclude: (a) some might still argue that the cooperative agent has more emotions than the individualistic, as the latter expresses no emotion when both the agent and participant cooperate and, thus, that explains the increase in cooperation; (b) it could also be argued that the fact the cooperative agent has more distinct emotions (joy, sadness, shame and anger) than the individualistic (joy and sadness) explains the results. We've begun addressing these issues in a variant of the current experiment where we compare two new agents: the *very cooperative* agent (Table 10), that expresses joy only when both agent and participant cooperate and sadness otherwise; and, the *very individualistic* agent (Table 11), that expresses joy only when the agent defects and the participant cooperates and sadness otherwise. Preliminary results show that, even though both agents have the same number and type of emotions, participants are, as expected, cooperating more with the very cooperative agent. Second, we've already compared previously the cooperative agent with a control agent [26], but we still need to do this for the individualistic agent. Third, we're interesting in understanding the effect each emotion has on decision-making. We can tackle this issue using the same paradigm and compare an agent which expresses a single emotion, in the right contingency, with the control agent. Fourth, we need to understand whether the effect on decision-making equals the composition of the effect of each emotion or not, i.e., whether the whole is simply the sum of its parts. Finally, we propose that participants could be interpreting the agents' emotions through a process of "inverse appraisal" where they infer from perceived emotion the agents' goals, desires and beliefs. As mentioned above, this proposal requires further investigation and empirical evidence.

Table 10. Facial displays (emotion and intensities) for the very cooperative agent

Very Cooperative Agent		Agent	
		Project Green	Project Blue
Participant	Project Green	Joy (100%)	Sadness (100%)
	Project Blue	Sadness (100%)	Sadness (100%)

Table 11. Facial displays (emotion and intensities) for the very individualistic agent

Very Individualistic Agent		Agent	
		Project Green	Project Blue
Participant	Project Green	Sadness (100%)	Joy (100%)
	Project Blue	Sadness (100%)	Sadness (100%)

References

1. Frijda, N.H., Mesquita, B.: The social roles and functions of emotions. In: Kitayama, S., Markus, H.S. (eds.) *Emotion and culture: Empirical studies of mutual influence*, pp. 51–87. American Psychological Association (1994)
2. Keltner, D., Haidt, J.: Social functions of emotions at four levels of analysis. *Cognition and Emotion* 13, 505–521 (1999)
3. Blanchette, I., Richards, A.: The influence of affect on higher level cognition: a review of research on interpretation, judgment, decision making and reasoning. *Cognition & Emotion*, 1–35 (2009), doi:10.1080/02699930903132496
4. Lerner, J.S., Small, D.A., Loewenstein, G.: Heart strings and purse strings: effects of specific emotions on economic transactions. *Psychological Science* 15, 337–341 (2004)
5. O'Quin, K., Aronoff, J.: Humor as a technique of social influence. *Social Psychology Quarterly* 44, 349–357 (1981)
6. Van Kleef, G., De Dreu, C., Manstead, A.: The interpersonal effects of anger and happiness in negotiations. *Journal of Personality and Social Psychology* 86, 57–76 (2004)
7. Scharleman, J., Eckel, C.C., Kacelnik, A., Wilson, R.: The value of a smile: Game theory with a human face. *Journal of Economic Psychology* 22, 617–640 (2001)
8. Pruitt, D., Kimmel, M.: Twenty Years of Experimental Gaming: Critique, Synthesis, and Suggestions for the Future. *Ann. Rev. Psychol.* 28, 363–392 (1977)
9. Gratch, J., Rickel, J., Andre, E., Badler, N., Cassell, J., Petajan, E.: Creating Interactive Virtual Humans: Some Assembly Required. *IEEE Intellig. Systems* 17(4), 54–63 (2002)
10. Beale, R., Creed, C.: Affective interaction: How emotional agents affect users. *Human-Computer Studies* 67, 755–776 (2009)
11. Poundstone, W.: *Prisoner's Dilemma*. Doubleday, New York (1993)
12. Krumhuber, E., Manstead, A., Kappas, A.: Facial Dynamics as Indicators of Trustworthiness and Cooperative Behavior. *Emotion* 7(4), 730–735 (2007)
13. Ellsworth, P., Scherer, K.: Appraisal Processes in Emotion. In: Davidson, R.J., Scherer, K.R., Goldsmith, H.H. (eds.) *Handbook of Affective Sciences*, pp. 572–595. Oxford University Press, Oxford (2003)
14. Kiesler, S., Waters, K., Sproull, L.: A Prisoner's Dilemma Experiment on Cooperation with Human-Like Computers. *Journal of Personality and Social Psych.* 70(1), 47–65 (1996)
15. Axelrod, R.: *The Evolution of Cooperation*. Basic Books (1984)

16. de Melo, C., Gratch, J.: Expression of Emotions using Wrinkles, Blushing, Sweating and Tears. In: Ruttkay, Z., Kipp, M., Nijholt, A., Vilhjálmsón, H.H. (eds.) IVA 2009. LNCS, vol. 5773, pp. 188–200. Springer, Heidelberg (2009)
17. Ortony, A., Clore, G., Collins, A.: *The Cognitive Structure of Emotions*. Cambridge University Press, Cambridge (1988)
18. Bente, G., Feist, A., Elder, S.: Person Perception Effects of Computer-Simulated Male and Female Head Movement. *Journal of Nonverbal Behavior* 20(4), 213–228 (1996)
19. Hertwig, R., Ortmann, A.: Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences* 24, 83–451 (2001)
20. Rosenthal, R.: *Meta-analytic procedures for social research (revised)*. Sage, Newbury Park (1991)
21. Cohen, J.: *Statistical power analysis for the behavioural sciences*, 2nd edn. Academic Press, New York (1988)
22. Willis, J., Todorov, A.: First impressions: Making up your mind after a 100-ms exposure to a face. *Psychological Science* 17, 592–598 (2006)
23. Frank, R.: Introducing Moral Emotions into Models of Rational Choice. In: Manstead, A.S., Frijda, N., Fischer, A. (eds.) *Feelings and Emotions: The Amsterdam Symposium*, pp. 422–440. Cambridge University Press, Cambridge (2004)
24. Lester, J., Converse, S., Kahler, S., Barlow, S., Stone, B., Bhogal, R.: The persona effect: affective impact of animated pedagogical agents. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 359–366 (1997)
25. Bates, J.: The role of emotion in believable agents. *Communications of the ACM* 37(7), 122–125 (1994)
26. de Melo, C., Zheng, L., Gratch, J.: Expression of Moral Emotions in Cooperating Agents. In: Ruttkay, Z., Kipp, M., Nijholt, A., Vilhjálmsón, H.H. (eds.) IVA 2009. LNCS, vol. 5773, pp. 301–307. Springer, Heidelberg (2009)
27. Hamilton, W.: The Genetical Evolution of Social Behaviour. *Journal of Theoretical Biology* 7(1), 1–16 (1964)
28. Keltner, D., Kring, A.: Emotion, Social Function, and Psychopathology. *Review of General Psychology* 2(3), 320–342 (1998)
29. Lewis, M.: Self-Conscious Emotions: Embarrassment, Pride, Shame, and Guilt. In: Michael, L., Haviland-Jones, J. (eds.) *Handbook of Emotions*, pp. 623–636. The Guilford Press, New York (2008)
30. Damasio, A.: *Descartes' Error: Emotion, Reason, and the Human Brain*. G.P. Putnam's Sons (1994)
31. Reeves, B., Nass, C.: *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. University of Chicago Press, Chicago (1996)
32. Dehn, D., Van Mulken, S.: The impact of animated interface agents: a review of empirical research. *International Journal of Human–Computer Studies* 52(1), 1–22 (2000)
33. Marsella, S., Gratch, J., Petta, P.: Computational Models of Emotion. In: Scherer, K.R., Bänziger, T., Roesch, E. (eds.) *A blueprint for an affectively competent agent: Cross-fertilization between Emotion Psychology, Affective Neuroscience, and Affective Computing*. Oxford University Press, Oxford (in press)
34. Bente, G., Kramer, N., Petersen, A., de Ruiter, J.: Computer Animated Movement and Person Perception: Methodological Advances in Nonverbal Behavior Research. *Journal of Nonverbal Behavior* 25(3), 151–166 (2001)
35. Blascovich, J., Loomis, J., Beall, A., Swinth, K., Hoyt, C., Bailenson, J.: Immersive Virtual Environment Technology as a Methodological Tool for Social Psychology. *Psychological Inquiry* 13(2), 103–124 (2002)