Powerful agents

enhancing an artificial social agent's persuasiveness by increasing its perceived social power

van Iersel, B.M.

Award date:
2012

Eindhoven, January 2012

Powerful Agents:

Enhancing an Artificial Social Agent's Persuasiveness by Increasing Its Perceived Social Power

By B.M. van Iersel

Identity number 0578231

In partial fulfillment of the requirements for the degree of

Master of Science
In Human Technology Interaction
IE&IS – Eindhoven University of Technology

Supervisors:

Dr. J.R.C. Ham (IE&IS)

Prof. Dr. C.J.H. Midden (IE&IS)

Abstract

In increasing the persuasiveness of artificial social agents, earlier research on persuasion in interactions between humans suggests that increasing the *perceived social power* of these agents might also increase their persuasiveness. However, we argue that explicit attributions of social power to artificial social agents will not be functional, because the process by which people make judgments of *social power* is automatic in nature. Therefore, the current research employed an artificial humanlike social agent to influence people's attitudes towards a device that helps save energy in the household environment. The *perceived social power* of the agent was manipulated by integrating *social power cues* into its nonverbal communication. Specifically, we manipulated the *loudness of communication* and the *postural openness* of this artificial agent, two factors that have been found to influence *perceived social power* in interactions between humans. Perceived social power is the degree to which people attribute a sense of power, control and confidence to another person, and this factor has been shown to affect the process of persuasion between humans. Unexpectedly, results indicated that people attribute more *social power* to agents that communicated *louder*, and that used a *closed* body posture. Furthermore, whereas results provided no evidence that the agent's *loudness* and *body posture* influenced the agent's *persuasiveness*, results suggested that the artificial agent was more persuasive when it was perceived as powerful. Our results supported a relationship between the perceived social power of agent and its persuasiveness, and suggested that different manipulations of the agent's social power are needed.

Introduction

Recent research suggested that artificial social agents (e.g., robots, avatars) can function as Persuasive Technology—Technology designed to influence human thinking, attitudes and/or behavior through Persuasion and Social Influence (Fogg, 2003). Artificial social agents have been shown effective in changing our attitudes (Bailenson, Yee, Patel, & Beall, 2008; Ham, Midden, & Beute, 2009; Ham et al., 2011) and our behavior (Iersel, 2011; Vossen, Ham, & Midden, 2010).

We argue artificial social agents will be likely to become increasingly abundant in the future. In our daily lives, we are becoming increasingly accustomed to interacting with artificial social agents (e.g. avatars). We watch animated movies starring digital actors that go on adventures, and that become alive by showing their own personality and sense of free will. We browse through websites with embedded avatars that explain their functionality, and play video games in which we ourselves take on the role of an avatar and interact with the artificial social agents around us. At the same time, technologies that can display artificial social agents, such as computers, smartphones, tablets, and different kinds of displays, are becoming more widespread every year.

Because artificial social agents can function as persuasive technology, and, as described in the previous paragraph, are likely to become abundant in the future, we argue that it is of great importance to understand the determinants of the persuasive power of these artificial social agents. This will allow us create to artificial social agents that are more persuasive, and therefore are more effective in inducing beneficial changes of our attitudes and behavior, for instance

promoting energy conservation behavior. Furthermore, research of persuasive artificial social agents and the determinants of their effectiveness will help us understand the power these agents can have over influencing our attitudes or behavior. This may help us identify possible ethical challenges, and indentify areas where legislative boundaries should be put into place in order to protect people from these new technologies being used in unethical or harmful ways.

A crucial determinant of the persuasive power of artificial social agents could be their *perceived social power*. Classical work in human-human interaction has suggested that power plays an important role in persuasion occurring in human interaction. That is, high-power persuaders (individuals and groups) are more successful at persuasion (Festinger & Thibaut, 1951; French & Raven, 1959).

However, the perspective offered by these classical studies is rather unclear, because the definition of power used in these classical works is about power differences between *groups* of people. That is, these classical works study persuasion processes in group conversations, where subgroups of varying size carry out a joint attempt at persuading other members of the conversation group to change their standpoint to that of the majority of group members. Higher power sources are defined in that context as subgroups consisting of more members carrying out a joint attempt to change the standpoint of a deviating member of the conversation group.

In the current study we use a *different* definition of power, that is, we define power as a *subjective perception of power* that a person has about another individual. We base this definition on a large body of recent research that defines power as either (1) an individual's relative capacity to modify the states of other people by providing or withholding rewards or

administering punishments (Briñol, Petty, Valle, & Rucker, 2007) or as (2) an individual's capacity to modify the states of other people through control or influence over the outcome of relevant situations (Fiske, 1993; Copeland, 1994). We argue that these two definitions are largely similar, because having the ability to modify the states of other people through control over which rewards and punishments are provided or withheld (definition 1) can be seen a specific instance of having control over the outcome of relevant situations (definition 2). In the remainder of this paper we very specifically define (in line with the second definition) **the power of an individual (A)** as *the degree to which another individual (B) perceives individual (A) to have outcome control, which is (A)'s relative capacity to control the outcome of situations which are relevant to (B)*. It is relevant to notice here that this definition describes power of (A) as a *subjective perception* which (B) has about (A). Furthermore, it is important to note that people are often not *consciously aware* that they make these judgments of power, and that the process of making these judgments of *social power* is often a process that people do not spend conscious attention to (see Hall, 2005).

To the best of our knowledge, no recent studies have evaluated the relationship between *power* (in line with the definition presented in the previous paragraph) and *persuasion*. We argue this absence of knowledge on the relationship between *power* and *persuasion* is scientifically relevant to address. We argue so, because classical studies provides us with initial indications that this relationship might exist, and because confirmation of a relationship between *power* and *persuasion* would provide us with a relevant determinant of persuasive power that has not yet been studied in the context of artificial social agents. Exploring the implications of this determinant of this novel field of persuasion furthermore allows us to protect society as a whole

from unethical or harmful forms of persuasion, by providing the insight into this novel technology that is needed to make a good ethical consideration. For these reason, this study aims at evaluating the relationship between *power* and *persuasion*.

We argue that the relationship between *power* and *persuasion* that classical studies suggest for interactions between humans (Festinger & Thibaut, 1951; French & Raven, 1959), could also occur in human-agent interactions through the dynamics described by the Media Equation hypothesis. The Media Equation hypothesis suggested that the dynamics by which humans respond to and interact with technology (e.g. artificial social agents) are comparable to the dynamics by which humans interact with other humans (Reeves & Nass, 2003). To illustrate, Reeves and Nass pointed out a social judgment dynamic that in normal human-human interaction causes an individual who *criticizes* other people to be perceived as *less likable* and as *more intelligent* than an individual who praises others. They then showed that this dynamic also occurs in interactions between humans and computers, causing people to perceive a computer that *criticizes* another computer to be perceived as *less likable* and as *more intelligent* than a computer that praises another computer (Reeves & Nass, 2003). If this process of the Media Equation also causes the dynamics through which *perceived social power* affects *persuasion* in human interaction (we will refer to these dynamics as *social power dynamics)* to occur in human-agent interactions, then it should be possible to improve the persuasiveness of artificial social agents though manipulating their perceived social power. Therefore, we now look at how this perceived social power of an agent could be manipulated.

Studies manipulating *power* commonly use either *explicit manipulations* (e.g. Copeland, 1994; Overbeck & Park, 2001; Briñol & Petty, 2007) or *implicit manipulations* (e.g. Chen &

Bargh, 2001) of power. We define explicit manipulations of power as manipulations, which can be consciously perceived by participants, for instance (1) by engaging participants in a power-related roleplaying exercise, or (2) by providing some participants with greater choice (and therefore greater *outcome control)* than other participants. Implicit manipulations of power on the other hand cannot be consciously perceived by participants, for instance by (1) using semantic cues in a word-completion task that lead to activation of *power*-related concepts, or (2) by using naturally occurring, environmental cues. We will now describe both *explicit* and *implicit* manipulations of social power in greater detail.

Explicit manipulations of power can be consciously perceived by participants. For instance, in Overbeck and Park (2001), power was *explicitly* manipulated by assigning participants to either a role of 'professor' (high power) or to a role of 'student' (low power) in a simulated e-mail interaction. Participants in the professor condition (high power) were presented with emails, supposedly by students, that presented them with a situation and asked them to make a decision on the outcome. For instance a student emailing "I am having a lot of trouble figuring out what you have in mind for the paper due on Monday, and it would really help if I could meet with you to get these points cleared up". With each email, the potential cost of complying with the request in the email was presented, for instance, "Your costs are: Time". Participants in this condition had a high degree of *outcome control*, sending out an email with their decision on the outcome each time. Participants in the student role (low power) were presented with a scenario for which they had to write an e-mail to a professor, for instance, "Professor Smith has assigned a paper that is due on Monday. It is now Thursday afternoon. Several aspects of the paper are unclear; you really need to meet with the professor to clear these

points up". Participants in this condition had no *outcome control*; they had to write an email making a request and subsequently wait for an automated reply. This reply was supposedly by the professor, and decided the outcome of the situation, for instance whether or not they would get the meeting they asked for. Another study using a similar *explicit* manipulation of power was conducted by Briñol and Petty (2007), who manipulated power by engaging pairs of participants in a roleplaying exercise, in which one participant played the role of 'boss' and the other participant played the role of 'employee'. Participants in the 'boss' condition were asked to act as if they had complete *control* over the work processes, the evaluation of employees, and the division of rewards. Participants assigned to the 'employee' role on the other hand were asked to act as if they had no control over these things. Both the manipulations used by Overbeck and Park (2001) and those used by Briñol and Petty (2007) were *explicit,* in the sense that participants could consciously perceive them. In the next paragraphs we will discuss *implicit manipulations*; manipulations to which participants cannot spend conscious attention, and discuss which of these two types of manipulations is more suitable for use in current study.

Implicit manipulations of power *cannot* be consciously perceived by participants. For example, Chen and Bargh (2001) present three studies in which they *implicitly* manipulated power either by (1) using semantic cues in a word-completion task that lead to activation of *power*-related concepts, or by (2) using naturally occurring, environmental cues. In the two studies that manipulated power by using naturally occurring, environmental power cues, participants were led to a professor's office to fill in the consent forms for the study, under the pretense that there was a scheduling conflict with the regular lab room. The professor's office had a large wooden desk with a high, cushioned, comfortable professor's chair behind it, and a

simple wooden guest chair in front of it. Power was manipulated by asking participants to sit down in the professor's chair for the high-power condition, and by asking participants to sit down in the simple wooden guest chair for the low-power condition. Results suggested perceptions of *power* were indeed higher for the participants that had been seated in the professor's chair, as opposed to participants that had been seated in the guest chair.

We argue that implicit manipulations of power are more suitable for manipulating perceived power of *artificial social agents* than explicit manipulations, because explicit manipulations give users the opportunity to form conscious thoughts about the power held by the agent. This could be problematic, because the Media Equation Hypothesis suggests that people *automatically* respond to technology in a social manner (Reeves & Nass, 2003). Even if people's *automatic* social response to our agent would have caused them to respond to the *power cues* communicated by this agent, *conscious* considerations of the level of power held by an avatar could still cause people to conclude that ascribing *power* to an artificial social agent does not make sense. Therefore, we used an *implicit* manipulation of power in this study, which manipulated power outside of conscious attention of users, and therefore, we argue, caused the least conscious considerations about the power of the agent.

Previous studies have indicated that power can also be manipulated *implicitly* by manipulating aspects of nonverbal communication that have a relationship with power (Hall, Coats & LeBeau, 2005). Hall and colleagues presented an elaborate meta-analysis of studies on the relationship between perceived social power and various aspects of nonverbal behavior. Their results suggested that the following four aspects of nonverbal behavior increase an individual's power: (1) *higher loudness of communication*, (2) *more postural openness,* (3) *more*

(successful) interruptions of others in a conversation, and (4) *smaller interpersonal distance between the individual and others* (Hall et al., 2005).

Based on the Media Equation hypothesis (Reeves & Nass, 2003), we argue that when we manipulate an agent's *social power cues* to be consistent with cues used by high-power individuals, users will likewise perceive this agent to be more powerful. Implicitly manipulating power by manipulating these two *social power cues* to be consistent with either high or low power therefore allows us to evaluate whether *social power dynamics* indeed occur in human-agent interactions, and whether agents that are perceived to have higher power are indeed more successful at persuasion.

In the current study, we will use an artificial social agent that employs multiple social cues (serving as persuasive mechanisms) that have been shown to influence people's judgments of power in human-human interactions, specifically the *loudness of communication* and the *postural openness* of the agent (Hall et al., 2005). These specific two *social power cues* were selected because they were the most practical to reliably manipulate in our experimental setup. We evaluate the influence of these two social power cues on the *perceived social power* of the artificial social agent, and whether *this perceived social power acts* as a *mediator* affecting the *persuasiveness* of the artificial agent.

Furthermore, we will evaluate whether using several powers cues in combination might lead to additive effects, because earlier research suggested that various persuasive mechanisms and social cues employed by an artificial social agent employs can influence one another's effectiveness. That is, these studies suggested that the effects of different social cues in the

communication of an agent *interact* (Ham et al., 201; Vossen et al., 2010). To illustrate, Ham and colleagues (2011) used a storytelling robot to present a story about to aversive consequences of lying to participants. After the agent finished presenting the story, participant's attitude towards the lying individual in the story was measured to determine the degree to which the story successfully persuaded users that lying was bad. The robot either use persuasive gestures, or not, and either gazed at participants, or did not. Results indicated that using persuasive gestures led to increased persuasiveness, but only when the robot also used the persuasive strategy of gazing at participants. The study by Ham and colleagues therefore suggested that multiple social cues used by an artificial social can sometimes lead to aversive interaction effects, such as the effect of one social cue disappearing in the absence of another. Therefore, this study will also evaluate whether the used *social power cues* interact, that is, whether the presence of one cue influences the effect of another like previous studies suggest (Ham et al., 2011; Vossen et al., 2010).

We manipulated the *loudness (sound volume level)* of the communication presented by the artificial social agent and the *postural openness* of this agent, evaluating the following main research question: *"What is the influence of the loudness of communication and the postural openness of an artificial social agent on its perceived social power, and what is the effect of this perceived social power on the agent's persuasiveness?"* We will evaluate the following hypotheses:

(1a) *"Louder communication of the artificial social agent presenting a persuasive message results in this agent being perceived as more powerful"*

(1b) *"This increase in perceived social power mediates an increase in the agent's persuasiveness"*

(2a) *"More postural openness of the artificial social agent presenting a persuasive message results in this agent being perceived as more powerful, and results in more persuasion"*

(2b) *"This increase in perceived social power mediates an increase in the agent's persuasiveness"*

To the best of our knowledge, no literature exists on which we could base a hypothesis about what kind of interaction effect could be occur when simultaneously manipulating the *loudness of communication* and the p*ostural openness* of our artificial social agent. Therefore, we do not include a hypothesis about the interaction effect between these two factors, and explore whether or not such an interaction happens to exist.

Method

Participants and design We used a 2 (postural openness: closed body posture vs. open body posture) x 2 (loudness: low volume vs. high volume) between-subjects design. Ninety-one students (72 men and 19 women) were randomly allocated to the four experimental groups. The age of participants ranged from 17 to 48 years old (M = 20.95, SD = 3.73). The experiment lasted 20 minutes, and participants received € 5 for participating in the study.

Materials Participants were seated in an individual room and listened to a persuasive message that was presented by an artificial social agent. The persuasive message used promoted the usefulness of the 8-Port Energy Saving Surge Guard (Figure 1). This device helps save energy by turning off auxiliary devices when the 'main' device is turned off. For example, the Surge Guard can be used to automatically cut power from the printer and speakers when the computer is turned off, preventing these devices from continually using a small amount of electricity. The persuasive message was presented in Dutch by a female voice, and lasted 2 minutes and 24 seconds. For participants in the low loudness condition, our aim was to set the loudness of the persuasive message at a low level that was still *audible for all participants*. For participants in the high loudness condition, our aim was to set the volume of the persuasive message at a high level that

Figure 1: The 8-Port Energy Saving Surge Guard

did not cause annoyance or reactance. A calibrated Decibel meter was used to verify the loudness was equal across the individual rooms used in the study.

The artificial social agent used was a female full-body model. Depending on *body posture condition*, our artificial social agent appeared on the screen in the posture shown in Figure 2 or the posture shown in Figure 3. This *postural openness* manipulation was

Figure 2: Closed Body Posture

Figure 3: Open Body Posture

created based on previous studies which suggest that *high postural openness* is equivalent to (1) legs far apart, and (2) arms far apart, while *low postural openness* is equivalent to (1) legs close together and (2) arms close together (e.g., Cashdan, 1998).

The avatar software used was developed by Haptek, and allowed for the manipulation of the body posture of the avatar. It also allowed for automatic generation phoneme information present a recording of human speech that the software used to give the agent realistic lip synchronization, that is, moving the lips of the agent consistent with the message it pronounces.

Furthermore, participants were asked to fill out a questionnaire that included (1) measures of the perceived social power of the agent, and (2) measures of participants' attitude towards the 8-Port Energy Saving Surge Guard (Figure 1). In current study, participant's attitude

towards the Surge Guard and the agent's perceived social power served as dependent variables. Furthermore, our questionnaire also included exploratory measures for (1) the anthropomorphism of the agent, (2) liking of the avatar, (3) participant's reactance towards the message, (4) the amount of attention paid to the message, and (5) the 'naturalness' *of the posture* of the agent. Please refer to Appendix I for an in-depth description of these exploratory items.

We measured the degree of *perceived social power* that participants attributed to the artificial social agent by taking the mean of a total of *twenty* items, each representing one question from our questionnaire (Crohnbach's Alpha = 0.924). *Six* of these twenty questions asked participants to give their impression of the agent's character on power-related items, answering on a 7-scale ranging between lowest and highest power. For example, *"rate your impression of the agent on the following scale"*: insecure-confident, powerful-powerless, or measly-impressive.

Two of these questions twenty asked participants to rate the feeling they experienced during the presentation of the message, again answering on a 7-point scale ranging between lowest and highest experience of power. For instance, "rate the way you felt while the agent was presenting the message": unmoved-intimidated, or unmoved-impressed. The agent was perceived to have the least power when the participant was unmoved, and the most power when the participant was to some degree impressed or intimidated.

Twelve of these twenty questions asked participants to respond to power-related statements on a 7-point disagree-agree scale, for example, *"rate your level of agreement to the following statement: the agent exuded a sense of power"*. Refer to Appendix II for a more in-depth description of the processes that led to the selection of these particular twenty items.

We measured participant's attitude towards the 8-Port Energy Saving Surge Guard by taking the mean of a total of fifteen items, each representing one question from our questionnaire (Crohnbach's Alpha = 0.941). These items asked participants to give their opinion on to which degree the 8-Port Energy Saving Surge Guard (1) is necessary, (2) is positive, (3) is useful, (4) provides added value, and (5) should be used in more places. Each of these five items was repeated in three different contexts, for instance "indicate to which degree you find the Surge Guard necessary…"; (1) "…in general", (2) "…for others", and (3) "…for yourself: . This resulted in a total of fifteen items.

Procedure Participants were invited to take part in the experiment via an email sent to a selection of the subscribers of the JFS Participant Database, a participant database maintained by the IE&IS faculty of the Technical University of Eindhoven. Upon arrival, participants were seated in a cubicle in front of a computer screen. All instructions were presented to participants through this screen. Participants were informed about the artificial social agent used in this study, and were asked to pay close attention to her message. The agent then proceeded to present the persuasive message, which promoted the usefulness of the Energy Saving Surge Guard (Figure 1). After the agent presented the persuasive message, participants completed a series of questionnaires that measured (1) participants' attitude towards the 8-Port Energy Saving Surge Guard, and (2) the perceived social power of the agent. Furthermore, the questionnaire also measured items intended for exploratory analysis, including measures for (1) the anthropomorphism of the agent, (2) liking of the avatar, (3) reactance towards the message, (4) the amount of attention paid to the message, and (5) the 'naturalness' *of the posture* of the agent.

After completing the questionnaire participants were debriefed, thanked, and paid for participation.

Results

We submitted our measures for *power* and *persuasion* into a GLM with two dependent variables (*power*, and *attitude* towards the Surge Guard) and two binominal independent variables: (1) body posture condition (open vs. closed) and (2) loudness condition (low vs. high), as well as one covariate (*attention* of participant during the message). Because our manipulations of power were implicit and were therefore rather subtle manipulations, we argue that it was crucial that participants paid attention to the message presented by the avatar. Therefore, *attention* was used as a *covariate* in our GLM analysis.

Hypothesis 1a holds true We hypothesized that louder communication of the artificial social agent presenting a persuasive message would result in this agent being perceived as more powerful. Results indicated that sound condition indeed had an effect on *power*, $F(1, 86) = 6.90$, $p < 0.05$, with *low* volume resulting in *lower power* ($M = 4.00$, $SD = 0.92$) than *high* volume ($M = 4.39$, $SD = 0.79$). This indicates with Hypothesis 1a holds true.

Hypothesis 1b cannot be confirmed We further hypothesized that this increase in *perceived social power* mediated an increase in the agent's persuasiveness. Results did not provide evidence for an effect of sound condition on Attitude ($F < 1$) [1] and was thereby not able to support Hypothesis 1b.

The antithesis of Hypothesis 2a holds true Our second hypothesis was that more postural openness of the artificial social agent presenting a persuasive message causes this agent to be perceived as more powerful. Results indicated *posture* condition had an effect on power,

[1] Conducting the same GLM without using Attention as Covariate produced results comparable to the ones mentioned in the paragraphs on Hypothesis 1a and 1b.

$F(1,86) = 6.47$, $p < 0.05$, with *closed* posture condition resulting in *higher power* ($M = 4.40$, $SD = 0.79$) than *open* posture condition ($M = 3.99$, $SD = 0.91$). This means the antithesis of Hypothesis 2b holds true, as *closed body posture* resulted in the agent being *more* power and persuasive (instead of *less* powerful and persuasive, as hypothesized).

Hypothesis 2b cannot be confirmed We further hypothesized that this increase in *perceived social power* mediated an increase in the agent's persuasiveness. Results did not provide support for an effect of posture condition on participant's attitude towards the Surge Guard ($F < 1$), and therefore did not provide support for Hypothesis 2b. A small, non-significant effect of posture condition was found, with *closed* body posture resulting in slightly higher attitude towards the Surge Guard ($M = 5.17$, $SD = 0.95$) than *open* body posture ($M = 5.04$, $SD = 0.86$).

We did not have a hypothesis about the existence of an interaction effect between our *sound condition* and *posture condition*. Results of the GLM indeed did not indicate an interaction between sound condition and posture condition ($F<1$)[2].

Correlation analysis suggest that power and attitude towards the Surge Guard correlate, $r = .24$, $p < .05$. This suggests that the *perceived social power* of an artificial social agent can indeed be a determinant of its persuasive power. We will discuss the implications of this finding in detail in the discussion section.

Furthermore, our measure for naturalness of the posture of our artificial social agent had a mean value on the lower half of the 7-point scale ($M = 3.35$, $SD = 1.27$). Results indicated *posture condition* had an effect on perceived *naturalness*, $F(1, 84) = 5.97$, $p < 0.05$, with participants in the *open* posture condition perceiving our agent as *more natural* ($M = 3.71$, $SD =$

[2] Conducting the same GLM without using Attention as a Covariate produced comparable results.

1.11) than participants in the *closed* body posture condition (M = 3.00, SD = 1.33). We will discuss the implications of these findings in the discussion section.

Discussion

In this thesis we studied the effects of the *perceived social power* of an artificial social agent on its persuasive effectiveness. In our study, two implicit *social power cues* were used to manipulate the *perceived social power* of an artificial social agent. As a consequence, we expected the agent to be more persuasive, that is, more effective in changing participant's attitudes.

Our results indicated an effect of our *loudness* manipulation on the *social power* of our artificial agent, with the agent being perceived as more powerful when it presented its message more loudly. This is consistent with our Hypothesis 1a, and replicates earlier research on social power in human-human interactions (Hall et al., 2005). Results also indicated an effect of our postural openness on the power of the artificial social agent. Participants however perceived the agent with *closed* body posture as *more* powerful than the agent with *open* body posture. This is opposite to the effect in Hypothesis 2a. This finding supports the idea that the *perceived social power* of an artificial agent can be influenced by its body posture, but also suggests that our *body posture* manipulation may have been unsuccessful at providing the right social cue to participants.

Results did not provide support for the *social power* of our artificial agent *mediating* an effect on the agent's persuasiveness. The persuasiveness of our agent was not significantly influenced by our manipulations of its *social power cues.* This means our results cannot support Hypothesis 1b and 2b.

While results did not provide evidence for an effect of the *social power cues* used by our agent on how *persuasive* this agent was, results *did* indicate an effect of the agent's actual

perceived social power on its persuasiveness. This suggests that an agent that is perceived to be more powerful can indeed be more persuasive, even though the process of manipulating this *social power* seems to be more complex than we had hypothesized based on the Media Equation Hypothesis (Reeves & Nass, 2003). Our results indicated that our agent was slightly more persuasive when it presented its persuasive message more loudly, even though this effect was not statistically significant. These findings indicate that a more effective manipulation of the *perceived social power* of our artificial social agent could possibly be successful at increasing the persuasiveness of an artificial social agent.

We argue that one possible reason that our manipulations of *loudness* and *postural openness* did not directly affect *persuasion* can be found in the *implicit nature* of these manipulations, combined with the nature and presentation of our *persuasive message*. We argue so, based on the Elaboration Likelihood Model (Petty & Cacioppo, 1986), that describes that people process information either (1) *centrally*, or (2) *peripherally*. When people process a message *peripherally*, they are more sensitive for *peripheral cues* providing information about the truth and value of this message than when they process this message *centrally*. Two main factors affecting whether a message is processed centrally or peripherally are the perceiver's (1) available cognitive resources and (2) motivation to pay attention to the message (Petty & Cacioppo, 1986). Because participants were placed in cubicles with nothing else to do than listen to our persuasive message, it is likely that they had a large amount of cognitive resources available, and paid moderately high amounts of attention to our message. It is therefore likely that most of our participants processed the persuasive message *centrally*, making them less sensitive to *peripheral cues* included in this message, such as our implicit manipulations of

social power. This is one possible explanation for why our manipulations of the *social power cues* used by the artificial agent did not affect how well it performed at persuading participants about the advantages of the Surge Guard.

To the best of our knowledge, the current study has been the first to show that an artificial social agent that communicates with a *louder* volume is perceived to have higher *social power*. This finding is similar to effects of power in interactions between humans (e.g. Hall et al., 2005). This result is also in line with the Media Equation Hypothesis, confirming that the *dynamic* through which people attribute higher *power* to conversation partners who talk more loudly, which usually only occurred in interactions between humans, could also occur in interactions between humans and artificial agents.

Our finding that participants attributed *more* power to an artificial agent that displayed a *closed* body posture compared to an *open* body posture was the opposite of the effect we hypothesized, and does not replicate research on interactions between humans. One possible explanation for this unexpected finding is that *high social power* is a very strong claim to make for an artificial social agent. This means that the manipulations of *postural openness* would have to be exceptionally realistic and natural, because too little realism of the agent could cause its claim of high social power to be an obviously false claim, causing a backfire effect. Participants however rated both postures of our artificial social agent low on *naturalness*. Therefore, we argue that it is possible that the version of our artificial social agent that had an *open body posture*, and therefore made a claim of *high social power*, backfired because the agent was not perceived to be sufficiently natural and realistic to back up such a claim. This effect is similar to the effect described by the Uncanny Valley (Mori, 1970), which describes that when artificial

social agents pass a certain threshold of *similarity* to a human (e.g., when they have a humanlike face), people become much more perceptive of small differences between the agent and a real human. This sensitivity is reflected in the low rating of *naturalness* of our agent. This backfiring effect could explain our finding that our agent that had a *closed* body posture (that did not make a claim of high social power, and was thereby not likely to backfire) was rated as more powerful than our *open* body posture.

Our results provided no evidence for an interaction effect between the effects of our sound and body posture conditions, neither on a participant's attributions of power to the avatar nor on the extent to which a participant was persuaded by that avatar. This suggests that the effects of these two social power cues are independent of each other, at least when these social power cues are communicated by an artificial social agent. This result is inconsistent with the findings of Ham and colleagues (2011), who found that using several social cues led to aversive affects, that is, to the absence of one social cue causing another social cue to no longer be effective. We argue this finding relevant for future studies looking at the effect of social cues of artificial social agents to verify, in order to establish whether or not different kinds of social cues cause aversive interaction effects.

We suggest that future research investigating the relation between an agent's *perceived power* and its effectiveness at persuading users that uses an artificial social agent to present a persuasive message, and that uses *implicit* manipulations power, should aim to design and present this persuasive message in a manner that makes it likely that the message is processed *peripherally*. We argue so, based on the Elaboration Likelihood Model (Petty & Cacioppo, 1986) that was described earlier. Specifically, care should be taken to (1) reduce the amount of

cognitive resources participants have available and (2) to minimize the personal relevant that the message has to participants. When participants only have a part of their cognitive resources available to listen to the message, and are not motivated to spend too much attention to this message, *peripheral* cues such as *implicit* manipulations of *social power* are likely to have the greatest effect.

For future research into increasing the persuasiveness of artificial social agents that manipulates *social cues* such as postural openness, we suggest pre-testing the perceived *naturalness* of the social cues used in the study, as we argue low rating of naturalness could possibly disturb the effect of the manipulation, or even cause the manipulation to backfire. This suggestion is especially relevant when the artificial social agent used is highly similar to a human (e.g., when it has a humanlike face), and therefore may cause the 'uncanny valley' effect described earlier to occur. We point out that it is likely that such a backfiring effect would occur in a different manner than a backfire effect caused by *reactance.* We argue so, because our results indicated that the degree to which our persuasive message caused negative emotions in participants was not affected by whether the artificial agent had an *open* or *closed* body posture. Future studies could evaluate this claim by controlling for the naturalness and realism of the used *social cues*.

We further suggest future research into increasing an artificial social agent's perceived social power could use a manipulation of the loudness of communication, as our results suggest this manipulation can successfully affect perceptions of power. We argue one interesting direction for future research in this area can found in a paper by Shiomi and colleagues (2010), who evaluated the effectiveness of direct manipulations of *loudness of communication* of an

artificial social agent compared to manipulations of the *whispering cues* used by the agent. Shiomi and colleagues (2010) use an artificial social agent that employs a number of social cues that humans commonly use when whispering. Their results indicated that an agent could successfully employ these *whispering cues* to evoke a sense of confidentiality in people, which allowed the agent to more effectively increase people's motivation to engage into an annoying task. We suggest future research could look at the effect of manipulating both loudness of communication and whispering cues to evaluate the combined effect of these manipulations on persuasion by an artificial social agent.

Another recommendation for future research is including measures for participant's level of attention to the artificial social agent, because our results indicated that people's attention significantly influenced their perceptions of *power*, as well as the degree to which they were *persuaded.*

One more recommendation for future research is that the *social power cue* of interpersonal distance could be a valuable addition to our methods of manipulating *perceived social power* (see Hall et al., 2005). We suggest manipulations of interpersonal distance would likely be most effective in experiments using an Immersive Virtual Environment, because perceptions of the distance to the agent are most similar to real-life perceptions of distance when people are immersed in an Immersive Virtual Environment.

Finally, we stress that the ethical implications of using artificial social agents to influence people's attitudes and behavior are important to consider. Especially when the used persuasion is carried out in (partially) subconscious ways, it is important to weigh the value of people's free

will against the benefits to the individual and to society as a whole of creating a change in attitude or behavior. Therefore, we argue this kind of persuasive technologies should only be applied to create changes in attitude and behavior that benefit both the individual making the change, and society as a whole. The question whether or not people should always explicitly be made aware of (subconscious) persuasion attempt is more difficult to answer, because raising this awareness could render certain persuasion strategies ineffective, and eliminate the benefit that would otherwise have been created through the persuasion. As the sophistication of our persuasive technology increases, this question becomes increasingly important, and we argue that it should be answered by our society as a whole. This is one reason why studies into the possibilities of persuasive technology such as current study could be of vital importance, as they provide us with the information we need to make a good ethical consideration.

In conclusion, we suggest carrying on the line of research initiated by this paper would add value to the fields of persuasive technology and persuasive artificial social agents. Our results suggested that the *perceived social power* of an agent indeed affected its *persuasiveness*, even though affecting this *power* might be a complex venture about which there is still much to learn. Our results suggested one effective way of increasing *power* is increasing the *loudness of communication.* This replicated studies on interactions between humans (Hall et al., 2005 for an overview) in the context of interactions between *humans and agents*. Therefore we showed that, like the Media Equation Hypothesis suggests (Reeves & Nass, 2003), dynamics that occur in interactions between humans can also occur in interactions between humans and agents. The possibility of increasing an artificial social agent's persuasiveness by using *social cues* to enhance its *perceived social power* provides opportunities of creating more humanlike,

persuasive and engaging artificial social agents. These enhanced artificial agents could for instance be employed on the many displays present in our environments, for example to persuade people to adopt pro-social attitudes and behavior, such as stimulating energy conservation behavior or discouraging littering behavior. Using sophisticated, humanlike social cues will allow these artificial social agents of the future to change attitudes and behavior for the better with a powerful, convincing persuasion style that has so far only been achieved by real humans.

Appendix I: Exploratory Measures from Questionnaire

Our questionnaire also included exploratory measures for (1) the anthropomorphism of the agent, (2) liking of the avatar, (3) reactance towards the message, (4) the amount of attention paid to the message, and (5) the 'naturalness' of the posture of the agent.

We measured the anthropomorphism of the agent with a total of twelve items (Crohnbach's Alpha = 0.87). Five of these items asked participants to rate the agent on human-like properties on a 7-point scale. For instance, "Rate your impression of the avatar during the message": fake-natural, artificial-lifelike. Seven of these items asked participants to answer questions rated to anthropomorphism on a 7-point scale ranging between not at all-very much. For example, "Please rate the agents on humanlike characteristics it might have": to which degree does the agent have its own thoughts?

We measured participant's liking of the agent by taking the mean of a total of six items (Crohnbach's Alpha = 0.89). These items asked participants to indicate their opinion of the agent on a 7-point scale related to liking, for instance: boring-interesting, unkind-kind, or unfriendly-friendly.

We measured participant's reactance to the message presented by the agent by taking the mean of a total of eight items (Crohnbach's Alpha = 0.80).These items asked participants to respond to indicate their level of agreement to a reactance-related statement on a 7-point disagree-agree scale. For example, "rate your level of agreement to the following statement": "the agent threatened my freedom of choice", or "the agent made me angry".

We measured the Attention of participants by taking the mean of a total of four items (Crohnbach's Alpha = 0.92). These items asked participants to rate the way their felt during presentation of the message on a 7-point scale related to attention. For instance, "rate the way you felt while the agent was presenting the message": distracted-focused, bored-fascinated, or indifferent-attentive.

We measured the Naturalness of the Posture of the agent by taking the mean of a total of four items (Crohnbach's Alpha = 0.87). These items asked participants to respond to indicate their level of agreement to a statement related to the naturalness of the posture of the agent. For example, "rate your level of agreement to the following statement": "the posture of the agent came across as natural", or "the agent was sat down in a natural posture".

Appendix II: Obtaining our measure for Power

A total of twenty-four items in our questionnaire were aimed at measuring the perceived power of the artificial social agent. Seven of these questions asked participants to give their impression of the agent's character on power-related items, answering on a 7-scale ranging between lowest and highest power. For example, "rate your impression of the agent on the following scale": insecure-confident, powerful-powerless, and measly-impressive.

Five of these items asked participants to rate the feeling they experienced during the presentation of the message, again answering on a 7-point scale ranging between lowest and highest experience of power. For instance, "rate the way you felt while the agent was presenting the message": unmoved-intimidated, unmoved-impressed. The agent was perceived to have the least power when the participant was unmoved, and the most power when the participant was to some degree impressed or intimidated. Twelve of these questions asked participants to respond to power-related statements on a 7-point disagree-agree scale, for instance "rate your level of agreement to the following statement": the agent exuded a sense of power, the agent had a submissive voice. All twenty-four items were recoded such that the lowest end of the answering scale represented low power, and the highest end of the scale high power.

These twenty-four items were submitted to a Factor Analysis in order find the items that together formed the strongest measure of power. From this factor analysis, three factors emerged with eigenvalues of 8.83, 2.84, and 2.24 respectively. Inspection of the items these factors were built up from suggested that these factors represented (1) power and confidence in general, (2) powerlessness and submission, and (3) dominance. A total of twenty items contributed to one of

these three factors, that is, had factor loadings of 0.35 or higher for one or more of these factors in the Rotated Component Matrix obtained from our Factor Analysis.

Our final measure for power was the mean of these twenty items. The Crohnbach's Alpha value for these twenty items was 0.924.

References

Bailenson, J. N., Yee, N., Patel, K., Beall, A. C. (2008). Detecting digital chameleons. *Computers in Human Behavior, 24*, 66-87.

Briñol, P., Petty, R. E., Valle, C., Rucker, D. D. (2007). The effects of message recipients' power before and after persuasion: A self-validation analysis. *Journal of Personality and Social Psychology, 93*, 1040–1053.

Cashdan, E. (1998). Smiles, speech, and body posture: how women and men display sociometric status and power. *Journal of Nonverbal Behavior, 24*, 209-228.

Chen, S., Lee-Chai, A.Y., Bargh, J.A. (2001). Relationship Orientation as a Moderator of Effects of Social Power. *Journal of Personality and Social Psychology, 80*, 173-187.

Copeland, J. T. (1994). Prophecies of power: Motivational implications of social power for behavioral confirmation. *Journal of Personality and Social Psychology, 67*, 264-277.

Festinger, L., & Thibaut, J. (1951). Interpersonal communication in small groups. *Journal of Abnormal and Social Psychology, 46*, 92–100.

Fiske, S. T. (1993). Controlling other people: The impact of power in stereotyping. *American Psychologist, 48*, 621–628.

Fogg, B. (2003). *Persuasive technology: Using Computers to Change What We Think and Do*. San Francisco: Morgan Kaufman.

French, J., & Raven, B. (1959). The bases of social power. D. Cartwright (Ed.), *Studies in social power* (pp. 150–167). Ann Arbor, MI: Institute for Social Research.

Hall, J.A., Coats, E.J., LeBeau, L.S. (2005). Nonverbal Behavior and the Vertical Dimension of Social Relations: A Meta-Analysis. *Psychological Bulletin, 131*, 898–924.

Ham, J., Bokhorst, R., Cuijpers, R. van der Pol, D. Cabibihan, JJ. (2011). *Making Robots Persuasive: The Influence of Gazing and Gestures by a Storytelling Robot on its Persuasive Power*. Proceedings of ISS November 2011.

Ham, J.R.C., Midden, C.J.H. & Beute, F. (2009). Can ambient persuasive technology persuade unconsciously? Using subliminal feedback to influence energy consumption ratings of household appliances. In Chatterjee, S. & Dev, P. (Eds.), *Proceedings of the 4th International Conference on Persuasive Technology 2009*, Claremont, USA, (pp. article no : 29-1/6).

Iersel, B.M. (2011). Can Subliminally Primed Feedback Influence our Behavior? From Vickary's Fantasies to Energy Consumption Behavior Change. *Unpublished Master's Thesis*, Eindhoven University of Technology.

Mori, M. (1970). Bukimi no tani: The uncanny valley (Translated from Japanese by K. F. MacDorman & T. Minato). *Energy*, 7, 33–35.

Shiomi, M., Nakagawa, K., Matsumura, R., Shinozawa, K., Ishiguro, H., Hagita, N.: "Could I have a word?": Effects of robot's whisper. In IROS(2010) 3899-3904.

Overbeck, J.R., Park, B. (2001). When Power Does Not Corrupt: Superior Individuation Processes Among Powerful Perceivers. *Journal of Personality and Social Psychology, 81*, 549-565.

Petty, R.E., Cacioppo, J.T. (1986). The Elaboration Likelihood Model of Persuasion. In L. Berkowitz (Ed.), Advances in experimental social psychology (Vol. 19, pp. 123-205). New York: Academic Press.

Reeves, B., Nass, C. (2003). The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places, Center or the Study of Language and Information – *Lecture Notes*, Centre for the Study of Language and Inf.

Vossen, S., Ham, J.., Midden, C. (2010). What Makes Social Feedback from a Robot Work? Disentangling the Effect of Speech, Physical Appearance and Evaluation. *Lecture Notes in Computer Science*, Vol. 6137/2010, 52-57. DOI: 10.1007/978-3-642-13226-1_7