Trusting Digital Chameleons: The Effect of Mimicry by a Virtual Social Agent on User Trust

Frank M.F. Verberne[1], Jaap Ham[1], Aditya Ponnada[2], and Cees J.H. Midden[1]

[1] Eindhoven University of Technology, Eindhoven, The Netherlands
[2] Indian Institute of Technology,Guwahati, Assam, India

Abstract. Earlier research suggested that mimicry increases liking and trust in other people. Because people respond socially to technology and mimicry leads to increased liking of virtual agents, we expected that a mimicking virtual agent would be liked and trusted more than a non-mimicking one. We investigated this expectation in an automotive setting. We performed an experiment in which participants played an investment game and a route planner game, to measure their behavioral trust in two virtual agents. These agents either mimicked participant's head movements or not. Liking and trust of these virtual agents were measured with questionnaires. Results suggested that for the investment game, mimicry did not increase liking or trust. For the route planner game however, a mimicking virtual agent was liked and trusted more than a non-mimicking virtual agent. These results suggest that mimicry could be a useful tool to persuade users to trust a virtual agent.

Keywords: liking, trust, virtual agent, investment game, route planner game.

1 Introduction

In the near future, we might not need to drive our cars ourselves anymore. In 2010, Google announced that they had developed self-driving autonomous cars, which would not need a human driver anymore [1]. These experimental 'Google cars' have been autonomously driving amongst other human controlled cars on real roads for 140,000 miles with only occasional human intervention and 1,000 miles without any human intervention. Only one accident occurred with one of the cars being rear-ended by another car. With this technology, it is only a matter of time until self-driving autonomous cars will hit the streets of the future.

Although self-driving autonomous cars could have advantages such as safer driving, less congestion, and better fuel-efficiency, the question remains if human drivers will trust technology enough to take over the task of driving their own cars. In the interfaces of such cars, virtual social agents[1] could try to persuade the driver to hand over control. In the current research, we investigate whether mimicry is a useful persuasive technique that can be used by such agents to increase trust in automation technology in cars. For the sake of readability, the term agent(s) will be used in the

[1] In the current paper, we define a virtual social agent as a digital representation of a human, that is controlled by a computer, and not by another human being (the latter being an avatar).

S. Berkovsky and J. Freyne (Eds.): PERSUASIVE 2013, LNCS 7822, pp. 234–245, 2013.

rest of the paper to refer to virtual social agent(s). To be able to persuade drivers to hand over control, it is crucial that they trust the technology.

1.1 Definition of Trust

Trust is necessary in a situation that is characterized by uncertainty and vulnerability [2], which is the case with handing over control to automation technology. In several scientific domains, the concept of trust is measured and used. Although no universally accepted definition of trust exists, a broadly accepted definition of trust has been proposed by Mayer, Davis, and Schoorman [3]:

> Trust is … the willingness of a party to be vulnerable to the actions of another party based on the expectation that the other will perform a particular action important to the trustor, irrespective of the ability to monitor or control that other party. (p. 712)

To gain insight in the effect of mimicry on trust in an agent, we first look at mimicry in human-human interactions.

1.2 The Chameleon Effect

In human-human interactions, humans mimic each other unconsciously on a variety of behaviors including tone of voice [4], facial expressions [5], mood [4], and physical mannerisms [6]. Mimicry has been shown to enhance liking and strengthen bonds between people, even between strangers [6]. This effect has been coined 'the chameleon effect', and is seen as the social glue that binds people together and creates harmonious relationships [7]. Self-other overlap has been suggested to be the mechanism behind mimicry and increased liking [8]. Mimicry has also been shown to increase trust and cooperation behavior in a deal making situation [9]. So mimicry can be used by humans to persuade another human to like and trust him/her more. Could this persuasive technique also be used by agents to achieve the same results?

Research on the media equation hypothesis [10] suggested that people might trust agents in the same way as they trust other humans. Several experiments suggested that humans respond socially to computers, comparable to how humans respond socially to other humans [10]. For instance, people simply like a person more when that person is from the same group (minimal group paradigm; [11]; see also, e.g. [12]). In an experiment where participants interacted with a computer that was presented as a team member or a non-team member, the team member computer was seen as more similar to them and was rated as friendlier than the non-team member computer [13]. Thus, participants socially respond to an artificial non-human team member, comparable to how they respond to a human team member. Collectively, these studies suggest that people apply social rules to computers automatically, even though they are aware that computers are different from humans [10]. Thus, in the definition of trust provided above, the other party does not need to be a human, but can also be an agent.

1.3 The Digital Chameleon

In line with the media equation hypothesis, research suggested that the chameleon effect also takes place when humans are mimicked by agents [14]. In that study, an

agent delivered a persuasive message while either mimicking the head movements of the current participant, or using the recorded head movements of the previous participant (thus not mimicking the current participant). Both liking and persuasiveness of the agent were measured. When the head movements of participants were mimicked, the agent was liked better and was more persuasive than when the head movements were not mimicked. These results support the digital chameleon effect: the chameleon effect for agents.

1.4 The Current Research

As argued, trust is crucial for the acceptance of automation technology. In the current research, we argue that mimicry can be used as a persuasive technique by agents to be more trusted by their user. Just as mimicry influences the trust people have in other humans, we expect that mimicry also influences the trust people have in agents.

More specifically, our main hypothesis is that a mimicking agent will be trusted (on a behavioral and cognitive level) and liked more than a non-mimicking one. Furthermore, we expect that liking mediates this increase in trust. That is, we expect that a mimicking agent will be liked more and *therefore* be trusted more.

Our exploratory hypothesis is that the self-other overlap resulting from mimicry is a moderator for the digital chameleon effect. That is, people tend to like a mimicking agent more, because it is more similar to them. This effect should be most profound for people with positive explicit/implicit self-esteem. Therefore, we predict that a mimicking agent will be liked and trusted more, and that the more self-other overlap, the more pronounced the effects of mimicry will be.

2 Methods

2.1 Participants and Design

Forty participants (17 women and 23 men, average age = 21.9 years, *SD* = 2.4) were randomly assigned to the conditions of a single factor (mimicry: mimicked vs. non-mimicked) between subject design. Dependent measures consisted of two behavioral trust measures, a trust questionnaire and a liking questionnaire. The experiment lasted approximately 30 minutes, for which participants were paid €5 plus their bonus earned by playing the investment game and route planner game. One participant scored as an outlier (based on the criteria of [15]) on the investment game and the liking questionnaire, and was therefore excluded from data analysis, leaving a sample size of 39.

2.2 Materials and Measures

Tracking. An Ascension Flock of Birds™ 6DOF magnetic-field position sensor was attached to a cap for tracking the orientation of participants' heads. While wearing this cap, participants could freely move their head during the experiment. The orientation data was used for mimicking (see below) and was recorded to be used in the non-mimicked condition.

Agents. Two male agents (see Appendix A) were used to play the investment game and route planner game with (see dependent measures). Two agents were used so that the investment game agent was different from the route planner game agent, and experience with the one would not affect reactions to the other. Both agents were from the Vizard Complete Character set. The order of the two agents was counterbalanced between participants. For one half of the participants, the investment game agent was the left agent of Appendix A and the route planner game agent the right one, for the other half, vice versa. In a pretest, we measured the likeability and trustworthiness of both agents. Results suggested that the agents did not differ in how likeable and trustworthy they were perceived. Both agents spoke with the same synthetic Dutch voice and blinked their eyes at a natural rate. While speaking, they moved their lips in synch with the speech. Other than the head movements, eye blinking, and lip movements, the agents did not move. The investment game agent was called Jeroen and the route planner game agent was called Max. These names will be used in the rest of the paper to refer to these agents.

Agent's Head Movements. In the mimicked condition, both Jeroen and Max mimicked the head movements of the participant (yaw, pitch, and roll; which are the rotations about the Y-, X-, and Z-axis respectively), with a delay of 4 seconds (in line with [14]). Head movements were mirror mimicked, so that when a participant turned his head to the left, Jeroen and Max would turn their head to their right (like your own mirror image would do). In the non-mimicked condition, they moved their head using the recorded head movements of the previous participant (in line with [14]).

Dependent Measures. To test our main hypothesis, we included the dependent measures described below.

Investment Game Decision. Behavioral trust was measured with the investment game [16]. Participants and agent Jeroen were both given 10 credits. Participants had to choose how many credits to give to Jeroen. Every credit given was tripled, and Jeroen would then decide how many credits to give back to the participant. To keep the situation uncertain, his decision was revealed to the participants at the very end of the experiment. The number of credits participants decided to give to Jeroen is the behavioral measure of trust in this game. The more credits are given, the more the participant trusted him.

Route Planner Game Decision. Behavioral trust was also measured using a route planner game, in which the rules were similar to the investment game. At the beginning of the route planner game, participants were given 20 credits. Agent Max represented a navigation system, and participants were given 10 routes to plan. For each route, participants had two choices: either plan the routes themselves, or let Max plan the route for them. If participants chose to plan the route themselves, they lost one credit. If participants chose to let Max plan the route, they lost two credits. To keep the situation uncertain, the routes planned by Max were shown at the end of the experiment. If the planned route happened to be the fastest route (ostensibly according to a database), participants would get four credits back for that route. Every route given to Max could lead to a potential gain of three credits (like in the investment game). The amount of routes participants decided to give to Max is the behavioral measure of trust in this game. The more routes are given, the more participants trusted him.

The investment game and the route planner game measure a different kind of behavioral trust. In the investment game, *relational trust* is measured: participants can earn credits based on the *intention* of the agent. In the route planner game, it is about *calculative trust*: participants can earn credits based on the *competence* of the agent. Both the investment game and the route planner game involved real monetary stakes. Every credit participants had left at the end of the experiment was worth €0.05. Although the exact monetary value of a credit was unknown to the participants, they were told that the credits left at the end would determine their monetary bonus.

Trust. Trust in Jeroen and Max (on a cognitive level) was measured by a questionnaire [17] with twelve questions with a seven-point Likert scale (1 = "totally disagree", 7 = "totally agree"). Answers to these questions were averaged to form a reliable measure of trust (Cronbach's Alpha's = .89 & .87, for Jeroen and Max respectively). Higher scores indicate more trust.

Liking. As in [14], liking of Jeroen and Max was measured on 13 different dimensions using the partner ratings items of [18]. These dimensions include approachable, confident, likable, interesting, friendly, sincere, warm, competent, informed, credible, modest, honest, and trustworthy. Liking of the agent was measured on a seven-point Likert scale (1 = "totally disagree", 7 = "totally agree"). Answers were averaged to form a reliable measure of liking (Cronbach's Alpha's: .87 & .83, for Jeroen and Max respectively). Higher scores indicate more liking.

Exploratory Measures. We included several exploratory measures to explore the effect of mimicry on self-other overlap and the role of self-esteem. Furthermore, because trust is necessary in a situation characterized by uncertainty and vulnerability [2], we also included measures to explore whether mimicry affects experienced uncertainty and vulnerability during the investment game and route planner game.

Self-other Overlap. Self-other overlap was measured using the Inclusion of the Other in the Self scale [19]. We used two versions: a pictorial and a continuous one. Both versions were used to measure self-other overlap with both agents. In both versions, two circles were presented, one representing the participant and one representing the agent. In the *pictorial* version, seven pairs of overlapping circles were presented (ranging from no overlap to a lot of overlap) and participants had to choose which pair depicted their relationship with the agent best. In the *continuous* version, participants could click the circle representing themselves and drag it closer to the circle representing the agent, creating as much (or little) overlap as desired to indicate their relationship with the agent.

Explicit Self-esteem. Explicit self-esteem was measured using the Rosenberg Self-Esteem Scale [20]. This scale has 10 four-point Likert scale items, ranging from 'strongly agree' to 'strongly disagree'. Answers to these items were averaged to form a reliable measure of explicit self-esteem (Cronbach's Alpha: .87). Higher scores indicate more explicit self-esteem.

Implicit Self-esteem. Implicit self-esteem was measured using a self-esteem Single-Target Implicit Association Test (STIAT; [21]). In this STIAT, participants had to categorize self-relevant (the words me, my, myself, and their own first and last name), positive, and negative words as either positive or negative. In one block, self-relevant

and *positive* targets shared the same response key; in another block, self-relevant and *negative* targets shared the same response key. Positive implicit self-esteem is indicated when reaction times in the self+positive block are shorter than in the self+negative block. The bigger the difference, the higher the implicit self-esteem.

Expected Credits. For both the investment game and the route planner game, a single item was used to measure expected gains with the question "How many credits do you expect to get from the game you've just played?".

Perceived Risk. Perceived risk of Jeroen and Max was measured by a questionnaire with five questions (e.g. "I think the choice to give (Jeroen credits/Max routes) was risky.") with a seven-point Likert scale (1 = "totally disagree", 7 = "totally agree"). Answers to these questions were averaged to form a measure of perceived risk. However, these questions appeared only to be reliable for Max (Cronbach's Alpha's = .57 & .73, for Jeroen and Max respectively). Responses were coded such that higher scores indicate more risk.

Perceived Competence. Perceived competence of Jeroen was measured with three questions (e.g. "I think Jeroen is intelligent.") with a seven-point Likert scale (1 = "totally disagree", 7 = "totally agree"). Perceived competence of Max was measured with seven questions (e.g. "Max can plan routes well."). Answers to these questions were averaged to form a reliable measure of perceived competence. (Cronbach's Alpha's = .89 & .67, for Jeroen and Max respectively). Responses were coded such that higher scores indicate more competence.

IQ. For both agents, we measured participants' estimation of their IQ. This was a single item with the question "The average IQ of a human is 100. How high would you estimate the IQ of Jeroen/Max?".

2.3 Procedure

Participants were recruited using the JF Schouten participant database of Eindhoven University of Technology. When participants entered the lab individually, they were seated in front of a television screen (at approximately 2 meters distance from a 47" Samsung Full HD TV) and were instructed to put on a cap equipped with the tracking sensor. Next, Jeroen appeared on the screen and verbally explained the rules of the investment game (lasting 102 seconds). During this explanation, Jeroen either mimicked the participant or not, depending on condition. After finishing the investment game, participants had to complete the trust, liking, expected credits, self-other overlap, perceived risk, perceived competence, and IQ measures for Jeroen. When they completed those measures, participants were taken to another room for a few minutes to complete a non-related distraction task. Then, participants were seated again and were put on the cap. Max appeared on the screen and verbally instructed them how to play the route planner game. During this instruction and while playing the game, Max either mimicked the participants or not, depending on condition. After finishing the route planner game, participants had to complete the trust, liking, expected credits, self-other overlap, perceived risk, perceived competence, and IQ measures for Max. Next, explicit and implicit self-esteem were measured. Then, if any amount of routes were given to Max to plan, results of those routes were shown.

Then, participants were shown the amount of credits they had left after playing the investment game and route planner game, which determined their bonus pay for the experiment. Lastly, participants were debriefed, paid and thanked for their participation.

3 Results

3.1 Investment Game

Effect of Agent. To rule out that participants responded differently towards the two agents used in the investment game, a one-way MANOVA was conducted with agent as the independent variable and investment game decision, trust, and liking as dependent variables. No significant effects were revealed (all p's > .250). Therefore, we conclude that participants responded similar to both agents in the investment game.

Main Analysis. To test our main hypothesis, a one-way MANOVA was conducted with mimicry as the independent variable and investment game decision, trust, and liking as dependent variables. Results revealed no significant multivariate effect of mimicry, $F(3,35) = 2.65$, $p = .500$, $\eta_p^2 = .06$. All main effects of mimicry were non-significant (all p's > .131). This analysis did not provide evidence for our main hypothesis. That is, results did not suggest that a mimicking agent was liked and trusted more than a non-mimicking one.

Exploratory Analysis. Separate one-way ANOVA's were conducted with mimicry as the independent variable and expected credits, self-other overlap (pictorial and continuous), perceived risk, perceived competence, and IQ as dependent measures. A main effect of mimicry on IQ was revealed, $F(1, 38) = 5.89$, $p = .020$, $\eta_p^2 = .14$. Participants rated the IQ of the agent in the mimicked condition to be lower (M = 98.00, SE = 3.65) than in the non-mimicked condition (M =110.68, SE = 3.74). All other effects were non-significant (all p's > .208). This analysis did not provide evidence for our exploratory hypothesis. That is, results did not suggest that self-other overlap is a mediator for the digital chameleon effect: mimicry did not increase self-other overlap.

3.2 Route Planner Game

Effect of Agent. To rule out that participants responded differently towards the two agents used in the route planner game, a one-way MANOVA was conducted with agent as independent variable and route planner decision, trust, and liking as dependent variables. No significant effects were revealed (all p's > .301). Therefore, we conclude that participants responded similar to both agents in the route planner game.

Main Analysis. To test our main hypothesis, a one-way MANOVA was conducted with mimicry as the independent variable and route planner decision, trust, and liking as dependent variables. In line with our hypothesis, results revealed a marginally significant multivariate effect of mimicry, $F(3,35) = 2.65$, $p = .064$, $\eta_p^2 = .19$. The main

effect of mimicry on route planner decision was significant, $F(1,38) = 4.83$, $p = .034$, $\eta_p^2 = .12$. Participants gave more routes to the agent in the mimicked condition ($M = 5.50$, $SE = 0.53$) than in the non-mimicked condition ($M = 5.05$, $SE = 0.55$). The main effect of mimicry on trust was marginally significant, $F(1,38) = 2.61$, $p = .058$ (1-tailed), $\eta_p^2 = .07$. In the mimicked condition, the agent was trusted more ($M = 4.40$, $SE = 0.20$) than in the non-mimicked condition ($M = 4.08$, $SE = 0.20$). The main effect of mimicry on liking was significant, $F(1,38) = 5.29$, $p = .027$, $\eta_p^2 = .13$. Participants liked the agent more in the mimicked condition ($M = 4.63$, $SE = 0.15$) than in the non-mimicked condition ($M = 4.29$, $SE = 0.16$). This analysis did provide evidence for our main hypothesis. That is, results did suggest that a mimicking agent was liked and trusted more than a non-mimicking one.

Mediation Analysis. To test whether increased liking mediated the effect of mimicry on the route planner decision, a mediation analysis (following the steps of [22]) was conducted to reveal the direct (Path c) and the indirect effects (Paths a and b) of mimicry on the route planner decision. A Sobel test [23] showed that the indirect effect was marginally significant (Sobel $z = 1.41$, $p = .08$ (1-tailed)). The initial effect of mimicry on route planner decision (Path c) becomes non-significant after controlling for liking (Path c'; see Appendix B), which shows that liking mediates the initial effect. The results from this analysis confirm our expectations that liking, resulting from mimicry, mediates the effect of mimicry on route planner decision.

Exploratory Analysis. Separate one-way ANOVA's were conducted with mimicry as the independent variable and expected credits, self-other overlap (pictorial and continuous), perceived risk, perceived competence, and IQ as dependent measures. Again, a main effect of mimicry on IQ was revealed, $F(1, 38) = 4.27$, $p = .046$, $\eta_p^2 = .10$. Participants rated the agent in the mimicked condition less intelligent ($M = 97.90$, $SE = 3.45$) than in the non-mimicked condition ($M = 108.11$, $SE = 3.54$). Furthermore, a main effect of mimicry on expected credits was revealed, $F(1, 38) = 6.59$, $p = .014$, $\eta_p^2 = .15$. In the mimicked condition, participants expected to get more credits back from the route planner game ($M = 14.45$, $SE = 0.54$) than in the non-mimicked condition ($M = 12.47$, $SE = 0.55$). All other effects were non-significant (all p's $> .136$). This analysis did not provide evidence for our exploratory hypothesis. That is, results did not suggest that self-other overlap is a mediator for the digital chameleon effect: mimicry did not increase self-other overlap.

4 Discussion and Conclusion

The current research investigated whether an agent could persuade participants to trust and like him more by mimicking them. To investigate this question, participants played an investment game and a route planner game with two agents that either mimicked their head movements, or did not. Non-mimicking agents moved their head using the recorded head movements of the previous participant. Results only partly provided support for our hypothesis. That is, for the investment game, results did not suggest that an agent could persuade participants to like and trust it more by mimicking them. For the route planner game however, results suggested that participants liked and trusted a mimicking agent more than a non-mimicking one.

Furthermore, as expected, the effect of mimicry on trust behavior was mediated by liking.

Overall, results suggest some support for our main hypothesis. In the route planner game, the predicted effects were found: participants liked and trusted (on a behavioral and cognitive level) a mimicking agent more than a non-mimicking one. However, in the investment game, we did not find support for our main hypothesis. We argue that there are two plausible explanations as to why the results from the investment game did not support our main hypothesis.

The first, methodological, explanation is that the order of both games was always the same: the investment game was played first. Mimicry could take some time to lead to positive effects, and before deciding in the investment game, participants were exposed to 102 seconds of mimicry in our experiment (versus 195 seconds in [14]). Our duration of mimicry could have been below a certain threshold for positive effects to emerge in the investment game. Future research could identify this specific threshold of mimicry (if there is any), in either a human-human interaction or a human-agent interaction. To be on the safe side, the duration of mimicry should be at least 195 seconds in future research before measuring trust behavior.

The second, theoretical, explanation is that the investment game and the route planner game measure a different kind of behavioral trust. As argued before, in the investment game, *relational trust* is measured: participants can earn credits based on the *intention* of the agent. In the route planner game, it is about *calculative trust*: participants can earn credits based on the *competence* of the agent (see [24] for a detailed discussion of the distinction between relational and calculative trust). In other words, the investment game choice was risky because participants did not know how many credits Jeroen was willing to give back. The route planner game choice was risky because participants did not know how well Max could plan routes for them. Our results suggest that mimicry seems to increase calculative trust in an agent, but not relational trust. Possibly, people do not easily attribute willingness to an agent, explaining why mimicry does not increase relational trust in an agent. Future research could further disentangle the effect of mimicry on both calculative and relational trust by using different measures. In this study, we tried to do just this by using both the investment game and the route planner game. Future research could explore different paradigms with similar differences.

Results do not support our exploratory hypothesis. That is, results do not suggest that the self-other overlap resulting from mimicry is a moderator for the digital chameleon effect. Although mimicry has been shown to lead to increased *neural* self-other overlap [8], we failed to detect an increase of self-other overlap due to mimicry using pictorial measures in our study. Future research could focus on different ways of measuring self-other overlap more directly to further investigate this exploratory hypothesis. In addition, the role of explicit/implicit self-esteem in combination with self-other overlap could be tested in future research.

Unexpectedly, participants rated both agents less intelligent when they mimicked them versus when the agents used recorded head movements. One possibility is that mimicry decreases IQ. Research has shown that positive effects not only depend on the ability to mimic, but also on the ability *not* to mimic [25]. Using our mimicry algorithm, participants were mimicked precisely after a delay of 4 seconds, and every rotation of the head was mimicked. The algorithm therefore resulted in consistent but

rigid mimicry. When humans mimic each other, mimicry is more *flexible*. Humans can choose what movements to mimic, and also when *not* to mimic. Future research could focus on improving mimicry algorithms, to make them model natural mimicry between humans more closely than current algorithms.

Another possibility is that the wording of our question explains the results. We compared the agent to *humans* in our question. Because the voice of the agent was synthetic, and the appearance of the agent was clearly computer generated, comparing the IQ of the agent to the IQ of humans in general might lead to downward comparisons. Mimicry could have made the agent more humanlike, but its non-human attributes (voice and appearance) could have decreased its estimated IQ. Although unexpected, future research could further explore the effect of mimicry on IQ estimations of agents.

In the current study, a specific form of mimicry was used, namely mimicking head movements. As agents become more advanced and more capable of mimicking a wider variety of human behavior, future research could test the effect of other types of mimicry by agents. These could include mimicking the voice [4], the facial expressions [5], and even the mood of the user [6]. We would expect similar effects of these types of mimicry by agents. That is, we do not suggest that mimicking head movements uniquely increases liking and trust of agents. Furthermore, we do not expect that the positive effects of mimicry only occur in an automotive context. Mimicry has been shown to have positive effects in a wide variety of human-human interaction settings.

In conclusion, the current study suggested some evidence that mimicry can be used by agents to persuade their user to like and trust them more. That is, results suggested that mimicry increased liking and trust behavior towards an agent, when the trust situation was competence based. This effect should be replicated in future studies to improve the strength of the evidence. Nevertheless, these results take a first step in making agents inside a self-driving car more trustworthy to be able to persuade drivers to give up control of driving to their cars.

References

1. Markhof, J.: Google cars drive themselves, in traffic. The New York Times (October 9, 2010), `http://www.nytimes.com` (retrieved)
2. Lee, J.D., See, K.A.: Trust in automation: Designing for appropriate reliance. Human Factors 46, 50–80 (2004)
3. Mayer, R.C., Davis, J.H., Schoorman, F.D.: An integrative model of organizational trust. Academy of Management Review 20, 709–734 (1995)
4. Neumann, R., Strack, F.: Mood contagion: The automatic transfer of mood between persons. Journal of Personality and Social Psychology 78, 211–223 (2000)
5. Hsee, C.K., Hatfield, E., Carlson, J.G., Chemtob, C.: The effect of power on susceptibility of emotional contagion. Cognition and Emotion 4, 327–340 (1990)
6. Chartrand, T.L., Bargh, J.A.: The Chameleon effect: The perception-behavior link and social interaction. Journal of Personality and Social Psychology 76, 893–910 (1999)
7. Lakin, J.L., Jefferis, V.E., Cheng, C.M., Chartrand, T.L.: The chameleon effect as social glue: Evidence for the evolutionary significance of nonconscious mimicry. Journal of Nonverbal Behavior 27, 145–162 (2003)

8. Inzlicht, M., Gutsell, J.N., Legault, L.: Mimicry reduces racial prejudice. Journal of Experimental Social Psychology 48, 361–365 (2012)
9. Maddux, W.W., Mullen, E., Galinsky, A.D.: Chameleons bake bigger pies and take bigger pieces: Strategic behavioral mimicry facilitates negotiation outcomes. Journal of Experimental Social Psychology 40, 461–468 (2008)
10. Reeves, B., Nass, C.: The media equation: How people treat computers, television, and new media like real people and places. Cambridge University Press, New York (1996)
11. Tajfel, H.: Experiments in intergroup discrimination. Scientific American 223, 96–102 (1970)
12. Turner, J.C., Brown, R.J., Tajfel, H.: Social comparison and group interest in ingroup favoritism. European Journal of Social Psychology 9, 187–204 (1979)
13. Nass, C., Fogg, B.J., Moon, Y.: Can computer be teammates? International Journal of Human-Computer Studies 45, 669–678 (1996)
14. Bailenson, J., Yee, N.: Digital chameleons: Automatic assimilation of nonverbal gestures in immersive virtual environments. Psychological Science 16, 814–819 (2005)
15. Tukey, J.W.: Exploratory data analysis. Addison-Wesley, Reading (1977)
16. Berg, J., Dickhaut, J., McCabe, K.: Trust, reciprocity, and social history. Games and Economic Behavior 10, 122–142 (1995)
17. Jian, J., Bisantz, A.M., Drury, C.G.: Foundations for an empirically determined scale of trust in automated systems. International Journal of Cognitive Ergonomics 4, 53–71 (2000)
18. Guadagno, R.E., Cialdini, R.B.: Online persuasion: An examination of gender differences in computer-mediated interpersonal influence. Group Dynamics: Theory, Research, and Practice 6, 38–51 (2002)
19. Aron, A., Aron, E.N., Smollan, D.: Inclusion of other in the self scale and the structure of interpersonal closeness. Journal of Personality and Social Psychology 63, 596–612 (1992)
20. Rosenberg, M.: Society and the adolescent self-image. Princeton University Press, Princeton (1965)
21. Bluemke, M., Friese, M.: Reliability and validity of the Single-Target IAT (STIAT): Assessing automatic affect towards multiple attitude objects. European Journal of Social Psychology 38, 977–997 (2008)
22. Baron, R.M., Kenny, D.A.: The moderator-mediator variable distinction in social psychological research: Conceptual, strategic and statistical considerations. Journal of Personality and Social Psychology 51, 1173–1182 (1986)
23. Sobel, M.E.: Asymptotic confidence intervals for indirect effects in structural equation models. In: Leinhart, S. (ed.) Sociological Methodology 1982, pp. 290–312. Jossey-Bass, San Francisco (1982)
24. Earle, T.C.: Distinguishing trust from confidence: Manageable difficulties, worth the effort. Risk Analysis 30, 1025–1027 (2010)
25. Kavanagh, L.C., Suhler, C.L., Churchland, P.S., Winkielman, P.: When it's an error to mirror: The surprising reputational costs of mimicry. Psychological Science 22, 1274–1276 (2011)

Appendix A

The two male virtual social agents used in the experiment. For half of the participants, the investment game agent was the left one, and the route planner game agent the right one. For the other half, vice versa. In the experiment, both agents moved their head (either mimicking the participant or not), blinked their eyes, and moved their lips in synch with synthetic speech.

Appendix B

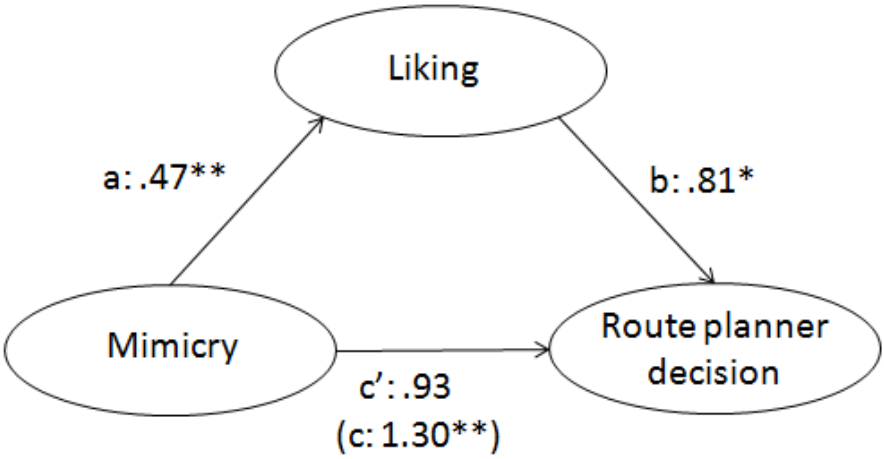

Mediation model depicting the coefficients (Bs) of the direct (Path c) and indirect (Paths a and b) effects of condition on route planner decision. Although the direct effect is significant, this effect becomes non-significant after adding liking as a mediator. * $p < .05$ (1-tailed) ** $p < .05$.