THE PROTEUS EFFECT:

BEHAVIORAL MODIFICATION VIA TRANSFORMATIONS

OF DIGITAL SELF-REPRESENTATION

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF COMMUNICATION

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Nick Yee

June 2007

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Jeremy Bailenson) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Daniel Schwartz)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Clifford Nass)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

(Byron Reeves)

THE PROTEUS EFFECT:

BEHAVIORAL MODIFICATION VIA TRANSFORMATIONS

OF DIGITAL SELF-REPRESENTATION

Nick Yee, Ph.D.

Stanford University, 2007

Adviser: Jeremy Bailenson

Digital media allows us to make both dramatic and subtle changes to our self-representations with an ease not available elsewhere. These changes can greatly affect how we interact with others in virtual environments. For example, facial and behavioral mimicry can make us more likeable and persuasive. In addition to gaining social advantages, our *avatars* (digital representations of ourselves) can also change how we behave. This occurs via conforming to expected behaviors of the avatar - a process referred to as the Proteus Effect.

I conducted a series of four pilot studies that explore the Proteus Effect. In the first study, I found that participants in attractive avatars walked closer to and disclosed more information to a stranger than participants in unattractive avatars. In the second study, I found that participants in taller avatars negotiated more aggressively in a

bargaining task than participants in shorter avatars. In the third study, I demonstrated that the Proteus Effect occurs in an actual online community. And in the final study, I showed that the Proteus Effect persists outside of the virtual environment. Placing someone in a taller avatar changes how they consequently negotiate in a face-to-face setting.

The two dissertation studies extended these pilot studies by attempting to clarify the underlying process that leads to the Proteus Effect. In the first dissertation study, I isolated and teased out the unique contribution of the Proteus Effect from an alternative explanation - priming. Priming is a process whereby visual stimulus (such as words or photographs) leads someone to behave in a semantically-consistent manner. In the second dissertation study, I extrapolated from existing theories of stereotype formation to examine the consequences of placing users in implausible bodies that fall outside the range of normal human variation (such as a very short or very tall body).

Approved for publication:

By: ____________________

For Department of Communication

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

INTRODUCTION

In today's world, digital media provides people with many ways to represent themselves. From 3D avatars (i.e., digital representations of the self) in online games to voice messages on our cell phones, not only are there a wide variety of self-representations available, but many of these digital self-representations are also easily alterable. For example, the gender or height of our video game avatars can be changed with the press of a button, and models on fashion magazines are routinely air-brushed and "enhanced" in Photoshop. Of course, neither the practice of creating representations of the self nor the act of transforming these representations is a novel or unique affordance of digital media. Statues and photography are examples of non-digital means of self-representation, and make-up is an example of a physical transformation of the self. What digital media does provide, however, is the ability to provide precise yet dramatic transformations with an ease (and reversibility) that was never possible before. As these self-transformations pervade our mediated and networked societies, it becomes important to understand the psychological processes underlying the alteration of someone's self-representation - both in terms of how it may impact one's own attitudes and behavior, as well as how it may influence our interactions with others.

In this dissertation, I will start by offering a framework for understanding self-representations in general. I will then discuss possible transformations of the self, and in particular, focus on a research and theoretical paradigm known as Transformed Social Interaction. After laying these foundations, I will present several lines of research in this area that illustrate how these self-transformations can lead to attitudinal and behavioral modification. In particular, I will describe a series of four pilot studies that focus on a line of research demonstrating that individuals conform to stereotypical behaviors inferred from their digital self-representations - referred to as The Proteus Effect. In particular, my findings demonstrated that digital self-representations have a dramatic impact on behaviors in virtual environments and in subsequent in face-to-face interactions.

In the two dissertation studies, I attempted to clarify the underlying process that leads to the Proteus Effect. In the first dissertation study, I isolated and teased out the unique contribution of the Proteus Effect from an alternative explanation - priming. Priming is a process whereby visual stimulus (such as words or photographs) leads someone to behave in a semantically-consistent manner. In the second dissertation study, I extrapolated from existing theories of stereotype formation to examine the consequences of placing users in implausible bodies that fall outside the range of normal human variation (such as a very short or very tall body).

REPRESENTATIONS OF THE SELF

Researchers in virtual environments have typically made a distinction between *embodied agents* - digital representations controlled by computer algorithms, and *avatars*, digital representations controlled by people. While there has been a great deal of research in social interaction in virtual environments (see Schroeder, 2002, for a review of the existing empirical work on avatars), much of the existing research paradigm has focused on embodied agents. One reason for this disparity, as Bailenson and Blascovich have discussed (2004), is because commercial technology that allows digital avatars to look and behave in real-time like an individual has only recently become available. Consequently, research on avatars, or the effects of visually or behaviorally transforming avatars, has received relatively little empirical attention until recently.

This emerging research interest in avatars rather than embodied agents makes it important to have a framework for considering human representations. My colleagues and I have proposed such a framework (Bailenson, Yee, Blascovich, & Guadagno, in press). We attempted to address several issues with the framework. First of all, we wanted the framework to articulate the underlying variables in the range of possible human representations, both digital and non-digital. Secondly, we wanted the framework to highlight the fact that digital avatars are not unique or novel as human representations. And finally, we wanted the framework to provide a sense of where digital avatars fit in the larger scheme of possible human representations. In virtual environments, an avatar is defined as “a perceptible digital representation whose behaviors reflect those executed, typically in real time, by a specific human being”

(Bailenson & Blascovich, 2004, pg. 65). Of course, when applied to other forms of computer-mediated communication, the definition of an avatar becomes less clear. For example, visual representation is present in a digital photograph and in an appropriate context, such as in an online dating website, seems like a plausible avatar of a person even there is no behavioral correspondence. And of course, textual representations of people in multi-user domains (MUDs), without any graphical representation, and relying only on textual behaviors have long been considered as avatars (for example, see Schiano & White, 1998).

In the framework, my colleagues and I presented a model with three underlying variables - behavioral similarity, form similarity, and whether the representation occurs in real-time or not (see Figure 1). The y-axis of both graphs denotes behavioral similarity - the degree with which the behaviors of the representation correspond with those of the actual person. The x-axis of both graphs denote how visually similar the representation is to the actual person. The graph on the left contains representations that always or usually occur in real-time, while the graph on the right contains representations that always or usually do not occur in real-time.

For example, a real-time audio message, such as over the telephone, conveys a moderate degree of behavioral information (i.e., tone, mood, gender, and so forth), but lacks visual form correspondence. A photograph of a person, on the other hand, is a non-real-time representation with a high degree of visual correspondence, but almost no behavioral correspondence. It is important to note that the digital or analog nature of the representation is largely irrelevant in this framework because many of the examples in the graphs have digital or analog counterparts. For example, photography can be either analog or digital. And textual avatars in MUDs are the digital counterparts of oral (and paper) avatars in table-top role-playing games (e.g., Dungeons and Dragons).

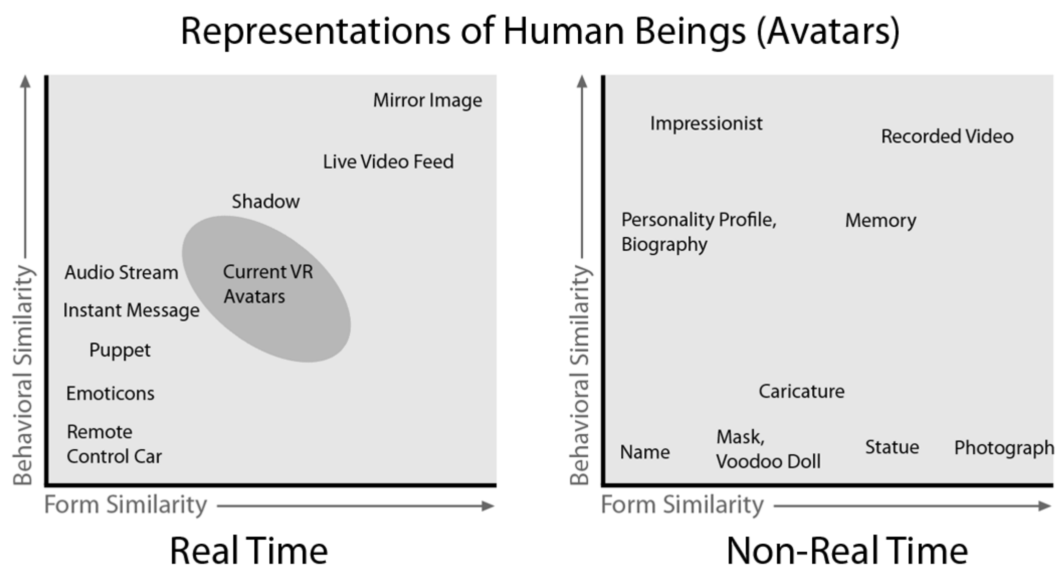

Figure 1. A Framework for Representations of Human Beings.

Mapping out the variables that underlie different forms of self-representation show that digital avatars are one small slice of the possible variation in human representation. Thus, even as the research in self-transformations emphasizes digital avatars, we must be mindful that digital avatars are not the only form of self-representations and that the theoretical findings from studies in virtual environments have much broader implications.

Transformations of the Self

While some human representations (such as photographs and video footage) are usually veridical (i.e., corresponds to reality), it is not hard to notice that there is a pervasive human desire to transform ourselves - to create non-veridical representations. The beauty and cosmetic aisles in our supermarkets - lined with hair coloring products, deodorants, lipstick, and anti-aging moisturizers - speak to this demand. Theorists have long noted that our self-representations are produced to be consumed by others around us. Goffman's approach to understanding identity emphasizes the theatrical and dramaturgical aspects of the self (1959). That is to say the presentations of the self must be understood as a constant performance in front of a social audience. Thus, we choose our clothing, gestures, and mannerisms to give off a desired impression. Psychologists have also noted that the primary function of self-presentation is instrumental. These forms of presentation involve making favorable impressions on others to gain social advantage - such as appearing pleasant or likeable (Jones, 1964; Jones & Pitman, 1982).

Many different kinds of self-transformations are currently available to us in the physical world. My colleagues and I (Bailenson et al., in press) provide a framework for understanding these transformations (see Figure 2). The framework delineates short term from long term transformations, and within each category, there is a further delineation based on whether the alteration is targeted at an individual's appearance, nonverbal behavior, or verbal behavior. For example, the application of make-up is a short term change to a person's appearance, while learning a new language is a long-term transformation of verbal behavior. As the matrix illustrates, there is a wide range of transformations that people can deliberately apply to themselves.

	Appearance	Nonverbal Behavior	Verbal Behavior
Short Term	Haircuts Makeup	Mimicking Ingratiating Gestures	Lying Word Choice
Long Term	Plastic Surgery Dieting	Habit Suppresion Table Manners	Oratory Training Language Acquisition

Figure 2. Categories of Transformations in the Physical World.

Of course, not all transformations are *consciously* controlled by an individual. For example, in the course of normal conversations, people tend to converge on similar speech rates and syntactic complexity - a process known as speech accommodation (Giles & Claire, 1979). The same is also true for postures in dyadic or group interactions (LaFrance, 1982). And as other researchers have argued, automatic mimicry - the process whereby interactants unconsciously mimic each other's gestures - is an innate human behavior that enhances social rapport (Chartrand & Bargh, 1999; Lakin, Jefferis, Cheng, & Chartrand, 2003). Thus, there is a set of automatic behavioral transformations that humans appear to possess.

While simple transformations are easily achieved via haircuts, clothing, or make-up, dramatic transformations (e.g., of gender, height, or weight) typically involve surgical procedures, long-term discipline, or are largely impossible. On the other hand, digital media provides the ability to dramatically and easily transform different aspects of our own digital representations. In the simplest cases, consider the range of avatar icons that users in instant messaging applications or online message boards can choose from. These icons typically offer a wide range of human and non-human choices (see Figure 3).

Figure 3. Examples of Avatar Icons.

Digital media also allows us to change actual photographs of ourselves. Teeth whitening, chemical peels, and cosmetic surgery are available to everyone with a click of the mouse button. The most compelling cases of avatar creation probably occur in online virtual environments where users can customize their avatars' bodies and faces. Every day, millions of people around the world interact via digital avatars in online games (Woodcock, 2006). Users in these online environments are often given a great degree of control over the appearance of their avatars. For example, in the online social world *Second Life*, users can alter their avatar's gender, age, height, weight, musculature, hair style and color, eye shape and color, lip fullness, cheek fullness, nose protrusion, freckles, baldness, and so on (Figure 4). Many other online environments offer a similar range of flexibility in character customization (see Figure 5 and Figure 6). Our physical bodies are no longer the only bodies we can have.

Figure 4. Example of Avatar Customization Panel in Second Life.

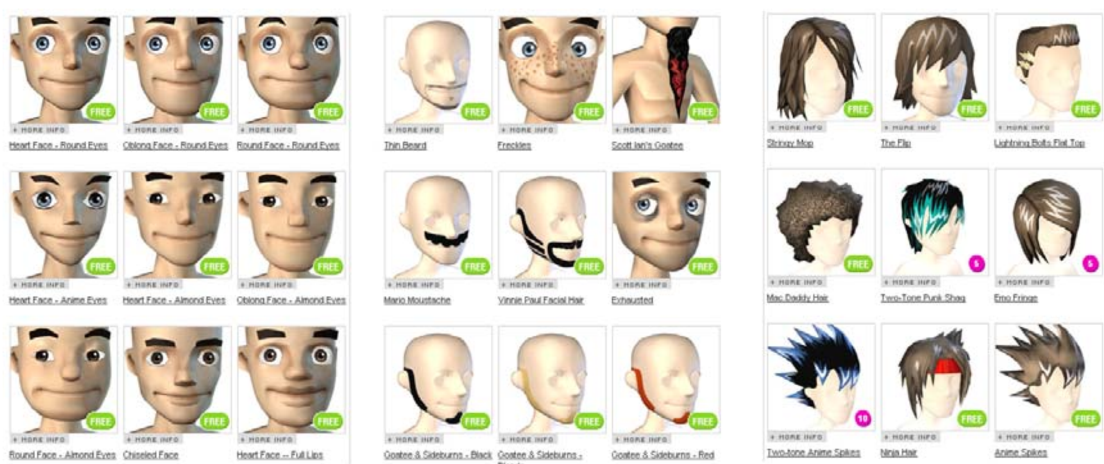

Figure 5. Examples of Customization Choices in the Online Social World Meez.com.

Figure 6. Examples of Avatars Created in the Online Game City of Heroes.

Transformed Social Interaction (TSI)

Digital environments allow us to transform our self representations dramatically, easily, and in ways that are not possible in the physical world. This is particularly true in Collaborative Virtual Environments (CVEs). CVE is a term for digital environments that allow geographically-separated individuals to interact via networking technology, oftentimes with graphical avatars. Thus, CVEs encompass both digital environments created for communication applications as well as online games created for entertainment purposes.

The mediated nature of interactions between users in Collaborative Virtual Environments (CVEs) provides the opportunity to render interactions differently to each participant. Unlike the real world where perceptions of the environment by different interactants are typically congruent, it is possible to create a CVE where each participant sees a different version of the world. As the information from each participant (position coordinates, head rotations, etc.) is relayed through the system, the system can transform that information before relaying it to other participants on a CVE. The term Transformed Social Interaction (TSI) has been coined to describe these real-time transformations (Bailenson & Blascovich, 2004).

TSIs of self-representation are directly relevant to our current theoretical interest. A basic TSI in this category might be the use of a very attractive digital avatar. A more interesting case would be the filtering of behaviors that might be distracting to others or perceived as negative. For example, students in a virtual classroom may want to filter out their actual eye-gaze behavior (i.e., staring out the window) and have a rendered eye-gaze that always looks attentively between the teacher and their notes.

PSYCHOLOGICAL CONSEQUENCES OF SELF TRANSFORMATION

As I've noted before, digital avatars are not the only form of self-representation; many physical forms of self-representation exist. Furthermore, transformations of the self pervade our societies already and are not unique to TSIs. On the other hand, the precision and ease of transformations afforded by digital media and virtual environments allow us to explore the psychological effects of such transformations in ways that would have been prohibitively difficult, expensive, or inhumane in a physical setting. In other words, virtual environments such as CVEs allow us to examine these effects in ways that were not possible before. In this section, I draw from communication and psychology theories, as well as relevant empirical studies, to lay out a framework for considering the consequences of transforming representations of the self.

In this framework, I consider three types of transformation effects. There is an emerging set of empirical studies that suggest these transformations can be used to gain social influence with others, particularly in terms of persuasion or likeability. But these transformations can also have an impact on the self. In particular, these transformations can have an effect on *attitude modification*; being in someone else's body can change how we think about the world. And finally, not only can these transformations change our attitudes, but they can also in fact lead to *behavioral modification*. Subtle changes in our self-representation can alter how we interact with other people. One important distinction between the social advantage effects from the latter two is that social influence effects mainly change the people we interact with while the latter two modification effects mainly change our own attitudes and behaviors.

Gaining Social Advantage

A great deal of research has been performed related to making conversational agents more likeable and persuasive in particular contexts, for example by matching agent personality or mood to the user (Nass & Lee, 2001; Nass, Moon, Fogg, Reeves, & Dryer, 1995; Nass, Robles, Bienenstock, Treinen, & Heenan, 2003). Instead of focusing on agents, I will explore how transformation of our avatars can make us more likeable and persuasive. In this section, I will consider two ways that transformations of visual and behavioral representations can have an effect on social interactions. In particular, I will consider transformations that can be used to gain a social advantage with people that an individual interacts with.

Facial Similarity

The first case study involves the manipulation of facial similarity. Faces appear to processed by specialized areas of the brain (Golby, Gabrieli, Chiao, & Ebenhardt, 2001), and the ability to detect and parse faces is well-developed in humans (Nelson, 2001). More importantly, faces are the primary means of conveying affect (Ekman, 1992; Zajonc & Markus, 1984), and affective cues are known to precede and influence cognitive processes (Zajonc, 1980). Thus, facial transformations might be a fruitful area to investigate in terms of social influence.

Two related well-known preference cues in the psychology literature are familiarity and self-similarity. People prefer familiar things over less familiar things (Zajonc, 1968, 2001) and this phenomenon can be induced in a short period of time even in a laboratory setting - the “mere exposure” effect. Our own faces are objects that we, over a lifetime, become intimately familiar with. And thus the ability to subtly transform our own faces to become more similar to another person’s face might trigger these familiarity preferences.

Similarity-based attraction and preference have also been well-documented. Individuals who are similar to us, whether in appearance or beliefs, are rated more positively, perceived as more attractive, and rated as more persuasive than individuals who are less similar to us (Berscheid & Walster, 1979; Brock, 1965; Byrne, 1971; Shanteau & Nagy, 1979). In fact, this similarity bias can be triggered with other seemingly arbitrary similarities. For example, altruistic behaviors can be elicited when another person shares the same birthday (Burger, Messian, Patel, del Prado, & Anderson, 2004). And in general, people express a higher likelihood of helping a hypothetical person with similar attitudes (Park & Schaller, 2005).

Thus, the ability to transform our facial similarity with a target individual would appear to have an effect on their judgment of us, based on both the familiarity and self-similarity effects. Of course, such a transformation would be difficult to achieve in the physical world. Digital tools of image manipulation, on the other hand, have over the past decade become widely available and affordable. Software, such as the commercially-available Magic Morph, (iTinySoft, 2002) allows the composition of two separate faces at a user-specified ratio. In other words, it has become possible to specify faces that contain a 20% or 40% contribution from another target face.

The impact of this facial-similarity manipulation has been tested in a series of studies in voting behavior (Bailenson, Garland, Iyengar, & Yee, 2006; Bailenson, Iyengar, Yee, & Collins, 2006). In both sets of studies, participants were asked to rate political candidates who either had or had not been morphed with their own photographs. Participants were also asked whether they would vote for the candidates. One study used an undergraduate sample, while the other study used a sample of nationally-representative voting-age citizens. The results of both studies showed that facial similarity is a powerful cue, changing voting behavior even in high-profile elections where a great deal of other information and partisan biases exist, such as the 2004 presidential election. Most importantly, in these two studies with over 700 participants altogether, no participants detected the self-similarity manipulation.

These studies suggest that subtle and largely undetectable transformations of visual representations can have a significant impact on how we judge another person. As surveillance cameras and digital repositories proliferate (NYCLU, 2006), the likelihood that any of our faces becomes captured in our daily lives becomes high. And while current technologies do not allow automated facial similarity manipulations, face recognition software will soon be able to perform that transformation without human monitoring. Thus, facial similarity may become a pervasive manipulation in the near future.

Behavioral Mimicry

Apart from transformations in visual representation, alterations in behavioral representations have also been shown to create a social advantage. This has been demonstrated with regards to behavioral mimicry. Research in interpersonal communication has shown repeatedly that unconscious synchronization of verbal and non-verbal cues are fairly pervasive. This was first described by Kendon (1970) as interactional synchrony, the pattern of hand and body gestures that are reflected and synchronized among pairs and groups of interactants. As mentioned earlier, a similar form of convergence is seen in different aspects of verbal communication, leading to similar speech rates, accents, and syntactic complexity (Cappella & Panalp, 1981; Giles & Claire, 1979; Levelt & Kelter, 1982). Other kinds of behavioral contagion are well-known. For example, these include synchrony in yawning (Provine, 1986), laughter (Provine, 1992), ice-cream consumption (Johnston, 2002), and general mood contagion (Neumann & Strack, 2000).

More recently, it has been shown that deliberately mimicking someone can lead to more favorable responses from them. For example, in a study by Chartrand and Bargh (1999), confederate mimickers were judged as more likeable by participants than confederates who did not engage in mimicry. Other studies have replicated this finding. Waiters who repeat their customer's orders receive larger tips (van Baaren, Holland, Steenaert, & van Knippenberg, 2003). Thus, there appears to be a link between mimicry and pro-social behavior.

Mimicry not only increases an individual's pro-social behavior towards the mimicker, but appears to increase an individual's pro-social behavior in general (van Baaren, Holland, Kawakami, & van Knippenberg, 2004). In that study, individuals who had been mimicked were more helpful and generous towards others than were non-mimicked participants. The reverse has also been shown. When there is a high need for affiliation and rapport, participants are more likely to mimic a confederate's gestures (Lakin & Chartrand, 2003). These lines of evidence support the claim that both unintentional (automatic) and intentional mimicry facilitates and expresses social affiliation and that the process is bi-directional - mimicry facilitates affiliation and prosocial behavior and affiliation goals increase mimicry (Lakin et al., 2003).

While the precision and detail of mimicry is limited in physical situations, CVE systems allow for highly precise levels of mimicry. Because the computer system is already tracking a variety of positional and orientation variables, sending those variables over a network, and then rendering those behaviors on an avatar in another location, it is trivial to playback those behaviors with a pre-determined delay. Moreover, a computer system could build up behavioral repertoire profiles of individual users as templates for general, rather than immediate, mimicry.

The impact of behavioral mimicry in digitally mediated environments was tested recently in an experimental setting (Bailenson & Yee, 2005). Participants interacted with an embodied agent which either mimicked their head movements at a four second delay, or played back the head movements of a previous participant. It was found that participants in the mimic condition were more likely to agree with the agent's argument than participants in the non-mimic (i.e., playback) condition. And much like the studies in facial similarity, levels of mimicry detection were very low (less than 5%). The findings of this study show that an automated program, leveraging data that is already being sent over the network, can produce measurable differences in attitudes via a transformation in behavioral representation. These findings have also been shown to extend to virtual handshakes (Bailenson & Yee, in press). In that study, a haptic device recorded the handshake of two participants. Both participants were told they would receive the other person's handshake via the haptic device; however, one participant always received their own handshake back (i.e., their handshake was mimicked). Results showed that male participants in the mimic condition liked their partner more than those in the normal condition.

Attitude Modification

The case studies in facial similarity and behavioral mimicry illustrate ways in which transformations of the self - both visual and behavioral - can have measurable effects on gaining social influence in interactions. These transformations allow us to become more likeable, more persuasive, and even more likely to be voted for. But apart from influencing others, transformations of self-representations can also change our own attitudes.

Stereotype Reduction

Research into social stereotypes has shown how easily prejudice emerges . For example, it has become well-known that strong beliefs about inter-group differences can be created with minimal, and oftentimes arbitrary, justification (Sherif, Harvey, White, Hood, & Sherif, 1961; Tajfel, Billig, Bundy, & Flament, 1971). Individuals randomly assigned to different groups on the basis of seemingly superficial reasons - such as accuracy of counting large number of dots - will have more positive expectations of members in their own group and more negative expectations of members of the other group despite the fact that there is no rational reason to differentiate members of one group over another.

Others have lamented at how intractable these stereotypes can be. These stereotypes can be triggered without conscious control (Bargh, 1996), and lead to prejudicial interactions unless conscious intervention is applied (P. Devine, 1989). And not only do negative stereotypes impact how people interact with members of other social groups, but these stereotypes in fact create cognitive burdens for members of the minority group as well. For example, negative stereotypes can lead to systematic underperformance via a mechanism known as stereotype threat (Levy & Langer, 1994; Steele, 1997; Steele & Aronson, 1995). In the landmark study (Steele & Aronson, 1995), Black students underperformed in a verbal test when told that the test was designed to measure their intellectual ability. Black participants who were told the test was about understanding different problem-solving strategies scored significantly higher. Steel and Aronson argued that the extra cognitive pressure from the fear of reinforcing a negative stereotype causes systematic underperformance in ability tests.

These findings have led other researchers to explore ways to decrease the accessibility or application of stereotypes. One early approach proposed by Allport was known as the *contact hypothesis* (1954). Allport suggested that social interaction between two different social groups would reduce antagonism between them. Unfortunately, a set of criteria need to be fulfilled for this to occur, such as equal status and mutual interdependence, criteria that seldom are true in our more serious social inequities. Another seemingly intuitive intervention method is known as *thought suppression* - the deliberate avoidance of applying negative stereotypes. But much like asking people to not think about white bears, thought suppression almost always back-fires in the form of repetitive priming (Macrae, Bodenhausen, Milne, & Jetten, 1994; Wegner, Schneider, Carter, & White, 1987).

One promising approach is known as *perspective-taking*. In social interactions, the fundamental attribution error (Jones & Nisbett, 1971) leads us to consider more situational factors when judging ourselves (i.e., I performed poorly because of lack of sleep) and more dispositional factors when judging others (i.e., He performed poorly because he’s not very bright). Thus, when people are forced to observe their own actions (via a video tape), they tend to make more dispositional rather than situational attributions (Storms, 1973). Furthermore, when participants are asked to take the perspective of the person they are observing, participants tend to make situational rather than dispositional attributions (Regan & Totten, 1975).

More importantly, it has been found that perspective-taking leads to an increased overlap between the self and other (Davis, Conklin, Smith, & Luce, 1996), leading to more favorable judgments of the other person. This effect in fact extends to the social group that individual belongs to. Asking individuals to consider everyday situations from the perspective of a member from a different social group has been shown to generate more positive attitudes towards that social group in general (Galinsky & Moskowitz, 2000). In their study, Galinsky and Moskowitz asked undergraduate students to imagine a day in the life of an elderly person. One could imagine that if it were possible to convincingly place an individual into the body of an elderly person instead of simply asking them to imagine it, the effect might be stronger. Of course, such a transformation is difficult to accomplish using traditional means in a laboratory setting.

On the other hand, such a transformation would be much easier to accomplish in a CVE and indeed has been performed in an experimental setting (Yee & Bailenson, 2006). Participants in an immersive virtual reality setting saw their altered representations via a virtual mirror. In the mirror, they saw themselves in the body of an elderly person. In the virtual environment, participants interacted with a confederate and performed a variety of memory and interview tasks. Outside of the virtual environment, participants were then asked to fill out some surveys on their attitudes towards the elderly. It was found that negative stereotyping of the elderly was significantly reduced when participants were placed in avatars of old people compared with those participants placed in avatars of young people. In other words, a transformation in visual representation led to a reduction in negative stereotyping of other social groups.

BEHAVIORAL MODIFICATION

In the previous section, I considered several different psychological consequences of self-transformations - both in terms of attitudinal changes and in gaining social advantage. In this final case study, I will consider how self-transformations can also lead to behavioral changes. After providing the theoretical framing for this effect, I will present four pilot studies in some detail as this line of research provides the motivation and rationale for the two dissertation studies.

Even though the plasticity of our self-representations is an important part of our online identities, the quantitative research in computer-mediated communication (CMC) has tended to focus on the impact of technical affordances on social interaction in online environments. For example, it has been argued that lack of social presence - the sense of interacting with other real people (Hiltz, Johnson, & Turoff, 1986; Short, Williams, & Christie, 1976) or the lack of social cues (Culnan & Markus, 1987; Kiesler, Siegel, & McGuire, 1984) creates an impoverished social environment, while others have shown that relationships develop slower in CMC but are not impoverished in the long term (Walther, 1996; Walther, Anderson, & Park, 1994). Other research has looked at how the narrow communication channels in CMC impacts impression formation (Hancock & Dunham, 2001; Jacobson, 1999; Trevino & Webster, 1992; Walther, Slovacek, & Tidwell, 2001). And while there has been research on self-representation in online environments, the focus has been on the impact of anonymity and authenticity (Anonymous, 1998; Flanagin, Tiyaamornwong, O'Connor, & Seibold, 2002; Jarvenpaa & Leidner, 1998; Postmes & Spears, 2002) - in other words, the gap between the real and virtual self and how that difference changes social interactions.

In the current work, I was instead interested in exploring how our avatars change how we behave online. As we change our self-representations, do our self-representations change our behaviors in turn? As we choose or create our avatars online and use them in a social context, how might our new self-representations change how we interact with others? Thus, I was interested in the impact of our actual self-representations on our behaviors in virtual environments rather than the effects of anonymity or authenticity.

There are several theoretical frameworks that could be used in considering these effects. For example, how we appear may greatly influence how others treat us and how we respond in turn. But beyond such a "reflexive" change, it may also be possible that we conform to stereotypes and assumptions of our own representations independent of how others perceive and interact with us. These various theoretical frameworks will be discussed below in much greater detail.

Behavioral Confirmation

One way in which our visual representations may change how we behave occurs via *behavioral confirmation*. Behavioral confirmation is the process whereby the expectations of one person (typically referred to as the *perceiver*) cause another person (typically referred to as the *target*) to behave in ways that confirm the perceiver's expectations (Snyder, Tanke, & Berscheid, 1977). In the seminal study by Snyder and colleagues (1977), male and female undergraduate students interacted over a telephone. Male perceivers who believed that a female target was attractive caused her to behave in a more charming and friendly manner regardless of how attractive the target actually was. The phenomenon of behavioral confirmation has been demonstrated with regards to many other expectancies and stereotypes, such as hostility (Snyder & Swann, 1978), gender stereotypes (Geis, 1993; Towson, Zanna, & MacDonald, 1989), and racial stereotypes (Word, Zanna, & Cooper, 1974). Thus, in an online environment, a perceiver interacting with a target using an attractive avatar may cause the target to behave in a more friendly and charming manner. In fact, the study by Snyder and colleagues itself occurred in a mediated context (i.e., over the telephone). It is important to note that the source of behavioral change from the effects of behavioral confirmation stem from the perceiver rather than the target. It is the perceiver's behavior that in turn causes a change in the target's behavior.

Self-Perception Theory and Deindividuation Theory

Behavioral confirmation provides one potential pathway for avatars to change how a person behaves online, but might our avatars change how we behave independent of how others perceive us? When given an attractive avatar, does a user become more friendly and sociable regardless of how others interact with them? Another line of research suggests that behavioral modifications can occur without interaction with a perceiver. Bem's self-perception theory (1972) argues that people infer their own attitudes from observing their own behaviors. For example, if a person is well-compensated for something they enjoy doing already, they may come to view the behavior as less intrinsically appealing to them. Other researchers have shown some unexpected implications of this theory. In Valins' study (1966), when participants were made to believe their heartbeat had increased while viewing a photograph of a person, they came to believe the person in the photograph was more attractive. In another study (Frank & Gilovich, 1988), participants given black uniforms behaved more aggressively than participants given white uniforms. The authors of that paper argued that wearing a black uniform is a behavior that the subjects used to infer their own dispositions - “Just as observers see those in black uniforms as tough, mean, and aggressive, so too does the person wearing that uniform” (pg. 83). This effect has also been replicated more recently in a digital environment, where users given avatars in a black robe expressed a higher desire to commit anti-social behaviors than users given avatars in a white robe (Merola, Penas, & Hancock, 2006).

Another line of research has shown that the impact of identity cues is particularly strong when people are deindividuated. Zimbardo (1969) originally used deindividuation theory to argue that urban or crowded areas cause deindividuation which leads to antisocial behavior, however it has also been shown that deindividuation can lead to affiliative behavior as well (Gergen, Gergen, & Barton, 1973). When dyads were placed in a darkened room for an hour, many deliberately touched or hugged the other person. On the other hand, dyads in the fully-lit room talked politely and did not engage in physical contact. Thus, the effects of deindividuation are not necessarily anti-social. The argument that deindividuation can lead to both pro-social and anti-social behavior has also been demonstrated in another well-known study. In a teacher-learner paradigm with electric shock as punishment, subjects in costumes that resembled Ku Klux Klan robes delivered significantly longer shocks than subjects in nurse uniforms (Johnson & Downing, 1979). It was also found that these effects were stronger when subjects were made anonymous in the study. Thus, deindividuation does not necessarily always lead to anti-social behavior as

Zimbardo originally argued, but may in fact cause a greater reliance on identity cues whether those cues are anti-social or pro-social.

In the computer-mediated communication literature, the Social Identity Model of Deindividuation Effects (SIDE model, see Postmes, Spears, & Lea, 1998; Spears & Lea, 1994) argued that factors that lead to deindividuation, such as anonymity, might thus reinforce group salience and conformity to group norms. In this light, deindividuation does not, in and of itself, always lead to anti-normative behavior, but rather, behavioral changes depend on the local group norms (Postmes, Spears, & Lea, 2000). More importantly, behavior that is typically seen as anti-normative, such as flaming (i.e., using profanity and rude language) on message boards (Lea, O'Shea, & Spears, 1992), may in fact turn out to be normative and expected in particular contexts (Postmes et al., 1998).

The Proteus Effect

In the same way that uniforms in the physical world are a transformation of visual representation, non-veridical digital avatars are a more complete self transformation. And thus, we might also expect that behavioral modification occurs via self-observational inference. Whereas the uniform is one of many identity cues in the studies mentioned earlier, the avatar is the primary identity cue in online environments. Online environments that afford anonymity are like digital versions of a darkened room where deindividuation might occur, and indeed, many researchers have suggested that deindividuation occurs online due to anonymity or reduced social cues (Kiesler et al., 1984; McKenna & Bargh, 2000). Thus, we might expect that our avatars have a significant impact on how we behave online. In other words, users in online environments may conform to the expectations and stereotypes of the identity of their avatars. Or more precisely, in line with self-perception theory, they conform to the behavior that they believe others would expect them to have. I will refer to this process as the Proteus Effect.

While the Proteus Effect is similar to SIDE theory, there are several important theoretical differences. Most importantly, SIDE theory emphasizes conformity to local group norms (e.g., becoming more hostile on a hostile message board). On the other hand, the Proteus Effect emphasizes conformity to individual identity cues (e.g., becoming friendlier in an attractive avatar). Thus, theoretically, it would also be possible to pit one against the other - i.e., having an attractive avatar on a hostile message board. I would also argue that having an attribute (e.g., "being attractive") is conceptually different from being amongst a group of individuals who have that attribute (e.g., "being in a group of attractive people"), while the SIDE literature tends to conflate the two. Thus, in a situation where person A in a black uniform interacts with person B in a white uniform, SIDE theory might predict that the social identity of person A would default to the black uniform (i.e., become more aggressive) or the combined colors of the group in question – in other words, gray (i.e., remain neutral). The Proteus Effect would only predict the former. Another point of differentiation is that while the SIDE theory operates on the basis of an existing local group and its social norms, the Proteus Effect should operate even when the user is alone, as self-perception theory isn't predicated on the actual presence of other people, but simply that a person evaluates him or herself from a third-person perspective (i.e., an imagined third party).

PILOT STUDIES

Overview of Pilot Studies

In designing the studies examining the Proteus Effect, it was crucial that the impact of the Proteus Effect was isolated from that of behavioral confirmation. If participants were perceived by others to be attractive and believed themselves to be attractive at the same time, it would be impossible for us to claim that the Proteus Effect occurred independently of behavioral confirmation. Implementing the experimental designs in a CVE allowed us to use TSI to exclude the potential confound of behavioral confirmation. Specifically, I was able to decouple how the participant and the confederate perceived the participant's avatar. Thus, I could manipulate the attractiveness of an avatar from the participant's point of view without changing how the confederate perceived that same avatar. In short, such a manipulation occurs only "in the head" of the participant.

I conducted four studies to explore the Proteus Effect[1]. In the first study, I manipulated the attractiveness of the participant's avatar while the confederate was blind to condition. I predicted that participants in more attractive avatars would become friendlier to confederates than participants in less attractive avatars. In the second study, I manipulated the height of the participant's avatar while the confederate was blind to condition. I predicted that participants in taller avatars would be more aggressive in a negotiation task than participants in shorter avatars. In the third study, I attempted to extend the generalizability of the Proteus Effect by hypothesizing that

[1] Pilot Studies One and Two are currently in press in Human Communication Research (Yee & Bailenson, in press). Pilot Studies Three and Four are currently under review at Communication Research.

taller and more attractive avatars in an actual online community would outperform shorter and less attractive avatars. And in the final study, I examined whether these effects extend beyond the virtual environment. I hypothesized that users given taller avatars in a virtual environment would subsequently be more aggressive negotiators in a face-to-face bargaining task.

Pilot Study One: Attractiveness Manipulation

In the first study, I manipulated the attractiveness of the participant's avatar while the confederate was blind to condition. Studies have shown that attractive individuals are perceived to possess a constellation of positive traits (Dion, Berscheid, & Walster, 1972), and are evaluated more favorably by jurors in courtrooms (Friend & Vinson, 1974).

Hypotheses

Interpersonal Distance

According to nonverbal expectancy violations theory (Burgoon, 1978), when attractive individuals violate nonverbal expectancies, such as moving too close to someone, the positive valence that is created can be socially advantageous (Burgoon & Walther, 1990; Burgoon, Walther, & Baesler, 1992). Given that attractive individuals have higher confidence (Langlois et al., 2000), I hypothesized that:

H1: Participants in the attractive condition would walk closer to the confederate than participants in the unattractive condition.

Self-Disclosure

Friendliness was one of the measures used in Snyder, Tanke, and Berscheid's original study (1977), and in this study I used self-disclosure as a behavioral operationalization. Because attractive individuals tend to be more extraverted and more friendly (Langlois et al., 2000), I hypothesized that:

H2: Participants in the attractive condition would exhibit higher self-disclosure and present more pieces of information about themselves than participants in the unattractive condition.

Method

Design

In a between-subjects design, participants were randomly assigned to have an avatar with an attractive or unattractive face of his or her own gender and then interact with a confederate. I followed the paradigm in the study by Snyder and colleagues (1977) and always used a confederate of the opposite gender. The confederate was blind to the attractiveness condition such that the participant's avatar appeared to the confederate with an untextured face – one which was structurally human but left uncolored.

Participants

Thirty-two undergraduate students (16 men and 16 women) participated in the study for course credit.

Materials

Facial Attractiveness Pretest. I ran a pretest to get subjective determinations of attractive and unattractive faces (for the participants), and also average-attractiveness faces (for the confederates). To minimize the chances that the findings would be driven by idiosyncrasies of a particular face, I chose two faces in each of these three attractiveness conditions. Thus, there were two attractive faces, two unattractive faces, and two average faces for each gender. In total, I used 12 faces in the study.

Fourteen undergraduates from a separate subject population from the main study used a web-based survey to rate the attractiveness of every screenshot's face on a unipolar 7-point fully-labeled construct-specific scale (from "Not Attractive At All" to "Extremely Attractive"). The 12 faces used in the study are shown in Figure 7. The means and standard deviations of their attractiveness ratings are shown in Table 1.

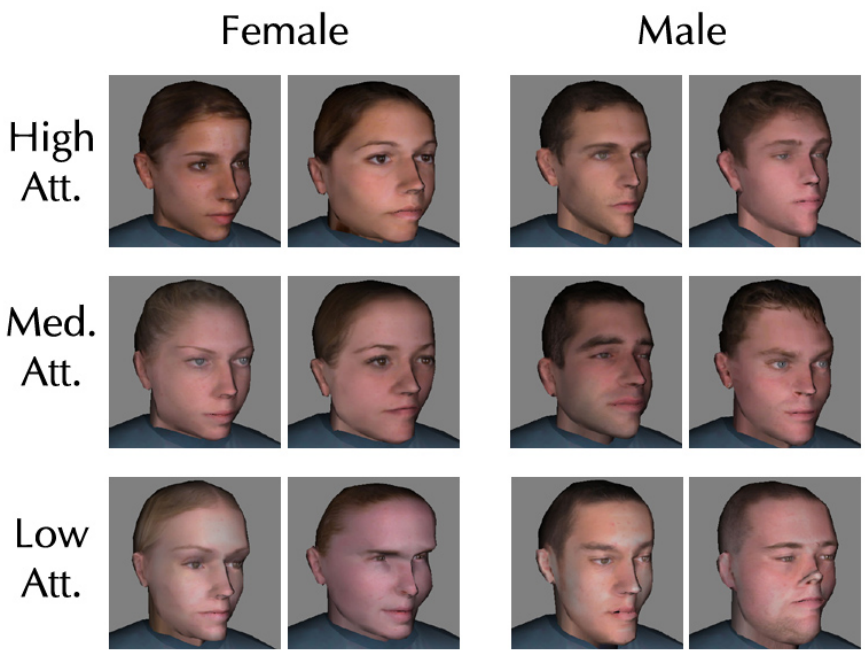

Figure 7. Faces Used in Study.

Table 1.

Means and Standard Deviations of Attractiveness Ratings for Avatar Faces.

	Female		Male	
Attractiveness	Face 1	Face 2	Face 1	Face 2
	M (*SD*)	*M* (*SD*)	*M* (*SD*)	*M* (*SD*)
High	5.50 (1.35)	4.32 (1.25)	4.64 (1.19)	4.04 (1.10)
Medium	3.39 (1.47)	3.50 (1.40)	3.11 (1.34)	2.93 (1.65)
Low	2.29 (1.15)	1.18 (0.55)	1.75 (1.11)	1.21 (0.50)

The Virtual Setting. The virtual setting was a white room that had the same exact dimensions as the physical room participants were in (see Figure 8). Two meters behind the participant was a *virtual mirror* that reflected the head orientation (rotations along pitch, yaw, and roll) and body translation (translation on X, Y, and Z) of the participant with the designated face (See Figure 8).

Apparatus

User Experience. The immersive virtual reality equipment allowed participants to wear a helmet-like head-mounted display and move naturally in the physical environment to move in a virtual environment. Unlike typical computer games where users use joysticks to move their avatars, an immersive virtual system allows users to walk in a virtual environment by walking in the physical environment. When a user moves in any direction in the physical environment, the virtual environment adjusts accordingly, and thus it appears to the user that they have moved the same amount in the virtual world at the exact speed that the user moved physically. The same is true for head movement. To look around the virtual environment, a user would tilt or turn their head in the physical world and the virtual environment would be updated accordingly. Technical details of the immersive virtual reality system are provided below.

Translation Tracking. An optical-tracking system (WorldViz PPT) of four cameras tracked an infrared-emitting light mounted on top of the head-mounted display. The system had a resolution of 1 part in 30,000 or approximately 0.2mm in a 5m square workspace at an update rate of 60 Hz.

Rotation Tracking. An Intersense IS300 cube mounted on to the head-mounted display was used to track the pitch, yaw, and roll of the participant's head (with an update rate of 180 Hz). The Intersense cube had an accuracy of 1 degree in yaw, and 0.25 degree in pitch and roll.

Display. Perspectively-correct stereoscopic images (i.e., separate images for each eye) were rendered at an average frame rate of 60 Hz. Participants wore an nVisor SX head-mounted display that featured dual 1280 horizontal by 1024 vertical pixel resolution panels that refreshed at 60 Hz. The system latency, or the amount of delay between a participant's movement and the resulting visual update in the head-mounted display, was a maximum of 65ms when accounting for the tracking and the rendering system combined. Figure 8 portrays the equipment setup.

Procedure

Three researcher assistants were present during each trial - the lead research assistant, the male confederate and the female confederate. The confederate in the trial was always the opposite gender of the participant and remained blind to condition. Once the virtual world was loaded, participants saw themselves in a room that was exactly the same dimensions as the physical lab room, as depicted in Figure 8.

Participants were then asked by the lead research assistant to turn around 180 degrees and asked to verify that they saw a mirror in front of them. After verbal affirmation, participants were then told that this is how they appeared to others in the virtual room. Several exercises (head-tilting and nodding in front of the mirror) were used to make sure participants had enough time to observe their avatars' faces. Every participant was thus exposed to the designated face for between 30 to 45 seconds.

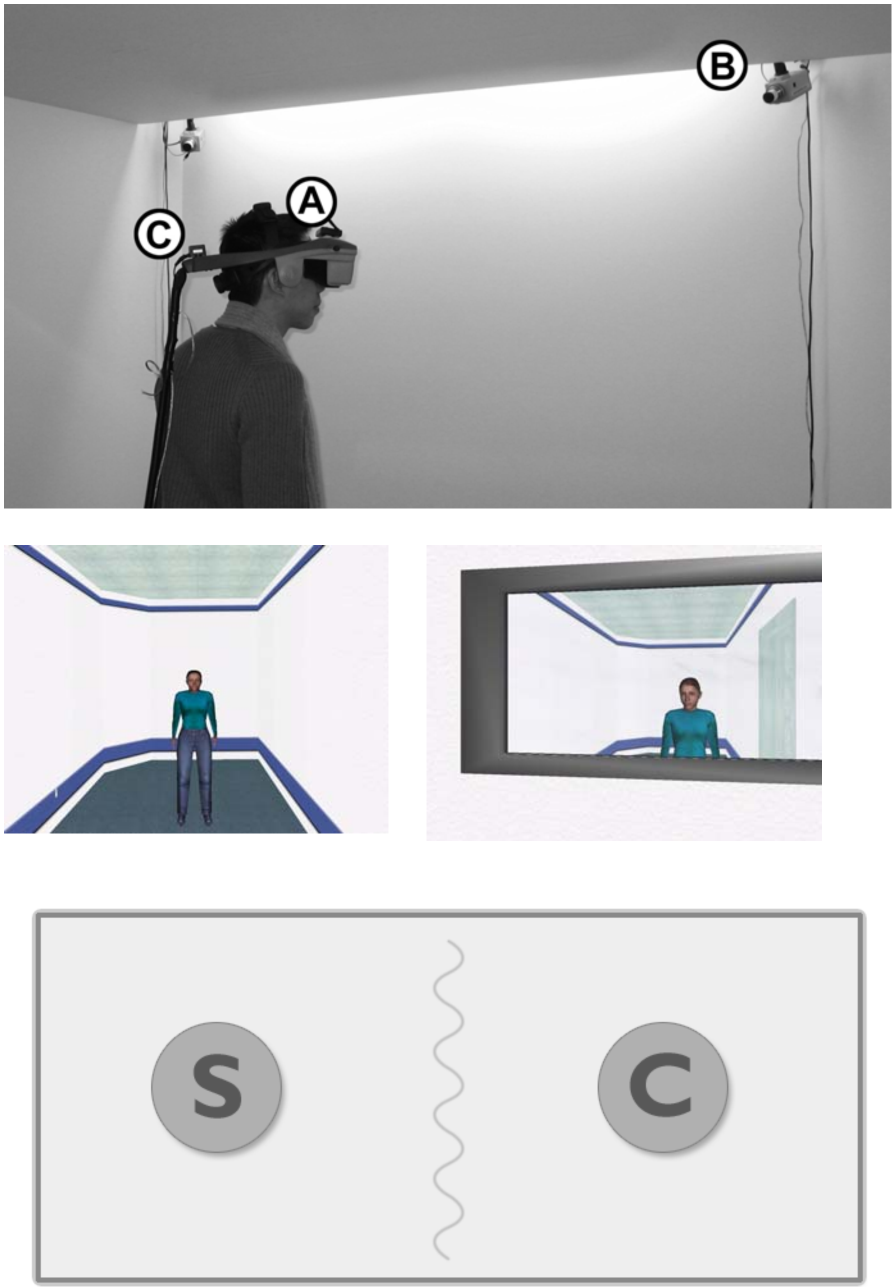

Figure 8. Equipment Setup.

The equipment setup is shown in the top panel. In the lab space, the participant wears the HMD (A). The orientation device (B) attached to the HMD tracks rotation while the cameras (C) are used for optical tracking of the participant's position in the in the room. The virtual room with the confederate is shown in the middle left panel. In the middle right panel is the participant's view of the mirror. In the bottom panel is a diagram showing the layout of the room, the position of the Subject (S), the position of the Confederate (C), and the curtain.

Participants were then asked to turn back around to face the front (i.e., their original orientation). The lead research assistant then introduced the confederate to the participant. The confederate followed a strict script that was displayed in their HMD so they could follow the specific verbal procedures while interacting with the participant inside the CVE. First, participants were greeted and asked to walk closer to the confederate. When the participant stopped or asked whether the distance was close enough, the confederate would then ask them to move a little closer. The confederate then asked the participants to introduce themselves. When the participants stopped or asked whether what they said was enough, the confederate asked the participants to say a little more. If the participants ever asked the confederate any other question, the confederate would reply with "I'm sorry. I can't answer that question. Let's continue".

Measures

Interpersonal Distance. The distance between the participant and the confederate was automatically tracked by the VR system. Previous research has validated the interpersonal distance measure inside CVEs (Bailenson, Blascovich, Beall, & Loomis, 2003).

Self-Disclosure. The amount of self-disclosure was measured by counting the number of pieces of information that participants gave during the two introduction prompts near the beginning of the conversational portion of the study (e.g., "Tell me a little bit about yourself" and "Tell me a little more"). Two blind coders were asked to count the number of pieces of information given by the participants. Every tape recording was coded by two blind coders and the coder inter-reliability was .84.

Results and Discussion

To ensure that the attractiveness manipulation was not so obvious as to elicit strong demand characteristics, participants were asked to write a paragraph and guess the intent of the experiment. Two coders blind to experimental condition read through these responses. Most participants guessed that the goal was to study conversational dynamics in VR as compared with face-to-face interactions. According to both coders, no participant mentioned attractiveness or mentioned that they thought the avatar's attractiveness was manipulated in the study.

Interpersonal Distance

I ran a t-test with attractiveness as the between-subject variable and the final distance as the dependent variable. Participants in the attractive condition walked significantly closer to the confederate (M = 0.98, SD = 0.36) than participants in the unattractive condition (M = 1.74, SD = 1.20), t[30] = -2.42, p = .02, d = .40.

Self-Disclosure

I performed a t-test using attractiveness as the between-subject variable and the self-disclosure count as the dependent variable. Participants in the attractive condition revealed significantly more pieces of information (M = 7.19, SD = 2.77) than participants in the unattractive condition (M = 5.42, SD = 1.56), t[30] = 2.23, p = .03, d = .38.

The results from the first experiment provided support for the Proteus Effect - that our self-representations shape our behaviors in turn. Participants in the attractive condition were willing to move closer to the confederate and disclosed more information to the confederate than participants in the unattractive condition. More importantly, this effect was measurable and significant immediately after only a brief exposure to the mirror task. The effect size in the current study—interpersonal distances changes of almost a meter—are quite large, much more so than effects found in previous studies on interpersonal distance (Bailenson, Blascovich, Beall, & Loomis, 2003) which were less than 15 centimeters. The reason the current manipulation produced such a drastic effect is most likely due to the more personal nature of the social interaction.

Pilot Study Two: Height Manipulation

In the second pilot study, I attempted to replicate the Proteus Effect with another manipulation - height. Because height is more often associated with self-esteem and competence rather than friendliness (Young & French, 1996), I employed a different behavioral measure. Instead of a proximity and self-disclosure task, a negotiation task - the “ultimatum game” (Forsythe, Horowitz, Savin, & Sefton, 1994) - was used as a behavioral measure of confidence. In the ultimatum game, two individuals take turns to decide how a pool of money should be split between the two of them. One individual makes the split and the other must choose to either accept or reject the split. If the split is accepted, the money is shared accordingly. If the split is rejected, neither of them gets the money. I hypothesized that participants with taller avatars would be more confident and be more willing to make unfair splits in their own favor than participants in shorter avatars.

Method

Design

In a between-subjects design, participants were randomly assigned to have an avatar that was shorter, taller or the same height as a confederate who was of the opposite gender. I relied on demographic data to assign the base height and height differences in the study. From the NHANES 2003-2004 data set (NCHS, 2004), I calculated the mean and standard deviation of height among Caucasians aged 18 to 22 in the US population. The mean height was 171.5 cm (or 5 feet and 7.5 inches) with a standard deviation of 10.2 cm. While men and women have different average heights, I decided to use the same base height across all conditions to avoid confounding height with gender in the experimental design.

In my study, the confederate had a base height of 172 cm. In the short condition, participants were 10 cm (one standard deviation) shorter than the confederate. In the tall condition, participants were 10 cm taller than the confederate. In the same height condition, participants were the same height as the confederate. Also, the confederate was blind to the height condition and the participant's avatar always appeared to the confederate as the same height. In cases where this would cause a mismatch in eye-gaze, trigonometry was used in the background to correct for the angle of eye-gaze so that the two could maintain eye contact. In other words, confederates did not know the experimental condition and always perceived the participant as the same height as themselves.

Participants

Participants were 50 undergraduate students who were paid ten dollars for their participation.

Materials

The physical lab and the virtual setting of this study were identical to the ones described in Pilot Study One except there was no mirror in the virtual room.

Apparatus

The apparatus used in Experiment Two was identical to the apparatus described in Pilot Study One.

Procedure

The initial procedure for this study was identical to that in Pilot Study One except for the virtual mirror task. After the participant had been introduced to the confederate, the researcher explained the money sharing task. A hypothetical pool of $100 was to be split between the confederate and the participant. One of the two would designate a split. The other would either accept or reject the split. If the split was accepted, the money would be shared accordingly. If the split was rejected, neither would receive any money. The participant was told there would be four rounds of this game and that the researcher would alternate as to who would be making the split for each round.

The participant always designated the split in the first and third rounds. The confederate was instructed to always accept a split as long as it did not exceed $90 in favor of the participant. The confederate always designated a split of 50/50 in the second round and 75/25 (in the confederate's favor) in the fourth round. These two ratios were chosen to represent a fair and unfair split. After the money sharing task, the participant was taken out of the virtual setting.

Measures

Monetary Splits. The split offers were recorded by the research assistant during the negotiation task.

Results and Discussion

To ensure that the height manipulation was not so obvious as to elicit strong demand characteristics, participants were asked to guess the intent of the experiment. Two coders blind to condition read through the responses. Most participants guessed that the goal was to study conversational dynamics in virtual reality as compared with face-to-face interactions. According to both coders, no participant mentioned height or guessed that height was manipulated in the study.

Negotiation Behavior

There were three measures of interest: amount offered by participant in the first round (from hereon referred to as *split one*), amount offered by participant in the third round (from hereon referred to as *split two*), and whether the participant accepted the unfair split by the confederate in the final round (from hereon referred to as *final split*). Three outliers (more than 3 standard deviations from the mean) in split one and split 3 were excluded from analysis. The cut-offs were 88.5 and 84.2 respectively.

I ran an ANOVA with height as the between-subject factor and split one as the dependent variable. The effect of height was not significant ($F[2, 47] = 0.63$, $p = .53$, $\eta^2 = .03$), see Table 2.

Table 2.

The Means and Standard Deviations of Interpersonal Distance and Split One across Height Conditions.

Height	Split One	Split Two	Final Split
Short	54.99 (12.47)	52.06 (7.30)	0.72 (0.46)
Normal	58.69 (15.85)	55.69 (8.10)	0.31 (0.48)
Tall	53.75 (10.25)	60.63 (6.55)	0.38 (0.50)

I then ran an ANOVA with height as the between-subject factor and split 3 as the dependent variable. There was a main effect of height ($F[2, 46] = 5.64$, $p = .006$, $\eta^2 = .20$). A post-hoc test using Tukey's HSD showed that participants in the tall condition split the money significantly more in their own favor ($M = 60.63$, $SD = 6.55$) than participants in the short condition ($M = 52.06$, $SD = 7.30$), p = .004. See Table 2 for all means and standard deviations of the splits by condition.

Finally, to test the effect of height on the acceptance rate of the final unfair offer, I ran a logistic regression using acceptance rate as the dependent variable and height (recoded short as 1, normal as 2, and tall as 3) as the independent variable. Height was a significant predictor of acceptance rate, $\chi^2(1, N = 50) = 4.41$, $p = .04$. Prediction success for acceptance of the unfair offer was 54% and it was 80% for rejection of the unfair offer. Participants in the short condition were about twice as likely to accept the unfair offer (72%) as participants in the normal (31%) and tall condition (38%).

I was surprised that the effect of height on negotiation did not emerge until the second split. Informal discussion with the research assistants and review of the recordings suggest that many participants were "testing the waters" in the first split, but became more bold in the second split. In any case, the effect of height on the second split was highly significant and suggests that the manipulation of height does affect negotiation behavior, but that these effects may emerge over time.

In summary, my findings from Pilot Study Two extended the Proteus Effect with a different manipulation. Participants in the tall condition were significantly more likely to offer an unfair split than participants in the normal and short conditions. At the same time, participants in the short condition were significantly more likely to accept an unfair split than participants in the normal and tall condition. Thus, my findings from the negotiation task support the Proteus Effect.

Pilot Study Three: Generalizability to Natural Setting

Introduction

While my laboratory findings helped establish causality and showed that manipulations to digital self-representation can indeed change a user's behavior, they failed to convincingly generalize to natural interactions in actual online environments.

First of all, social interactions in actual online environments are neither limited to dyadic interactions or to short, scripted responses. Secondly, showing that the Proteus Effect does occur in a natural setting would allow generalizing the effect beyond a small sample of undergraduate students. Specifically, as other studies have shown, the behaviors and attitudes of college students may not be representative of the general population (Bailenson, Shum, Atran, Medin, & Coley, 2002; Iyengar, 2002; Peterson, 2001; Shruptine, 1975). Thus, showing the Proteus Effect in a naturalistic online environment would also increase the generalizability of the finding.

To test the Proteus Effect in a naturalistic setting, I collected data from an actual avatar-based environment with thousands of daily users - the online game

World of Warcraft (WoW). In particular, based on my findings from Pilot Study One and Two, I hypothesized that taller and more attractive avatars in WoW would outperform shorter and less attractive avatars. It is important to note that neither height nor attractiveness have any direct functional benefit in the game—for example, tall avatars don't walk faster than shorter avatars and attractive avatars don't have special skills.

Method

Data Collection

The main advancement metric in WoW is known as *levels*, a rank of progression between 1 and 60. Game-play in WoW is largely structured around attaining higher levels by completing a variety of tasks. An automated script was implemented that performed a *census* every 15 minutes on three WoW servers[2]. The census gathered information about all unique characters such as their race, their location, and the level they were currently at within the game. To ensure that characters were not only sampled from a given time of day or day of the week, census data was collected from these three servers over a period of seven days (2016 independent censuses were conducted). All unique characters were then extracted from the census data points. Altogether, the sample yielded 76,843 individual characters. The average character success level was 23.6 (SD = 20.08).

Avatar Attractiveness

To calculate the average attractiveness of each of the eight races that serve as avatars in the game, I captured screenshots of each race from the character generation

[2] This research was conducted in collaboration with the following researchers at PARC: Nicolas Ducheneaut, Eric Nickell, and Robert Moore.

segment of the game. The character generation segment has a "random generation" option that allows a user to randomly select and combine one of the possible hair styles, face textures, skin tones and hair colors available of the selected race. I used this random generation feature to generate the sample avatars of each race. Because the gender of every character can also be selected, I captured the first four randomly generated avatars for every race and gender. Thus, altogether, 64 images of WoW avatars were generated.

Twenty-two undergraduate students were presented these 64 images in randomized order and were asked to rate the attractiveness of every avatar. They responded on a 7-point bi-polar fully-labeled scale ranging from "Extremely Unattractive" (1) to "Extremely Attractive" (7). The average attractiveness ratings of each of the eight races are shown in Table 3.

Table 3.

Average Attractiveness and Height Ratings of Races in World of Warcraft.

	Attractiveness Mean (SD)	Height (in pixels)
Dwarf	2.97 (1.70)	299
Gnome	3.48 (1.54)	194
Human	4.71 (1.53)	447
Night Elf	3.10 (1.83)	515
Orc	2.05 (1.68)	424
Tauren	2.60 (1.70)	495
Troll	2.47 (1.65)	462
Undead	2.13 (1.75)	457

Height

To calculate the height of every race, I measured the number of pixels of their avatar's height in a digital imaging application. The height of all avatars of each race is held constant in WoW. In other words, all Gnomes have the same height, all Night Elves have the same height, and so on. The heights of the eight races are presented in Table 3.

Results

A multiple regression with character height and attractiveness as the independent variables and character level as the dependent variable revealed the overall model was significant (*Adj.* $R^2 = .01$, $F[2, 76840] = 264.20$, $p < .001$). Both predictors were also significant. The more attractive characters were, the higher level they achieved ($\beta = .05$, $t[76840] = 14.77$, $p < .001$). Also, the taller characters were, the higher level they achieved ($\beta = .07$, $t[76840] = 18.60$, $p < .001$).

The direction of the coefficients supports my hypotheses. The model suggests that characters of the most attractive race - Humans ($M = 4.71$, $SD = 1.53$) - are expected on average to be 2.9 levels higher than characters of the least attractive race - Undead ($M = 2.13$, $SD = 1.75$). Also, characters of the tallest race - Night Elves (height = 515) - are expected on average to be 4.5 levels higher than characters of the shortest race - Gnomes (height = 194). Thus, the model accounts for a difference of 7 levels (out of 60) based solely on the height and attractiveness of the character. Given that it takes on average 480 accumulated hours of playing time (i.e., the equivalent of 3 normal work months) for a character to reach level 60 (PARC, 2005), 7 levels is a substantial difference. In sum, while I cannot account for the causal direction of this effect (i.e., more serious players might prefer taller, attractive avatars), the amount of variance that the appearance of an avatar covers in terms of the performance of a game player is quite high. Moreover, these results dovetail with findings from the experimental studies. As we saw in Pilot Studies One and Two, manipulations in an avatar's height and attractiveness made significant differences in short-term behavior in a dyadic interaction. The findings from the WoW data suggest that users do not habituate to their avatar's appearance and that these behavior changes are sustained over time.

Pilot Study Four: Transfer to Face-to-Face Interactions

Introduction

Findings that support the Proteus Effect also prompt the question of whether these effects may persist outside of the virtual environment. For example, given that the average user of online role-playing games spends 20 hours per week interacting with other people via their avatar (Yee, 2006), it is important to understand whether behavioral changes that occur due to the Proteus Effect in digital environments come to affect how these users interact with others in face-to-face settings.

This "media effects" inquiry has been studied in television (Bandura, Ross, & Ross, 1961; Gerbner & Gross, 1976; Krcmar & Vieria, 2005), music (Christenson & Roberts, 1998), and video games (Anderson & Bushman, 2001; Bushman & Anderson, 2002) among other media. Thus, it may be the case that the effects of the digitally-mediated experience may transfer to face-to-face interactions. Specifically, studies in Social Learning Theory (Bandura, 1977, 1986; Miller & Dollard, 1941) have shown that behavioral change can occur via modeling another person, also known as vicarious learning. When users are taken out from the virtual environment, they may model the modified behaviors of their avatars (i.e., themselves), in the same way that adolescents model the behaviors of characters they see on television shows (Larson, 1995). Thus, users with taller avatars may in turn become more confident in subsequent face-to-face negotiations because they model their avatar’s behaviors. Pilot Study Four tests this hypothesis.

Method

Design

I conducted a between-subjects design almost identical to Pilot Study Two. I assigned participants to two height conditions. In the *short* condition, participants were given avatars that were 10cm shorter than the confederate’s avatar. In the *tall* condition, participants were given avatars that were 10cm taller than the confederate’s avatar. After participants performed the negotiation task in the virtual environment, they were taken out of the virtual environment and performed the negotiation task again with the confederate face-to-face.

Participants

Participants were 40 undergraduate students (18 female, 22 male) who were paid ten dollars for their participation.

Materials and Apparatus

Materials and apparatus were identical to that used in Pilot Study Two.

Procedure

The virtual reality portion of this was conducted in the same way as that described in Pilot Study Two. After the virtual reality portion of the study, participants were taken out of the virtual environment and told they would perform the negotiation task again face-to-face with the same person (i.e., the confederate). Before their head-mounted displays were removed, the researcher drew the curtains. The participant and confederate were then seated in chairs (adjusted to be the same height) with the curtain between them. Once they were ready, the researcher drew open the curtain and guided the participant and confederate through the same negotiation task.

Measures

Negotiation Task Performance. The split offers were recorded by the research assistant during the negotiation task. As in Pilot Study Two, there were three decisions of interest in each negotiation task - the participant's first offer, their second offer, and whether they accept or reject the unfair offer.

Results

A series of repeated measures ANOVAs were conducted with Trial (in VR and outside VR) as the within-subject factor, Height Condition (short and tall) as the between-subject factor, and the three decisions (first offer, second offer, acceptance of unfair offer) were used successively as the dependent variables.

In the repeated measures ANOVA with the first offer as the dependent variable, there was a significant main effect of Height Condition ($F[1,38] = 4.75$, $p = .04$, $\eta^2 = .11$). Participants in the tall condition ($M = 54.56$, $SE = 1.04$) were significantly more likely to offer splits in their own favor than participants in the short condition ($M = 51.37$, $SE = 1.04$). There was also a significant main effect of Trial ($F[1,38] = 4.34$, $p = .04$, $\eta^2 = .10$). Participants interacting face-to-face were significantly more likely to offer splits in their own favor ($M = 53.75$, $SE = 1.00$) than participants inside of VR

(M = 52.18, SE = .60). The interaction was not significant ($F[1.38] = .16$, $p = .69$, $\eta^2 = .003$).

In the repeated measures ANOVA with the second offer as the dependent variable, the main effect of Height Condition was not significant ($F[1,38] = .13$, $p = .71$, $\eta^2 = .003$). There was, however, a significant main effect of Trial ($F[1,38] = 28.39$, $p < .001$, $\eta^2 = .23$). Participants interacting face-to-face were significantly more likely to offer splits in their own favor (M = 64.8, SE = 1.67) than participants inside of VR (M = 55.48, SE = 1.15). The interaction was not significant ($F[1,38] = .06$, $p = .81$, $\eta^2 < .001$).

And finally, to analyze the dichotomous measures of acceptance of the unfair offers, I ran two logistic regressions using the acceptance rates as the dependent variables and height as the independent variable. In the trial within VR, the regression model was not significant $\chi^2(1, N = 40) = .90$, $p = .34$, *Cox and Snell R-Square* = .02. Height was not a significant predictor of acceptance $B = -.61$, $p = .34$. In the trial outside of VR, the regression model also wasn't significant, $\chi^2(1, N = 40) = 1.62$, $p = .20$, *Cox and Snell R-Square* = .04. Height was also not a significant predictor of acceptance $B = -.82$, $p = .21$.

In sum, the effect of height was significant on the first split but not the second split or the final acceptance. This pattern—finding the largest effect only on the first split—is actually consistent with most research which utilizes this money splitting task (Bolton, 1991). More importantly, the effect of height persisted on the first task on trials in which participants negotiated face-to-face. Thus, the changes that occur to one’s conception of social identity due to being embodied within a specific avatar persist even after he or she leaves digital space.

Limitations and Areas for Future Study

The pilot studies provide several important foundations. First of all, the results strongly suggest that alterations to our digital self-representations can cause a significant difference in how we behave. In other words, there is support for the presence of the Proteus Effect. Furthermore, these effects generalize to natural virtual communities and even impact consequent face-to-face interactions. The studies also provide us with pre-tested stimulus materials (i.e., the pre-tested attractive and unattractive faces, specific height differences which are not explicitly detected) that produce measurable differences. And finally, the studies provide a set of dependent measures that appear to be sensitive to the manipulations of self-representation.

There are however several shortcomings to the pilot studies. Most importantly, the findings do not clarify the underlying psychological processes that drive the effect. Earlier I suggested that self-perception - inferring appropriate behavior as a third-party observer - leads to the observable differences in behavior, but the pilot studies do not directly provide evidence for this. To help clarify the underlying process driving the observed changes, two studies will be conducted.

In the first study, a potential alternative explanation for the observed effects is isolated and compared with the Proteus Effect. In social psychology, certain studies have shown that brief exposure to visual stimulus can lead to behavioral changes (Bargh, Chen, & Burrows, 1996; Dijksterhuis & van Knippenberg, 1998). Thus, it is possible that the findings from my pilot studies were simply a product of this psychological process. In the second study, I use the underlying assumptions of the Proteus Effect to understand what would occur when participants are placed in bodies that fall outside the range of normal human variation (e.g., a 8.6 m tall body).

DISSERTATION STUDY ONE: RULING OUT BEHAVIORAL ASSIMILATION

Introduction

Our bodies are wired with many automatic processes that can operate without our direct, conscious control. These include many bodily functions and reflexes, such as breathing, or the grasping and startle reflexes that can be observed in newborn infants. Automatic processes can also be found in human perception, such as our natural tendency to group and complete visual information that may be fragmented as demonstrated by Gestalt psychology. More relevant to the current work are findings that even cognitive processes can occur with an automaticity that is beyond conscious control. In the following section, I will describe several convergent lines of research in cognitive and social psychology detailing how presentation of words or trait terms can affect subsequent evaluations of stimuli as well as interpersonal behavior. This line of work provides a potential alternative explanation for the observed findings in my pilot studies into the Proteus Effect.

Cognitive Interference and Spreading Activation

One of the earliest studies that demonstrated automaticity in cognitive processes is the Stroop Effect (Stroop, 1935). When participants were asked to name the color a word was written in (always a non-congruent color word), the automaticity of the reading response increased task completion time due to cognitive interference. Stroop suggested that the differing association strengths between reading and color-naming were due to the disparity between training and frequency in reading and color-naming in everyday life. In other words, Stroop argued that we've been trained in life to always read words (instead of contemplating the shape of fonts for example), whereas colors can be appreciated or perceived without naming them. Thus, it is hard to suppress the automatic reading response even when it isn't needed.

This finding that certain cognitive processes can occur automatically and interfere with other cognitive processes generated a great deal of research into exploring memory retrieval and perceptual readiness in both cognitive and social psychology. In cognitive psychology, foundational work by Quillian (1962; 1967) provided a connectionist framework for showing how human semantic structure and processing could be built into a machine. This *spreading-activation model* proposed that concepts were represented as nodes, and nodes in turn were connected via relational links into a larger system of knowledge. A node becomes activated when it is used, such as being processed from a visual presentation of a certain stimuli. The term *semantic priming* is used to refer to the initial activation that produces a cascade via spreading-activation. And finally, the more closely-related two nodes are (i.e., *semantic relatedness*), the more energy that is carried through the link when either is activated, and thus the more likely the other node becomes activated as well. For example, seeing the word "fire" is more likely to activate "hot" rather than "nurse".

This model helped explain many effects found in semantic priming studies (Collins & Loftus, 1975). One important finding was that once nodes were activated via priming, they were also more likely to be retrieved in subsequent memory searches. Thus, in a task where participants had to decide as quickly as possible whether a sequence of letters was a word or not (i.e., a lexical decision task), participants were quicker to verify that "butter" was a word if "bread" preceded its presentation than if "nurse" preceded it instead (Meyer & Schvaneveldt, 1971). More recent research has shown that priming can aid word completion and recall even if the primes are flashed on a computer screen at a speed faster than can be visually processed and which participants do not consciously recall seeing (Jacoby & Witherspoon, 1982). This line of research has provided evidence that *implicit memory* exists and operates independently of explicit memory (Neely, 1977; Posner & Snyder, 1975a, 1975b; Tulving, Schacter, & Stark, 1982).

Priming and Social Perception

Research in social psychology drew upon and expanded on the research on semantic priming, but specifically with regards to how it impacted person perception. Early research showing that short verbal descriptions (i.e., warm vs. cold) could significantly impact subsequent evaluations of a person (Kelley, 1950) generated interest in exploring how priming could influence social perception. One potential confound in Kelley's study however was that participants might have been conforming to the judgment made by an authority figure rather than actually being primed by the trait label. In other words, participants may have relied on the experimenter's verbal assessment of the target person because they felt he provided a reliable description, a process more akin to conformity studies than to semantic priming.

One of the earliest studies to isolate the effect of priming on social perception (Higgins, Rholes, & Jones, 1977) manipulated the presentation of the terms "reckless" and "adventurous" in a lexical task (e.g., unscrambling or completing words) and then, in an ostensibly unrelated task, elicited participants' impression of a person who was planning to sail across the Atlantic in a sailboat. It was found that participants who had been exposed to the term reckless had a more negative impression of the person than participants exposed to the term adventurous. Thus, it is possible to create short-term associations that influence how people perceive others in social situations. The effect of priming on social perception was replicated in a series of studies by Srull and Wyer (1979; 1980) where they also found that the more a participant was primed with a trait in the first task, the more extreme their trait ratings were of the target person in the latter task.

One highly-regarded theoretical explanation for this effect is known as the *excitation-transfer model* (Higgins & King, 1981; Wyer & Carlston, 1979) and is similar to the spreading-activation model in many ways. The excitation-transfer model assumes that our likelihood of using a particular construct (e.g., hostile, intelligent, etc.) in a social judgment depends on the accessibility of that construct. Constructs can be chronically-accessible due to frequent individual exposure to certain kinds of social behaviors, such as kindness or aggression (Kelly, 1955); they can also be made temporarily more accessible due to recent use (Bruner, 1951). Specifically, Bruner argued that social information is typically ambiguous and accessible constructs may be used to interpret behaviors even though more relevant, but less accessible, constructs are available. Other researchers (Higgins & King, 1981; Higgins, King, & Mavin, 1982; Wyer & Carlston, 1979) have integrated these two kinds of construct accessibility by suggesting a general threshold model of construct activation. They argued that a construct's accessibility can be increased due to recent use or current goals and that over time gradually returns to a baseline value. Once a construct reaches a threshold level of accessibility, the likelihood that it will be used in a social judgment increases. And indeed, it has been shown that both chronic and temporarily accessible constructs influence social perception (Higgins et al., 1982). Thus, priming influences social perception by making certain constructs more accessible. To be more precise, processing a prime increases the excitation level of specific constructs beyond the threshold level.

While early studies in priming had users consciously process the semantic primes (e.g., reading or hearing prime-related words), other research has shown that participants don't need to have consciously perceived the prime for priming to be effective. Mirroring the cognitive psychology work mentioned earlier by Jacoby and Witherspoon (1982), it has also been shown that priming can affect social perception even if the primes are presented subliminally, outside of the participant's awareness (Bargh & Pietromonaco, 1982). To accomplish the subliminal presentation of primes, researchers implemented a perceptual vigilance task where participants had to identify the location of briefly-flashed stimuli on a computer screen. The flashed stimuli were actually words presented at a speed that precludes conscious identification. Bargh and Pietromonaco replicated Srull and Wyer's (1979) study but with subliminal primes, and found the same effect. Thus, even when participants were unaware of the presence of the primes, let alone the meaning of the primes, primed traits still had an effect on subsequent evaluations of another person.

Behavioral Assimilation

While a great deal of research has shown that priming can affect how we perceive and evaluate other people, perhaps the most provocative extension was research showing that priming can in fact change how a person behaves and treats other people. In other words, priming can affect not only social perception but also a person’s behavior, a phenomenon known as *behavioral assimilation* (Bargh et al., 1996; Dijksterhuis & van Knippenberg, 1998). As Bargh and his colleagues have argued, the mere act of thinking about a behavior increases a person’s tendency to engage in that behavior (Ansfield & Wegner, 1996; James, 1890; Wegner, 1994). In fact, an earlier study hinted at the behavioral consequences of automatic social perception (Carver, Ganellen, Froming, & Chambers, 1983). In that study, participants subliminally primed with the concept of hostility administered more shocks to a confederate learner than participants given neutral primes. Carver and his colleagues accounted for this effect by suggesting that perceiving hostility increased the likelihood of behaving in a more hostile manner.

In a series of experiments, Bargh and his colleagues (1996) demonstrated the varying contexts under which priming could lead to behavioral changes. For example, they found that participants primed with elderly-related traits walked slower than participants primed with neutral words. They also showed how stereotypical traits of a group could be primed, via spreading-activation, simply by presenting exemplars of the group. Thus, they hypothesized that flashing photographs of African-Americans would prime hostility. In their study, they found that participants primed with photographs of African-Americans behaved in a more hostile manner when a scripted computer error occurred than participants primed with photographs of Caucasians.

Another study (Dijksterhuis & van Knippenberg, 1998) demonstrated the effect of exemplar-priming on cognitive performance. Participants asked to imagine a typical professor (as a proxy for priming “intelligence”) performed better on a general knowledge task than participants asked to imagine a typical secretary. In a second experiment, participants asked to imagine a typical soccer hooligan (for priming "stupid") performed worse on a knowledge task than participants who received no prime. Dijksterhuis & van Knippenberg argued that the observed difference in cognitive performance was likely the outcome of changes in behavioral repertoire rather than in actual intelligence. In other words, participants primed with a typical professor may be more inclined to concentrate harder or to perform more thorough searches. In short, research in behavioral assimilation show that trait-priming can lead to differences in behavior and cognitive performance in addition to attitudes and social perception.

The work on behavioral assimilation provides an alternative explanation for my pilot study findings. For example, participants in the attractiveness study may simply have been primed with friendliness by processing an attractive face in the mirror in the attractive condition. In the same way that participants primed with "professor" performed better on a knowledge test, participants primed with "attractiveness" in my study may have behaved congruently with constructs activated via the perception of an attractive face, and vice versa for the unattractiveness condition.

To tease apart the unique contribution of the Proteus Effect from that of behavioral assimilation, I conducted an experimental study where the same visual stimulus was presented to two groups of participants, but where one group saw their own mirror reflection while the other group saw a virtual video playback on a "large screen" (i.e., the same framed area as the "mirror"). Similar to Pilot Study One, I manipulated the avatar's attractiveness as well. I hypothesized that behavioral changes from the Proteus Effect (i.e., mirror condition) would be significantly larger than that from behavioral assimilation (i.e., playback condition).

Method

Design

Study One was a replication of Pilot Study One with an additional level of manipulation. In a between-subjects design, participants were randomly assigned to one of two presentation conditions. In the *mirror* condition, participants were given an avatar and their avatar reflected their head movements in a mirror (as in Pilot Study One). In the *playback* condition, participants watched the recording of the previous participant in the mirror condition on a "large screen" (the same framed viewing area as in the mirror condition). In other words, participants in the playback condition saw the same visual stimulus as a participant in the mirror condition. Thus, differences between these two conditions could isolate the unique contribution of the Proteus Effect - the degree to which *being in* an attractive avatar changes one's behavior above the amount provided by priming alone.

As in Pilot Study One, I also varied the attractiveness of the avatar (or digital person in the playback condition). Participants were assigned to be in either the *attractive* or *unattractive* condition.

Participants

73 undergraduate students (37 female, 36 male) participated in the study for either course credit or $5.

Materials and Apparatus

The virtual faces and apparatus used in this study were the same as those described in Pilot Study One. In addition, I implemented a mock online dating website to explore another set of measures outside of the virtual environment after participants had been immersed in virtual reality.

Mock Dating Website Pretest. Photographs of 10 male and 17 female undergraduates from another university were rated by 13 undergraduates selected from a different pool than those in the study itself. All the photographs were frontal portraits taken under similar lighting conditions of each individual smiling and standing in front of a blue screen. These undergraduates were asked to rate the attractiveness of each photographed individual on a fully-labeled scale from 1 (extremely unattractive) to 7 (extremely attractive). Then for each gender, nine photographs were chosen that spanned as much of the attractiveness scale as evenly as possible. For the set of male photographs, the range was from 1.77 (SD = 0.73) to 6.08 (SD = 0.76), with a resulting mean of 3.84 and a standard deviation of 1.32. For the set of female photographs, the range was from 2.38 (SD = 0.87) to 6.23 (SD = 0.60), with a resulting mean of 4.06 and a standard deviation of 1.30.

Procedure

When participants arrived, they were told that they would be participating in two unrelated studies. They were then prepared for the first study which they were told would involve social interaction in a virtual environment. The procedure of this first study was very similar to Pilot Study One. In the mirror condition, after participants were immersed in the virtual lab setting, they were told the following:

> "In a moment, I will turn you around and behind you is a mirror. The way you appear in the mirror is how other people in this virtual environment will see you."

In the playback condition, participants were instead told:

> "In a moment, I will turn you around and behind you is a large TV screen. On that screen we will be playing back a recording of someone who was in this room some time ago."

To ensure that participants in both conditions had the same amount of experience with walking and moving their head prior to interacting with the confederate, participants in the playback condition were also asked to approach the "large screen" and tilt their heads as in the mirror condition (i.e., *the movement sequence*). Of course, in the playback condition, the digital person in the "large screen" did not reflect their head movements and this helped reinforce the fact that the person in the screen was not a reflection of them.

To facilitate the synchronization of the visual stimulus in both conditions, an automated algorithm was used to start the recording in the mirror condition when the participant had turned the first 90 degrees (of the 180) towards the mirror. The recording then ended when the participant had turned the first 90 degrees back around. Then, in the playback conditions, this recorded segment was played back when the participant had turned the first 90 degrees towards the mirror. Since it took roughly the same amount of time for all participants to complete the movement sequence, participants were asked in both conditions to turn back around as soon as they completed the sequence. The remainder of the virtual reality portion of the study is the same as that described in Pilot Study One.

After participants were taken out of the virtual environment, they were told that the second study involved studying interpersonal compatibility in romantic contexts. Participants were then asked to stand in front of a large blue screen and smile while the researcher took a photograph of them. This was done to increase the plausibility of the task when the participants later saw the photographs in the mock website (also taken in front of blue screens). After the photograph, the researcher explained the task in more detail to the participant:

> We’re studying how compatibility works on online dating websites such as Match.com. In this study, we’re interested in how well people can identify compatible partners from photographs alone. We’ll have you fill out some information about yourself, as if you were creating a profile on an online dating site. Then we’ll show you a panel of photographs from students from other colleges who have also created profiles. We’re interested in seeing whether people who have similar personalities are more likely to select each other or not.

Participants were also told that this was only a study of interpersonal compatibility and not an actual dating service, and thus, that they would not actually be meeting the people whose photographs they would be presented with.

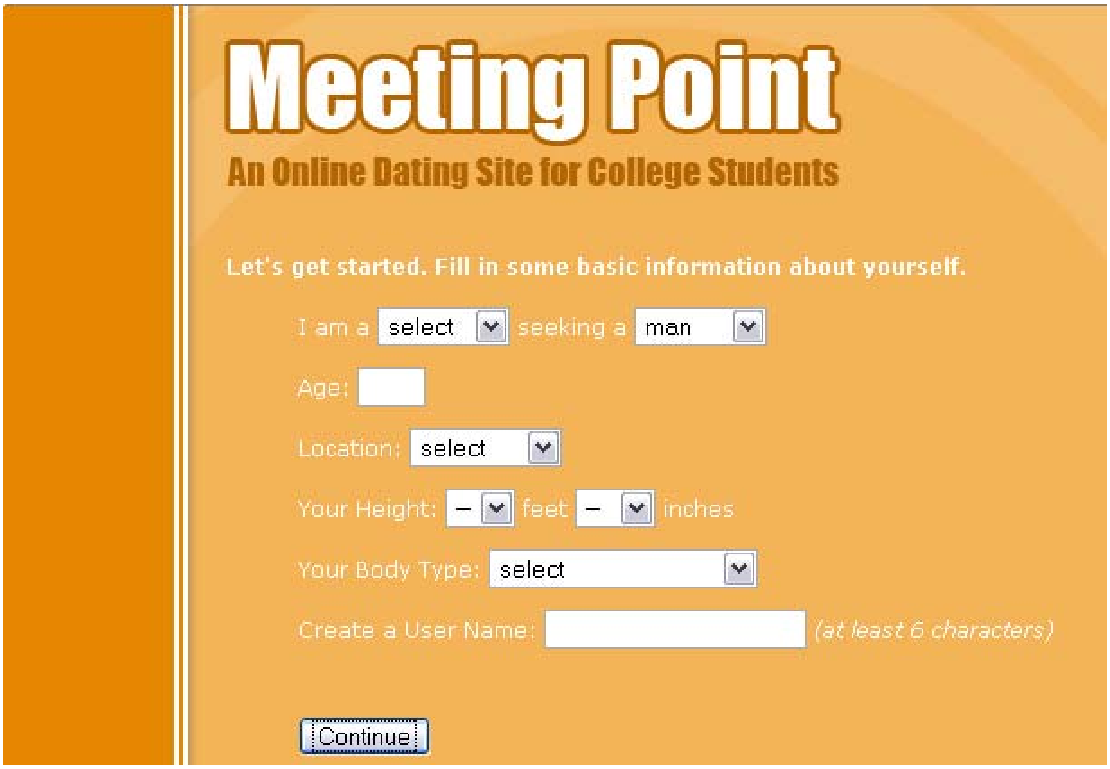

Figure 9. Profile Entry Page of Mock Dating Website.

After completing routine profile information about themselves (see Figure 9), participants were presented with nine photographs which they were told were chosen from a database based on their profile information. For each gender, the same nine photographs (as described earlier) were always shown (see Figure 10). The specific gender set was determined by a question in the profile regarding whether the participant was interested in men or women. Written instructions on the mock website asked the participants to:

> Pick the two people who you are most interested in and who you think would most likely be interested in you. We'll include your photo in their photo panel to see if there's a match.

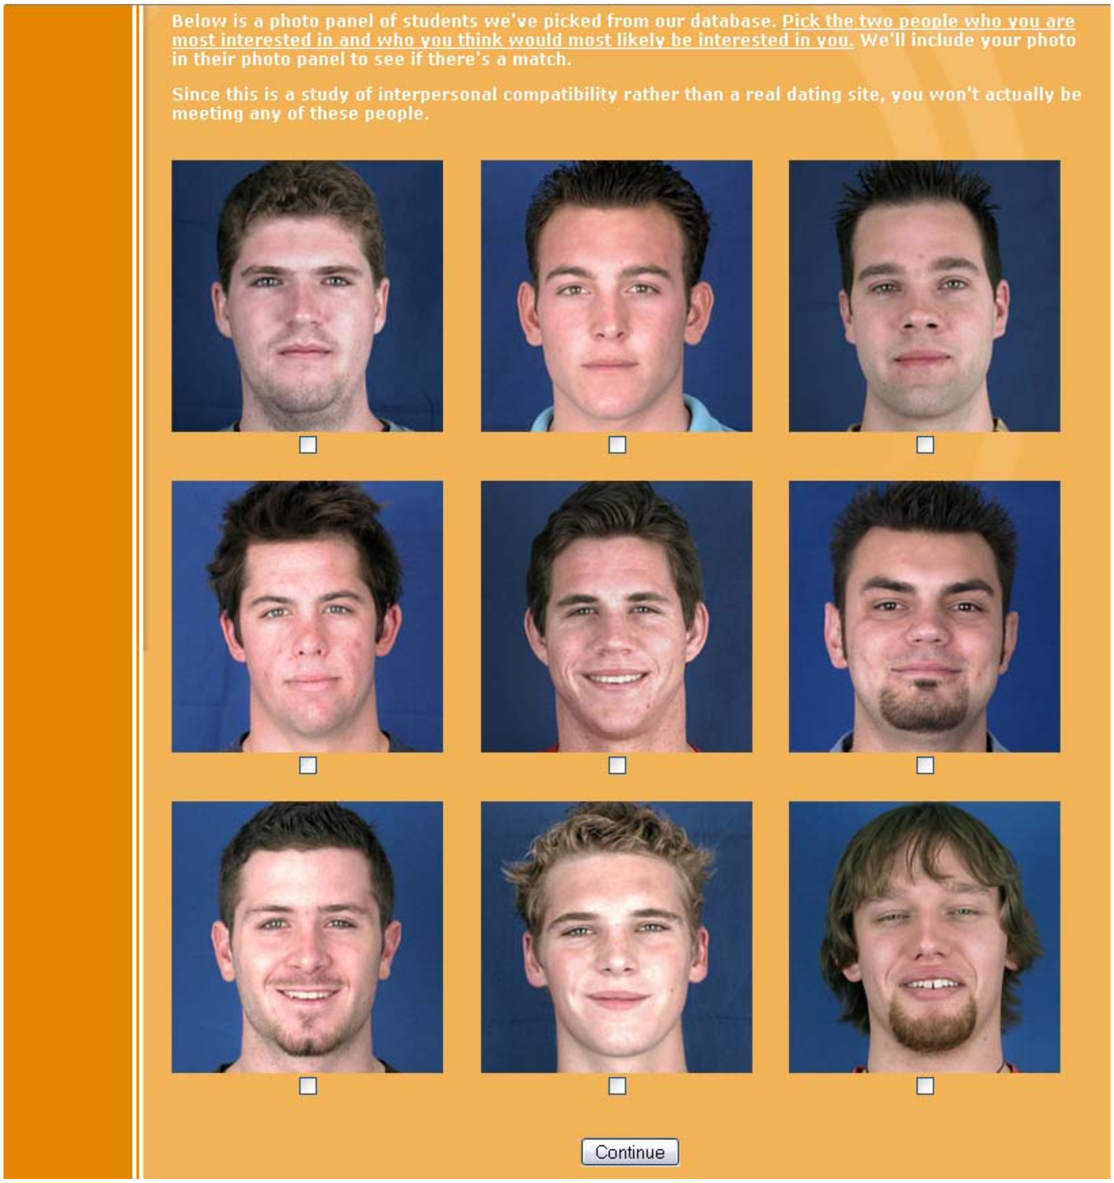

Figure 10. Partner Selection Page of Mock Dating Website.

Measures

Partner Choice. For each participant, I summed the attractiveness scores of the two individuals they chose from the photographs presented to them at the end of the mock dating website. In line my hypothesis, I predicted that there would be an interaction between the attractiveness and presentation conditions. Specifically, the difference between the attractiveness conditions in the mirror condition would be larger than the difference in the playback condition, such that subjects who saw themselves as attractive in the mirror would pick the most attractive choices.

Reported Height Difference. While completing their profiles for the dating website task, participants were asked to self-report their own height. Unbeknownst to them, I also measured their actual height using the tracking equipment on the head-mounted display. This allowed us to calculate a potential difference between their reported height and their real height. Previous research has shown that most people lie in their own favor on online dating profiles (Hancock, Toma, & Ellison, 2007). Thus, I predicted that participants would inflate their own height least in the attractive mirror condition where participants' increased self-confidence would lower their need to lie about their own height. In other words, a comparison between the attractive mirror condition and the other three conditions should be significant.

Participant Attractiveness. Each participant's photograph was rated by 12 individuals drawn from a separate population than from the participants in the study. Each participant was rated on a fully-labeled 7-point scale, ranging from 1 (extremely unattractive) to 7 (extremely attractive). Participant attractiveness was used as a covariate in my analyses.

And as in Pilot Study One, I collected information on interpersonal distance and self disclosure. With both these measures, I again predicted an interaction between attractiveness and presentation conditions similar to pattern described in the partner choice measure, such that participants in the attractive mirror condition would go closest to the confederate and reveal the most.

Results and Discussion

Partner Choice

I conducted an ANOVA with the attractiveness (attractive / unattractive) and presentation conditions (mirror / playback) as the independent variables, participant attractiveness as a covariate, and partner choice as the dependent variable. Participant attractiveness was not a significant covariate ($F[1, 68] = .02$, $p = .89$, $\eta^2 < .001$). The effect of attractiveness condition was not significant ($F[1, 68] = 1.81$, $p = .18$, $\eta^2 = .02$). The effect of presentation condition was also not significant ($F[1, 68] = .04$, $p = .84$, $\eta^2 < .001$). The interact was significant ($F[1, 68] = 5.31$, $p = .02$, $\eta^2 = .07$), as depicted in Figure 11. A post-hoc comparison using Fisher's LSD Test showed that in the mirror condition, participants in the attractive condition had a higher partner choice score ($M = 10.47$, $SE = .29$) than participants in the unattractive condition ($M = 9.42$, $SE = .29$), $p = .01$. In the playback condition, participants in the attractive condition ($M = 9.75$, $SE = .29$) did not have significantly higher partner choice scores than participants in the unattractive condition ($M = 10.03$, $SE = .29$), $p = .50$.

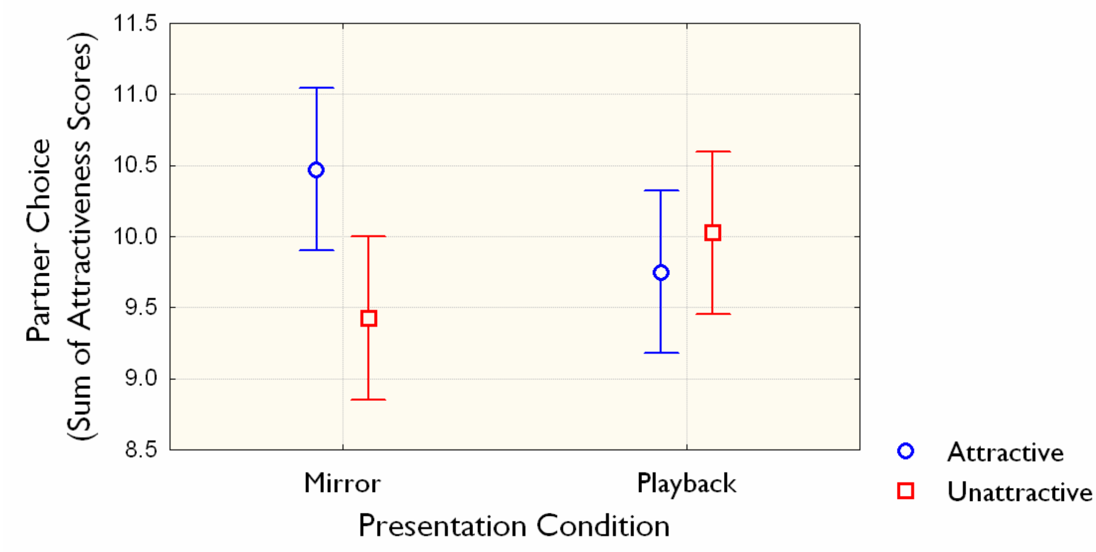

Figure 11. Partner Choice by Attractiveness and Presentation Conditions.

Reported Height Difference

I conducted an ANOVA with the attractiveness (attractive / unattractive) and presentation conditions (mirror / playback) as the independent variables, participant attractiveness as a covariate, and reported height difference as the dependent variable. Participant attractiveness was not a significant covariate ($F[1, 68] = .14$, $p = .71$, $\eta^2 = .002$). The effect of attractiveness condition was not significant ($F[1, 68] = .46$, $p = .50$, $\eta^2 = .007$). The effect of presentation condition was also not significant ($F[1, 68] = .55$, $p = .46$, $\eta^2 = .008$). The interaction was significant ($F[1, 68] = 4.26$, $p = .04$, $\eta^2 = .06$), as shown in Figure 12. A post-hoc comparison with Fisher's LSD test showed that in the mirror condition, participants in the unattractive condition were significantly more likely to increase their reported height ($M = 1.17$, $SE = .33$) than participants in the attractive condition ($M = .17$, $SE = .39$), $p = .05$. In the playback condition, participants in the attractive ($M = 1.20$, $SE = .37$) and unattractive conditions (M = .69, SE = .37) did not have significantly different reported height differences, p = .34.

Our planned comparison between the attractive mirror condition and the other three conditions was significant (t[68] = -2.04, p = .05).

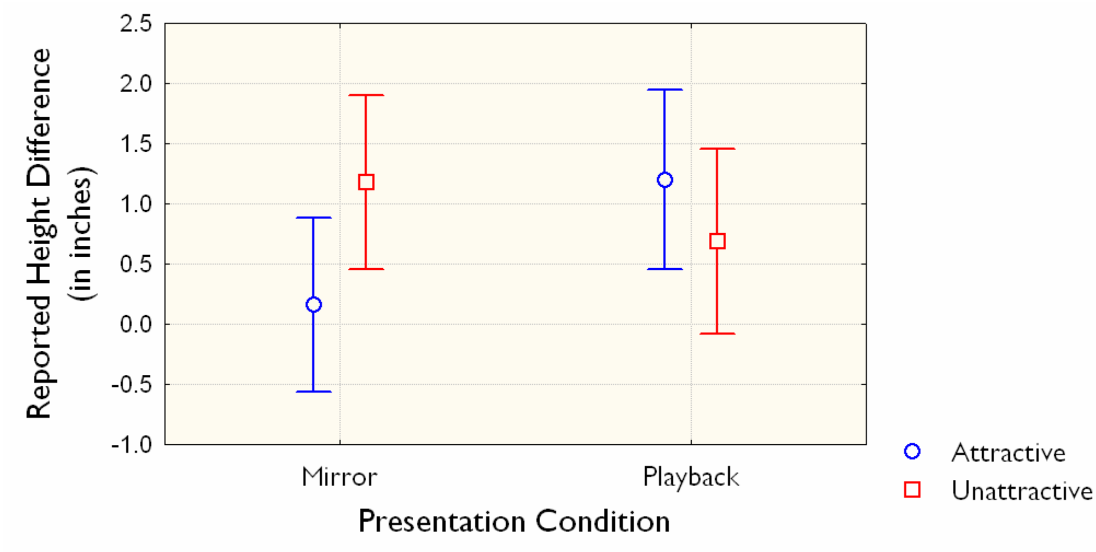

Figure 12. Reported Height Difference by Attractiveness and Presentation Conditions.

Interpersonal Distance

I conducted an ANOVA with the attractiveness (attractive / unattractive) and presentation conditions (mirror / playback) as the independent variables, participant attractiveness as a covariate, and distance from the confederate as the dependent variable. Participant attractiveness was not a significant covariate (F[1, 68] = .27, p = .60, η^2 < .001). The main effect of the attractiveness condition was also not significant (F[1, 68] = 1.34, p = .25, η^2 = .02). The main effect of presentation condition was also not significant (F[1,68] = .05, p = .83, η^2 < .001). The interaction was also not significant (F[1, 68] = 1.39, p = .24, η^2 = .02). While the interaction was not significant, a comparison of the means and confidence intervals showed that in the mirror condition, there was a trend where participants in the attractive condition (*M* = 1.94, *SE* = .28) walked closer to the confederate than in the unattractive condition (*M* = 2.53, *SE* = .24), p = .10 using Fisher's LSD Test, see Figure 13.

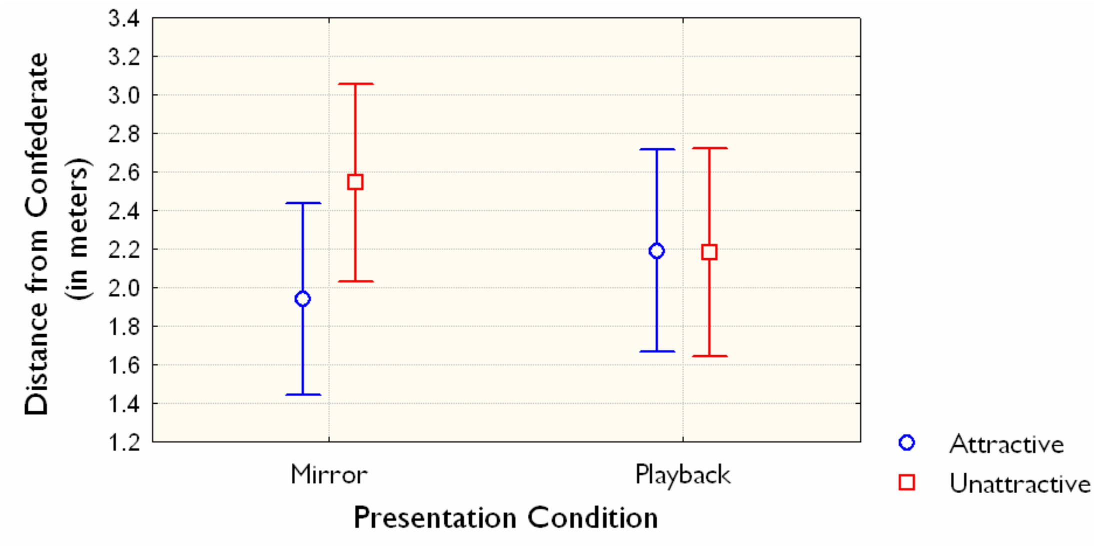

Figure 13. Distance from Confederate by Attractiveness and Presentation Conditions.

Self-Disclosure

I conducted an ANOVA with the attractiveness (attractive / unattractive) and presentation conditions (mirror / playback) as the independent variables, participant attractiveness as a covariate, and self-disclosure as the dependent variable. Participant attractiveness was not significant as a covariate ($F[1, 68] = .04$, $p = .84$, $\eta^2 < .001$). The effect of attractiveness condition was also not significant ($F[1, 68] = .31$, $p = .58$, $\eta^2 = .004$). The effect of presentation condition was also not significant ($F[1, 68] = 1.90$, $p = .17$, $\eta^2 = .03$). And the interaction was also not significant ($F[1, 68] = .05$, $p = .83$, $\eta^2 < .001$).

The interaction effects in partner choice and reported height difference provide support for my hypothesis that differences in the mirror condition would be larger than those in the playback condition. While the results in interpersonal distance and self-disclosure were less clear, the trend of the interpersonal distance findings was in the predicted direction. Together, these findings suggest that the Proteus Effect is not simply the result of priming.

More importantly, these findings suggest that embodiment plays a large role in the Proteus Effect. Perceiving the exact same visual stimulus in and of itself did not produce significant behavioral changes in the study. In other words, it is some combination of believing that you are really in a different body or the sense of agency or interactivity in a new body that leads to the effects observed. The latter notion also resonates with research that illustrates the importance of interactivity when creating engaging experiences (Liu & Shrum, 2003; Reeves & Nass, 1996; Schwartz, Blair, Biswas, Leelawong, & Davis, in press; Steuer, 1993). Thus, embodiment and digital self-representation are important factors leading to the observed behavioral changes in the Pilot Studies and the current study. Moreover, this implies that digital embodiment is a unique lever for behavioral change. This in turn helps support my theoretical framework that self-perception plays a large role in the Proteus Effect.

DISSERTATION STUDY TWO: CATEGORIZATION OF IMPLAUSIBLE BODIES

Introduction

One aspect of digital representation that has been mentioned repeatedly throughout the dissertation is the ease with which dramatic transformations can be made. The studies I have conducted related to The Proteus Effect, however, have so far involved only normal variations in human embodiment - slightly taller or more attractive avatars. On the other hand, most online environments allow users to take on avatars that are not within the range of normal human variation, such as diminutive gnomes or over-sized ogres. One provocative and relevant research question that stems from these “extreme” representations is whether and how the Proteus Effect may operate when a user is given an avatar that is outside the scope of normal human variation. Given that the causal explanation for the Proteus Effect is based on the user conforming to expected behaviors (or stereotypes) of the avatar representation, it makes sense to examine possible outcomes related to how people infer stereotypes from novel categories. Below, I will provide a foundation by briefly summarizing existing theories in categorization and then explore three different potential outcomes based on the literature in category-based induction, stereotype formation, and bias correction.

Models of Categorization

One fundamental problem our minds have to solve is how to categorize individual objects around us into meaningful groups. For example, given the natural variation in dogs and that no two dogs are exactly alike, how do we come to have a mental representation of dogs and know whether an animal is a dog or not? The classical explanation that dates back to Aristotle and Plato and also espoused by early cognitive psychologists (Katz & Postal, 1964) was that categories are based on rules that are singly necessary and jointly sufficient to determine whether an object belongs in a certain category. For example, a square is a two-dimensional shape with four sides of the same length where every angle is a right angle. This model of categorization suggests that strict boundaries exist in between all categories.

The classical model, however, soon became problematic as more and more studies showed that most natural categories do not have well-defined rules separating them from other categories (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Moreover, certain members are perceived as better exemplars of a particular category. For example, in the study by Rosch and colleagues, certain objects are judged to be more typical members of their category (i.e., a robin is a more typical member of the bird category than ostriches). On the other hand, these typical members may have salient attributes that, while characteristic for a category, are not category-defining (e.g., most birds can fly, but penguins and ostriches are birds even if they don't fly, (Smith, Shoben, & Rips, 1974).

Two main models have emerged for how categories are acquired and represented. Prototype-based models (Homa & Vosburgh, 1976; Posner & Keele, 1968; Reed, 1972) suggest that based on exposure to members of a category, people extract the common features of the members to form an abstract central tendency or prototype. This prototype is then used to judge whether other items belong in the category. By contrast, exemplar-based models (Hitzman, 1986; Medin & Schaffer, 1978; Nosofsky, 1986) suggest that categories consist of the exemplars (i.e., members of a category) themselves rather an abstracted set of features, and that categorization of new objects is based on comparison with specific exemplars within a category. The key difference between the two models is how the category is represented in the mind. According to prototype-based models, a category is an abstracted set of common features. On the other hand, exemplar-based models would say that a category is the set of learned exemplars. While others have suggested a mixed-model approach (Busemeyer, Dewey, & Medin, 1984; Elio & Anderson, 1981; Homa, Sterling, & Trepel, 1981), the general debate between the two types of models has continued (Minda & Smith, 2002; Zaki, Nosofky, Stanton, & Cohen, 2003).

Category-Based Induction

Let's consider an avatar that is outside the realm of plausibility in the normal range of human variation - a 0.34 m tall avatar (5 times shorter than the average height of 1.72 m). Given what we know about how users behave in a slightly shorter avatar, one intuitive hypothesis is that a 0.34 m person is just a very short "short person". And thus, a user in a very short avatar would incorporate an exaggerated version of the "short" stereotype - very shy, very passive, yields to demands. Ironically, this type of reasoning would fall into the classical model of categorization because we'd be arguing that a 0.34 m person falls into the category of "short people" due to good fit with a set of abstract rules for what constitutes a short person (i.e., shortness). In fact, both the prototype and exemplar models would argue that a 0.34 m person is a highly atypical member of the "short people" category because we don't meet 0.34 m people on a regular basis at all.

A well-described model of category-based induction (Osherson, Smith, Wilkie, Lopez, & Shafir, 1990) provides both a good explanation for the theoretical differences we just mentioned as well as a good framework for understanding how users might react. This model provides a way of understanding how good an argument is (going from a premise to a conclusion) in terms of *argument strength*. One example of such an argument is:

Premise A: Robins have sesamoid bones.
Conclusion: All birds have sesamoid bones.

Compare that to the following argument:

Premise B: Ostriches have sesamoid bones.
Conclusion: All birds have sesamoid bones.

In this induction model, the strength of an argument is based on the typicality of members in the relevant category. Thus, because robins are more typical birds than ostriches, the first argument is perceived as being stronger. This observed difference in argument strength is referred to as *premise typicality*. Our current investigation is the inverse of this observed argument strength difference. Since “0.34 m person” is an atypical member of “short people”, attributes from the category of “short people” do not transfer as strongly as compared with a more typical member of “short people”, such as a “1.50 m person”. Thus, according to this model, a 0.34 m person carries a weakened version of the stereotypes associated with short people because of its atypicality. Thus, one unintuitive hypothesis we might draw from this model is that participants in a 0.34 m tall avatar would in fact be less shy (i.e., more aggressive) than participants in a 1.50 m tall avatar.

Stereotype Formation

Understanding how stereotypes of social groups form is crucial to understanding what stereotypical traits user may attribute to novel social groups in digital environments. As we will see, the literature in stereotype formation has drawn on much of the literature in categorization that I reviewed earlier to understand a specific phenomenon in the domain of social psychology. Traditionally, social psychologists have believed that stereotypes of social groups were based on abstracted models that summarized the attributes of group members (Lippman, 1922). These generalizations have sometimes been referred to as schemata (Taylor & Crocker, 1981), prototypes (Brewer, Dull, & Lui, 1981), or Bayesian base rates (McCauley & Stitt, 1978), among other terms. Under this view, exposure to members of a social group produces abstracted representations over time that aggregates salient features (Shriffin & Schneider, 1977). More recent work has elaborated on this process by showing that an association between a social category and an attribute is formed when the association has a high central tendency and low variability - i.e., they co-occur consistently (Dijksterhuis & van Knippenberg, 1999).

Similar to the literature in general categorization, the abstraction-based model has been challenged by researchers who argue that exemplars play a large role in stereotype representation. This challenge arose due to several observed phenomena. For example, it has been observed that extensive sub-typing occurs with most social groups (Devine & Baker, 1991); in addition to a general stereotype of “Blacks”, people seem to have coherent subtypes for “Black athletes”, “welfare Blacks”, “streetwise Blacks”, and so forth. Other research has also found stereotypes to be more variable, granular, and context-dependent than would be suggested by a pure abstraction-based model (Hamilton & Sherman, 1994; Smith, 1990). The exemplar-based alternative (Smith & Zarate, 1992) argues that stereotypes are generated dynamically by the activation of particular exemplars stored in memory and aggregating their features. One interesting feature of this model is that the activated exemplars may not typically be considered to be from the same social group. For example, if a Chinese person reminded you of Tom Hanks, stored traits about the Caucasian actor may be used (in combination with other exemplars) to form the stereotype even though the Chinese person is clearly not Caucasian. Thus, while abstraction-based models propose that stereotypes are stored generalizations, exemplar-based models suggest that stereotypes are stimulus-specific aggregations created dynamically.

The exemplar-based model is not without its own theoretical weaknesses. For example, if exemplars were not joined together by some sort of categorization, people would be unable to activate the appropriate exemplars and make categorical judgments of social groups (i.e., “What do you think of lawyers?”), which we know they can (Medin & Wattenmaker, 1987). This suggests that abstracted categories of social groups do exist on some level in the mind. And a pure-exemplar model has also been argued to be too cognitively-inefficient (Hamilton & Sherman, 1994).

The weaknesses of both pure abstraction- and exemplar-based models has led some researchers to consider mixed models of stereotypes (Sherman, 1996), analogous to the adoption of mixed models of construct representation in cognitive psychology (Busemeyer et al., 1984; Elio & Anderson, 1981; Homa et al., 1981). Sherman tested a mixed model of stereotypes and found that at low levels of category member exposure and experience, people based their judgments on the activation of exemplars because no useful abstract knowledge had yet formed. At higher levels of exposure, an abstract representation formed, which then served as the referent for subsequent judgments.

The literature on stereotype formation provides several insights into my research question. At low levels of experience, individuals resort to exemplar-based stereotypes and may not produce any consistent stereotypes. In other words, while people will generate stereotypes based on personal exemplars, these stereotypes will not be consistent across people. This suggests that if a truly novel avatar were presented, no group-level differences would emerge. On the other hand, just because an avatar isn't in the realm of normal human variation doesn't mean it isn't heavily stereotyped. Television and movies have provided us with consistent and repeated exposure to a large set of non-existent creatures, such as vampires. Thus, if people have a great deal of previous exposure to their given avatar, we may expect consistent group-level differences to emerge, but if people have no or few exposures, then we would not expect group-level differences to emerge. With regards to my research question, this suggests that unless there are widely-adopted stereotypes of 0.34 m tall people, participants would fall back on individual exemplar-based stereotypes and there may be no consistent behavioral changes found.

Bias Correction

In retrospect, one unobvious benefit of using avatars that are within the normal range of human variation is that they do not make the manipulation salient and available for scrutiny and consideration. On the other hand, conspicuously non-human or extreme representations, such as being given a vampire avatar, invite scrutiny. A similar point was made in the priming literature with regards to the implementation of subliminal primes. As Bargh argued, “preconscious influences would play a stronger than usual role in subsequent behavior toward the target person, as the perceiver would not be aware of the interpretive bias and so could not correct for it” (pg. 230, Bargh et al., 1996). And indeed, several studies have shown that participants react very differently when particular contextual cues become salient.

In an interesting study, Higgins and King (1981) manipulated the salience of gender by changing the gender composition of participant groups. It was assumed that being the only member of his or her gender in a group would make the notion of gender more salient, which in turn would make the concept of gender more accessible for cognitive processing. Once the concept of gender became accessible for conscious processing, participants should be more likely to monitor their thoughts for traditional and stereotypical views of gender and consciously replace them with more “modern” views of gender equality. In their study, participants learned about a person with a list of traits. Results demonstrated that when gender was made salient, participants remembered more of the information that was less traditionally stereotypical than when gender was not made salient. In other words, when gender was not salient, information processing relied passively on ingrained and traditional views of gender, and only when gender was made salient and conscious effort was applied did participants behave differently.

Devine (1989) has further elaborated on the automatic and controlled components of stereotypes. In a series of studies, Devine first determined whether a participant was low- or high-prejudiced against Blacks using the Modern Racism Scale. It was found that in a task that involved subconscious processing (via subliminal priming), both high- and low-prejudiced people exhibited similar levels of prejudice. In that task, participants who had been subliminally primed with words related to Blacks rated the target person in an ambiguous story as being more hostile and aggressive than participants who had been subliminally primed with neutral words. In a separate task, participants were asked to list labels and thoughts they had of Blacks. It was only in this task where participants were able to consciously process and filter out prejudicial information did low-prejudiced participants produce significantly fewer negative descriptions of Blacks than high-prejudiced participants. Devine's study more vividly demonstrated the underlying arguments of Higgins and King (1981): unless made salient and accessible for conscious processing, people rely passively on traditional stereotypes. It is only when particular constructs (e.g., gender) or self-concepts (e.g., I am not racist) are made salient that we begin to monitor and filter out inappropriate thoughts.

It must also be noted that one other underlying assumption in both studies is that people try to present themselves in the most favorable light when possible, and this is consistent with the literature on self-presentation (Jones, 1964; Jones & Pitman, 1982). In terms of the current work, the line of research in bias correction suggests that as particular attributes of avatars deviate from the normal range of human variation, those particular attributes become salient, and thus more likely to become available for conscious monitoring, reflection, and change. Together with the self-presentation literature, this suggests that when given an avatar with attributes that have negative connotations (e.g., shy, slow, inarticulate) that participants may consciously try to behave in ways that counter those perceptions. On the other hand, this should less likely be the case if the attributes have positive connotations (e.g., confident, smart). With regards to the research question of interest, this suggests that participants in 0.34 m avatars would try to correct the negative bias associated with being shy and passive and perhaps behave similarly to participants in the tall condition, whereas participants in the 8.6 m avatars would feel less compelled to correct a positive bias. In this scenario, it is the participants in the short avatars who would behave differently from participants in the other three conditions. Of course, this assumes that there are no negative stereotypes of very tall people that participants might try to correct for. I will explore this in more detail with a pretest detailed in the next section.

Given that I have baselines and working dependent measures for manipulations in avatar height and it is easy to manipulate height outside of the normal human range of variation, participants in this study will be placed in either slightly shorter or taller avatars or much shorter or taller avatars (i.e., 0.34 m or 8.6 m, 5 times shorter or taller than the baseline 1.72 m respectively). The theories described above provide explanations for conflicting outcomes. As such, I will state the predictions derived from each explanation and see whether the data bears one of them out.

H1: Given that extreme avatars are less typical members of their categories, the effects of the stereotype should be weakened. Thus, I would expect that participants in a 0.34 m tall avatar to be less shy (i.e., more aggressive) than participants who are slightly shorter than the confederate. I would also expect that participants in the 8.6m tall avatar to be less aggressive than participants who are slightly taller than the confederate. Thus, in a planned comparison, I would expect a significant linear contrast with the following weights in the order of tall (2), very tall (1), very short (-1), and short (-2).

H2: Given that participants may not have stable or similar stereotypes of extreme avatars due to lack of exposure, participants may behave according to their own exemplar-based stereotypes and no group-level differences would emerge. Thus, I would expect a planned comparison using the following weights to be significant: tall (1), short (-1), very tall (0), and very short (0).

H3: Given the salience of the attribute in the extreme avatar, participants may consciously try to correct for negative biases. Thus, participants in the 0.34 m condition may behave significantly more aggressively than participants who are slightly shorter than the confederate. On the other hand, because being tall introduces more positive biases, I wouldn't expect participants in the 8.6m condition to correct any positive bias. Thus, I would expect a planned comparison using the following weights to be significant: tall (1), very tall (1), very short (1), and short (-3).

Pretest

Given that the Proteus Effect is a behavioral consequence of conforming to underlying stereotypes, understanding whether people have stable stereotypes of 0.34

m tall or 8.6 m tall people should be the precursor to studying the behavioral consequences of the Proteus Effect. Put another way, it would seem premature to hypothesize too much about behavioral outcomes of implausible bodies without first understanding stereotypes attributed to these bodies especially if the stereotypes are assumed to drive the behavioral consequences. Thus, for example, if people did indeed use exaggerated stereotypes of short people to conceptualize a 0.34 m tall person, this information would greatly aid in my prediction of the behavioral consequences of putting a user in a 0.34 m tall avatar. To explore stereotypes that people may have of these implausible bodies, I conducted a pretest using an open-ended elicitation approach similar to that used by Devine (1989). While I did not have specific hypotheses, I was interested in comparing the stereotype strength of different social groups, and in particular, stereotypes of 0.34 m and 8.6 m people compared with baseline groups.

Method

Design

In this pretest, participants were asked to list stereotypes that the general population may attribute to a particular social category. This instruction wording was borrowed from Devine's (1989) study where the emphasis was not on personal beliefs but the participant's knowledge of the content of the cultural stereotype. Aside from the specific social categories of interest (i.e., 0.34 m tall and 8.6 m tall people), I also included the "tall people" and "short people" categories as a baseline for comparison. As general baselines, I included highly-stereotyped groups (e.g., doctors and Blacks) and groups that were unlikely to be stereotyped (e.g., people with hazel eyes). And finally, we also included an assortment of well-known animals (e.g., tigers), less well-known animals (e.g., duck-billed platypus), and fantasy creatures (e.g., vampires) to explore stereotypes of non-human bodies. Altogether, there were 12 social groups that participants were asked to list stereotypes for: Blacks, doctors, tall people, people with hazel eyes, tigers, eagles, zebras, duck-billed platypus, vampires, unicorns, people who are 0.34 m tall, and people who are 8.6 m tall[3].

Participants were also asked to rate each category in terms of how difficult it was to generate the stereotypes as well as how consistent they thought the generated list would be from person to person. Both questions used a 5-point fully-labeled rating scale. The former had extremes labeled “Not Difficult At All” (1) and “Extremely Difficult” (5) while the latter had extremes labeled “Not Consistent At All” (1) and “Extremely Consistent” (5).

Participants

59 undergraduate students participated in this pretest for course credit.

Coding

Two coders independently coded the responses without an *a priori* coding scheme. The responses were coded by creating new codes for novel stereotype items. In cases where a given item was similar to an existing code (e.g., smart and intelligent, or criminals and dangerous), the coder used their own judgment as to whether a new code needed to be created. In cases where a participant provided multiple traits that were part of the same code (i.e., highly synonymous traits), those traits were all only

[3] In the actual pretest wording, imperial units were used instead of metric units in the case of the 0.34 m and 8.6 m categories because it was presumed that the participant pool was not familiar with metric measurements of height. The two measurements were phrased as 1-foot tall and 25-foot tall people respectively.

counted once under the appropriate code. Stereotype items for each social category were coded using their own set of codes. In other words, codes were generated for each social category independently. Across both coders, each social category had on average 56.6 codes (*SD* = 15.8). Coder inter-reliability is presented in the following sections after descriptions of the specific measures created from the coding.

Measures

Stereotype Strength. Similar to Devine's (1989) study, I created a measure of stereotype strength based on the proportion of respondents who listed the same stereotype. To calculate the stereotype strength of each social group in the pretest, I calculated the proportion of total respondents who provided each stereotype and took the average of the top five stereotypes. Thus, a stereotype strength of .30 would imply that the top five stereotypes were listed on average by 30% of the respondents. For this analysis, I excluded the trait "imaginary/fictional" for the fictional or implausible categories because this trait tended to be an outlier for those categories. The coder inter-reliability of stereotype strength across the 12 social groups was .95.

Stereotype Content. The codes also allowed us to explore the content of the stereotypes qualitatively. A high degree of semantic similarity was shared between the two coders regarding the top five stereotypes for each social category. For example, the top five coded stereotypes for Blacks are shown in Table 4. These findings are consistent with those in Devine's (1989) study where the top five codes were: poor, aggressive/tough, criminal, low intelligence, and uneducated. Athletic and rhythmic were rank 8 and 9 respectively in Devine's study.

Table 4.

Top Five Codes for Blacks Among the Two Coders.

	Coder One	Coder Two
1	Dangerous / Criminals / Gangs	Criminals / Dangerous / Problems with police / Gangs / Violent
2	Athletic	Athletic / Fast
3	Music / Rap / Hip-Hop	Unintelligent / Ignorant / Drop-Outs / Uneducated
4	Not Intelligent / Ignorant / Uneducated	Poor / Lower class
5	Poor	Great at music / Musical / Rappers / Singers

Results and Discussion

Stereotype Strength

The average stereotype strength for each social category is listed in Table 5 in descending order. As would have been predicted, the social groups that are most heavily-stereotyped had higher stereotype strengths, but more importantly, they provided a baseline for the social categories of interest. In particular, the two target social categories of interest (i.e., 0.34 m people and 8.6 m people) had lower stereotype strengths than the normal "Tall People" category.

It is also worth pointing out that fictional social groups do not necessarily have low stereotype strengths. For example, vampires have higher stereotype strengths than tall people and eagles. Thus, it is possible for fictional social groups to become heavily stereotyped through repeated exposure via television, movies, and other media.

Table 5.

Average Stereotype Strength for All 12 Social Categories.

Category	Stereotype Strength
Blacks	54%
Doctors	47%
Vampires	44%
Tigers	36%
Tall People	36%
Zebras	34%
Unicorns	33%
Eagles	26%
Short People	25%
8.6 m People	23%
Hazel Eyes	23%
Platypus	22%
0.34 m People	16%

Stereotype Content

The content of the stereotypes among the height conditions can also clarify the kinds of traits associated with them. The top 10 stereotypes (averaged across both coders) are shown in Table 6. Consistent with the literature in height as well as my findings from Pilot Study Two, traits related to confidence, dominance, and power are associated with tall people. And to a certain extent, we see the appropriate variations of these traits for the 8.6 m people (e.g., mean, powerful) and 0.34 m people (e.g., shy, timid, wimpy), but overall, the three categories appear to draw from distinct constructs (i.e., confident lanky athletes, barbaric giants, and small magical beings).

Table 6.

Top 10 Stereotype Traits Associated with Social Groups in the Height Conditions.

	Short People	Tall People	8.6 m People	0.34 m People
1	Physically Weak	Basketball players / Good athletes / Jocks	Mean / Unfriendly / Villains / Angry	Helpless / Easily Squished
2	Insecure / Self-Conscious	Handsome / Pretty	Violent / Destructive / Fierce / Brutal	Have High Voices
3	Inferiority Complex / Aggressive / Defensive	Awkward	Dumb	Reserved / Quiet / Timid / Shy
4	Unattractive	Strong / Powerful	Scary	Freak / Abnormal / Weird
5	Shy / Quiet / Not Confident	Skinny	Barbaric / Primitive	Weak / Wimpy
6	Different / Special	Gangly / Lanky	Big / Heavy	Cute
7	High Voices	More confident	Very tall / Huge	Unintelligent / Stupid
8	Annoying	Clumsy / Uncoordinated / Goofy	Strong / Powerful	Bearded
9	Less Mature	Intimidating / Powerful	Awkward / Uncoordinated / Clumsy	Child-Like
10	Not Good Leaders	In control / Dominant	Ugly	Enchanted / Magical

Subjective Ratings

The ratings of generative difficulty and perceived consistency were highly correlated (r = -.97). The means for each social category are listed in Table 7. The sum of the two ratings, once generative difficulty was reverse coded, was also highly correlated with the coded stereotype strengths listed in Table 5 (r = .79). Thus, participants' subjective ratings of stereotype consistency were highly congruent with the coded results.

Table 7.

Means and Standard Deviations of Generative Difficulty and Perceived Consistency of Each Category.

Category	Difficulty	Consistency
Blacks	1.56 (0.82)	3.73 (0.87)
Doctors	1.44 (0.69)	3.58 (0.82)
Vampires	1.92 (1.1)	3.23 (1.15)
Tigers	2.15 (1.25)	3.18 (1.03)
Tall People	1.87 (0.95)	3.16 (0.81)
Zebras	2.3 (1.16)	3 (1.04)
Unicorns	2.51 (1.21)	2.86 (1.01)
Eagles	2.13 (0.9)	2.77 (0.95)
8.6 m People	2.61 (1.35)	2.66 (1.09)
Short People	3.15 (1.41)	2.23 (0.83)
Hazel Eyes	2.89 (1.28)	2.1 (1.04)
Platypus	3.44 (1.23)	1.84 (1.03)
0.34 m People	3.48 (1.33)	1.81 (0.97)

Together, these three sets of findings suggest that most people do not seem to have strong stereotypes of 0.34 m and 8.6 m people. This resonates with category-based induction in that 0.34 m people appear to be atypical members of the short people category (and similarly for 8.6 m people in the case of tall people), and suggests that the initial intuitive outcome - the classical view that stereotypes of 0.34 m people are simply exaggerated stereotypes of short people - is highly unlikely to occur.

While the pretest data helps to reject the classical view of stereotypes, the data does not help much in suggesting which of the three stated hypotheses will most likely occur. Indeed, the pretest data supports the premises of all three hypotheses. With regard to H1, the pretest data supports the premise that the 0.34 m body is an atypical member of the short people category and draws weakly from its stereotypes (and similarly for the 8.6 m body). With regard to H2, the pretest data supports the premise that the extreme bodies are weakly stereotyped and thus users may rely on exemplars to generate stereotypes. And with regard to H3, the pretest data supports the premise that the social category of 0.34 m people elicits negative stereotypes of shyness and helplessness that may lead to bias correction. On the other hand, the pretest also showed that the 8.6 m person category also elicited negative stereotypes of anger and being mean. In light of this finding, and consistent with the bias correction literature that negative stereotypes will more likely be corrected, H3 was modified.

H3: Given that negative stereotypes seem to be present in both the 0.34 m and 8.6 m categories, I would expect bias correction to occur in both conditions. In other words, participants in the 0.34 m category may behave in a more confident manner (to offset the helplessness stereotype), while participants in the 8.6 m category may behave in a more fair manner (to offset the meanness stereotype). Thus, I would expect a planned comparison using the following weights to be significant: short (-1), very tall (-1), very short (1), tall (1).

To see how participants would behave in these extreme avatars, I conducted an experimental study.

Method

Design

Study Two was highly similar to Pilot Study Two. In this study, there were four height conditions. In all the height conditions, participants interacted with a confederate whose avatar was 1.72m tall. In the *very tall* condition, participants were given an avatar that was five times taller than the confederate's avatar (i.e., 8.6m). In the *tall* condition, participants were given an avatar that was 10cm taller than the confederate's avatar. In the *short* condition, participants were given an avatar that was 10cm shorter than the confederate's avatar. And finally, in the *very short* condition, participants were given an avatar that was 5 times shorter than the confederate's avatar (i.e., .34m). Participants engaged in a negotiation task (as in Pilot Study Two) with a confederate stranger.

Participants

60 undergraduate students (29 male, 31 female) participated in the study for course credit or $5.

Materials and Apparatus

The apparatus used was identical to that described in Pilot Study One. The mock dating website task was the same as that described in Study One. Within the virtual environment, participants were provided with a virtual body so that they knew that they were in a very large or very small body rather than simply being in the air or lying low on the ground.

Procedure

In the virtual environment, participants went through the same procedure as that described in Pilot Study Two with one additional instruction. To convince participants in the very tall and very short conditions that they were in very large and very small bodies respectively, participants were asked in all conditions to first look down and confirm that they could see their own bodies in the virtual environment. After this confirmation, participants were greeted by the confederate and then the researcher guided them through the negotiation task.

Measures

As in Pilot Study Two, negotiation performance was recorded: the splits offered and reaction to the final unfair split.

Participant Height. The optical tracking system was used to measure the height of the participant after they had donned the head-mounted display. The participant's height was then used as a covariate in the analyses.

Results & Discussion

As in Pilot Study Two and Four, we analyzed the three parts of the negotiation task separately. In cases where the omnibus test was significant, we then conducted the planned comparisons.

Split One

We conducted an ANOVA with height condition as the independent variable, participant height as the covariate, and the amount the participant split in their own favor in split one as the dependent variable. The covariate of participant height was not significant ($F[1,55] = .00$, $p = .99$, $\eta^2 < .001$). The effect of height condition was also not significant ($F[3,55] = .26$, $p = .85$, $\eta^2 = .001$). Because the ANOVA was not significant, we did not carry out any planned comparisons.

Split Two

We conducted an ANOVA with height condition as the independent variable, participant height as the covariate, and the amount the participant split in their own favor in split two as the dependent variable. The covariate of participant height was not significant ($F[1,55] = .00$, $p = .52$, $\eta^2 = .001$). The effect of height condition was significant ($F[3,55] = 3.16$, $p = .03$, $\eta^2 = .14$), see Figure 14. The planned comparison for H1 was not significant, ($t[55] = .51$, $p = .61$). The planned comparison for H2 was not significant ($t[55] = 1.35$, $p = .18$). The planned comparison for H3 was significant ($t[55] = 2.01$, $p = .05$).

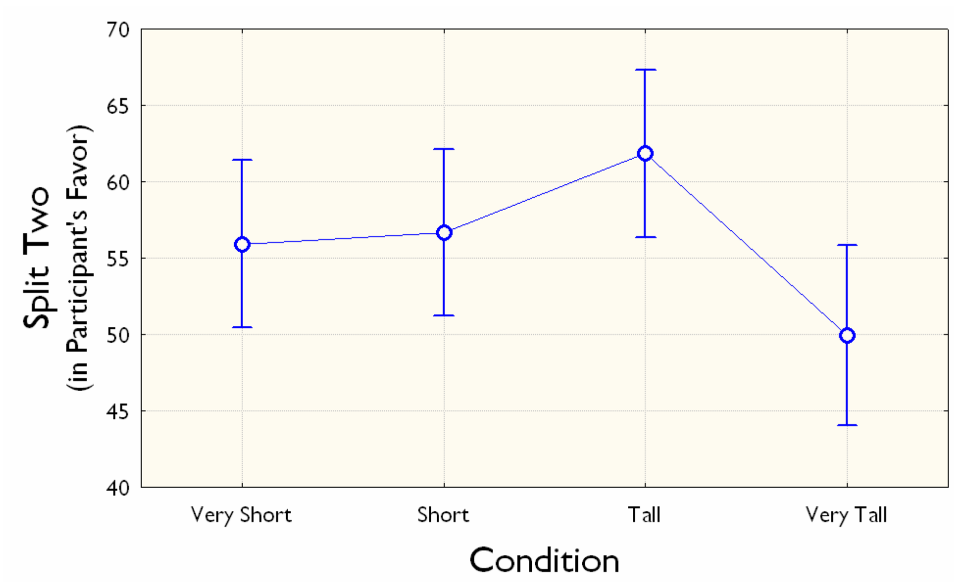

Figure 14. Means and 95% confidence intervals of height condition.

Acceptance of Unfair Split

We conducted a logistic regression with height condition as the independent variable, participant height as a covariate, and the acceptance of the unfair split as the dependent variable. The model was significant ($\chi^2[4] = 9.45$, $p = .05$, Cox & Snell r-square = .15). See Figure 15 for means and 95% confidence intervals of the four height conditions. We conducted the planned comparisons by creating dummy variables consisting of the specified weights for each condition. We then conducted a series of logistic regressions using the dummy variable as the independent variable, participant height as a covariate, and the acceptance as the dependent variable. In the planned comparison for H1, the model was significant ($\chi^2[2] = 6.89$, $p = .03$, Cox & Snell r-square = .11) and the dummy variable was a significant predictor ($p = .04$). In the planned comparison for H2, the model was also significant ($\chi^2[2] = 6.27$, $p = .04$, Cox & Snell r-square = .10) and the dummy variable was a marginally significant predictor ($p = .06$). And finally, the planned comparison for H3 was not significant ($\chi^2[2] = 3.07$, $p = .22$, Cox & Snell r-square = .05).

While the model was significant, the condition means (Figure 15) revealed that the only difference was between the short condition and the other three conditions. As such, the planned comparisons may be somewhat misleading in this case, as participants in the short condition accepted more than the other three conditions; and this single difference may be effecting the various planned comparisons.

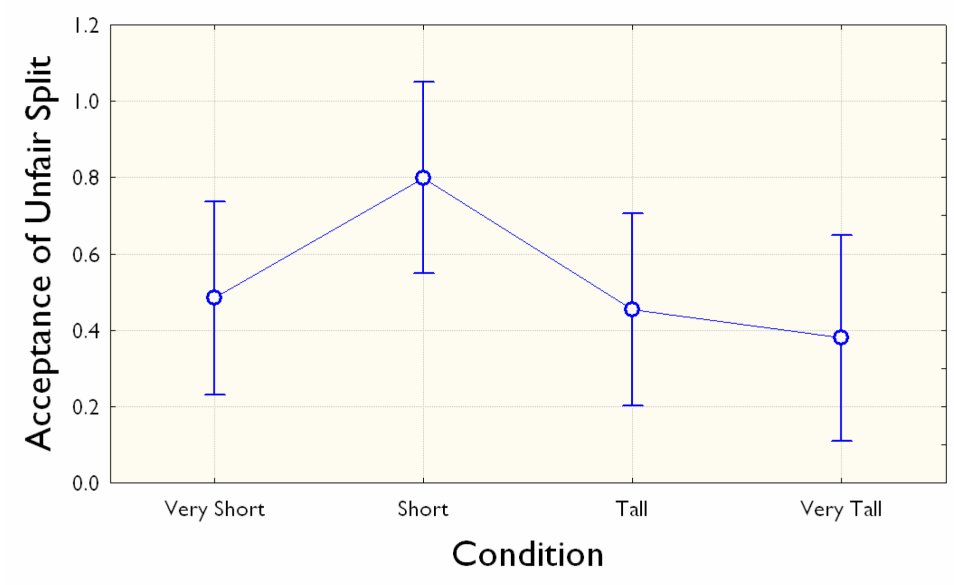

Figure 15. Means and 95% confidence intervals for height condition.

What is clear from the current study, is that in both split two and the acceptance of the unfair split, trends replicated the effects of Pilot Study Two. Participants in the tall avatars made larger splits in their own favor and were more likely to reject the unfair offer while participants in the short avatars were less likely to make splits in their own favor and more likely to accept the unfair offer.

Support for the three hypotheses was somewhat mixed. While the planned comparisons for H1 and H2 were significant in the acceptance analysis, only the planned comparison for H3 was significant in the analysis of split two. The means of the acceptance (in Figure 15) suggest that it is the replication of the short-tall difference that drove the significant planned comparisons. Thus, while the planned comparisons were significant, the means of the conditions do not provide convincing support for the claims of H1 and H2 (both of which would have predicted that the means of the very short and very tall conditions to fall in-between the short and tall conditions).

Understanding the pattern in the means of split two provides one potential explanation for these mixed findings. The means of split two by height condition suggest that it is bias correction in the very tall condition that drove the effect. Participants in the very short condition behaved no differently from those in the short condition. Given that participants in the very tall condition were the least aggressive in split two, we might have expected them to be the most yielding in the final unfair split, but they were in fact the least likely to yield. To a certain degree, this suggests that the bias correction in that condition was driven by a sense of fairness (rather than high or low aggression per se). In other words, to avoid fulfilling the expectation that a 28.6 m tall person could easily take advantage of others (i.e., being mean), participants in that condition tried to be as fair as possible. This meant splitting close to 50-50 as well as only accepting splits that were close to 50-50.

While there was a certain degree of support for bias correction in the very tall condition, we didn’t find support for a comparable bias correction in the very short condition. On the other hand, stereotype consistency for 0.34 m people was the weakest in both the coded results and the subjective perceptions of respondents as seen in the pretest. Thus, perhaps bias correction was not observed in the very short condition due to low stereotype consistency to begin with. Overall, the results of the pretest and the experimental study suggest that the Proteus Effect does not transfer to implausible bodies in a straight-forward way. Furthermore, the results of the experimental study suggest that understanding the consequences of the Proteus Effect in implausible bodies requires careful analyses of stereotype content to guide hypotheses of behavioral outcomes. Nevertheless, the present findings do suggest that bias correction appears to the likely outcome of being in implausible bodies that carry consistent negative stereotypes. Additional research needs to be conducted to clarify the consequences of being placed in implausible bodies as the results here were somewhat mixed. However, the findings in the current study are important in that they replicated the short-tall avatar differences in Pilot Studies One and Three.

CONCLUSION

Summary of Findings

Previous research in Transformed Social Interaction has shown that our behaviors and appearances in virtual environments can be strategically altered for social advantages. Thus, our avatars could employ behavioral mimicry and facial similarity to become more persuasive and likeable to others. In the current work, I sought instead to explore how our avatars can change our own behaviors. In particular, I hypothesized that users would conform to attitudes and behaviors expected of their avatars, a phenomenon I termed the Proteus Effect.

In my first two pilot studies, I found that changes to our digital self-representations can have a significant impact on how we behave online. Participants given attractive avatars became friendlier to confederate strangers than participants given unattractive avatars. I also found that participants given avatars that were 10cm taller than the confederate stranger were more confident and aggressive in a negotiation task than participants given avatars that were 10cm shorter than the confederate. Thus, how friendly or confident a person was in a virtual environment can be greatly influenced by manipulating certain features of their avatars. Moreover, these effects emerged after a fairly brief exposure to the virtual environment. In the case of the attractiveness manipulation, participants interacted with the confederate after about 30 to 45 seconds seeing themselves in the virtual mirror.

In the third and fourth pilot studies, I tried to explore whether the Proteus Effect would generalize to an actual online community and whether the effect would persist (even for a short period) outside of the virtual environment. Using a large data set from the online game *World of Warcraft*, I found that a character's height and attractiveness were both significant predictors of their performance level within the environment. Avatars that were taller or more attractive were significantly more likely to be higher level than avatars that were shorter or less attractive. This data suggests that the Proteus Effect occurs in natural virtual communities and is not limited to short-term, dyadic, interactions. In my final pilot study, I also found that the Proteus Effect persists even when the participants have been taken out of the virtual environment. Participants that had been given taller avatars continued to negotiate more aggressively face-to-face compared with participants that had been given shorter avatars. In other words, our avatars do not only significantly change how we behave in a virtual environment; our avatars in fact continue to influence our behaviors even when we step away from the virtual environment.

Our pilot studies helped to provide a foundational set of studies showing that the Proteus Effect does occur with two different manipulations and with replication in different settings. In the dissertation studies, I focused more on the underlying mechanisms of the Proteus Effect. In the first dissertation study, I tried to isolate the contribution of the Proteus Effect from that of priming. Indeed, if priming could explain the observed behaviors, then the observed findings have nothing to do with self-representation in a digital environment. In that study, I found that the observed differences were largely driven by the mirror condition (i.e., the Proteus Effect) rather than by the playback condition (i.e., behavioral assimilation). This suggests that the observed findings in the Proteus Effect have a great deal to do with transformed self-representation in a digital environment.

More importantly, the findings of the first dissertation study supported my claims that the effects of embodiment are unique and distinct from that of observing the visual stimulus alone. In other words, being placed in a new body and interacting with others via that new body can lead to rapid behavioral changes and this process is distinct from that of behavioral assimilation. Thus, there is something special about having an avatar and being embodied that produces behavioral changes. Seeing the same visual stimulus alone did not produce effects as strong as when participants were embodied.

In the second dissertation study, I sought to further understand the underlying mechanism of the Proteus Effect by exploring the behavioral consequences of placing a user in an implausible digital body. Thus, instead of using avatars that varied within the normal range of human variation, I in addition placed participants in avatars that were 8.6 m or 0.34 m tall. Contrary to the classical view that stereotype strength is a linear function of the extremity of the feature, I found some support instead for bias correction. Participants in the 8.6 m tall condition tried to be as fair as possible (to correct for the negative bias of being mean). They split close to 50-50 and rejected splits that deviated from 50-50. On the other hand, I did not see a comparable effect for participants in the 0.34 m condition. Given that stereotype consistency was lowest for the 0.34 m category, this suggests that it might have been hard for participants in that condition to correct for a consistent bias overall. These findings suggest that the consequences of the Proteus Effect do not simply follow a linear progression depending on the extremity of the manipulated features. In certain cases of negative stereotypes, these biases can become corrected by the user to produce the opposite of the associated behavior.

Together, these findings show that as we choose and customize our avatars, our avatars come to change us in turn. Our avatars can shape our behaviors both inside and outside the virtual environment. Thus, creating an avatar is very much a two-way process.

Limitations and Future Directions

There were several limitations to my studies which point to future research directions in this area. First of all, an important aspect of having an avatar in a digital environment is the element of choice. As was mentioned in the Introduction, users in many virtual environments are able to heavily customize their avatars. Users in virtual environments seldom have a pre-set avatar forced upon them as in my studies. It is unclear whether or how a user's personality or self-perception may lead to differences in avatar choice, and how these differences might interact with the Proteus Effect. For example, are short people likely to select shorter or taller avatars when they are given the choice? While my studies did not address the element of choice, my findings do suggest that changes in digital self-representation lead to behavioral differences and this is still an important finding. It is also important to bear in mind that avatar customization in virtual environments may not be as open-ended as possible. For example, in the social online world *There*, users can only create youthful avatars. Old people do not exist in *There*. In other words, there may be many features of our avatars that we actually don't have control over in online environments even when avatar customization is possible.

A second limitation was in my reliance on manipulations in feature ranges rather than categories. Specifically, both height and attractiveness are features that vary on a continuum. On the other hand, categorical changes may involve gender or ethnicity - social categories that have stronger boundaries. It would be interesting to explore whether the Proteus Effect extends to categorical changes in self-representation. In light of the findings of the second dissertation study, it is possible that a great deal of bias correction may occur especially when tasks seem to demand gender-stereotypical or race-stereotypical responses. On the other hand, previous studies in social perception theory that have employed strong identity cues (i.e., nurses uniforms or KKK-like uniforms) did not produce bias correction effects. Thus, it would be interesting to explore the behavioral consequences of categorical manipulations in self-representation.

A third limitation stems from my exclusion of the possibility of behavioral confirmation by keeping confederates blind to condition. While this experimental paradigm helped to ensure we were able to observe the unique contribution of the Proteus Effect, a social interaction in virtual environments is the outcome of both the Proteus Effect and behavioral confirmation working in tandem. On the one hand, this implies that behavior changes from avatar features may in fact be larger in a natural setting once behavior confirmation is factored in. On the other hand, this suggests that mismatches between self-perception and how others perceive us in a virtual environment may be a fruitful avenue of research. For example, what happens when we pit the Proteus Effect against behavioral confirmation? When users perceives themselves as tall, but other interactants perceive them as short, do the effects cancel each other out, or does one or the other tend to triumph?

Our findings also point at other directions for future research. The first concerns the duration of the Proteus Effect outside of the virtual environment. In the fourth pilot study, we saw that the effect of avatar height persisted to a face-to-face negotiation task. And in the first dissertation study, we saw that the effect of avatar attractiveness persisted on a dating website outside of the immersive virtual environment. In both cases, however, only a short time period elapsed between the virtual environment and the task itself. It would be interesting to examine more carefully how long the Proteus Effect persists for outside of the virtual environment. It is also important to keep in mind that the typical online game user spends on average 20 hours a week in their digital avatar, and that their online behavioral repertoires may shape and become part of their face-to-face behavioral repertoire over time.

While my pilot studies seemed to show a clear-cut linear consequence to manipulations in attractiveness and height, my study of the responses to implausible bodies showed that other processes come into play when the manipulations fall outside of the range of normal human variation and especially if they carry negative stereotypes. Once these manipulations were made salient, they became accessible for conscious bias correction. Our pretest on the content of different stereotypes also highlighted the fact that some categories have more stable and consistent content than others. Thus, future work in this area needs to take into account both bias correction and unstable stereotypes especially when exploring implausible bodies.

Implications

The set of studies presented in the current work makes clear that our self-representations have a significant and almost instantaneous impact on our behavior. The appearances of our avatars shape how we interact with others. As we choose our self-representations in virtual environments, our self-representations shape our behaviors in turn. These changes happen not over hours or weeks, but within minutes.

These findings hint at interesting implications and the potential for using avatars specifically to engineer behavior changes. For example, my findings suggest that regularly placing someone in a tall avatar may help to improve their self-esteem and confidence over time in face-to-face settings. On the other hand, this also implies that negative behavioral changes might also carry over to face-to-face interactions. Thus, playing a video game that places users in dark, menacing avatars may lead to increased aggression in face-to-face settings. While the literature on video game violence is framed by the general aggression model (Bushman & Anderson, 2002) - i.e., players are primed to interpret others as hostile in the physical world after playing a violent video game - the Proteus Effect suggests that embodiment may augment this transfer of aggression. In particular, Dissertation Study One highlighted the unique contributions of digital embodiment over the presentation of visual stimulus alone. Thus, the Proteus Effect highlights the potential for both positive and negative behavioral consequences from video games, and virtual environments in general.

On the level of the virtual community (Rheingold, 1993), these results imply that the design and implementation of avatars can have a significant effect on shaping the emergent social norms and interaction patterns. To put it more bluntly, these results suggest that friendliness in an online community can be significantly altered by changing how avatars are designed. For example, by skewing the distribution of avatar attractiveness in the customization settings towards more attractive avatars, users may behave in a more friendly and pro-social manner. Thus, we may think of friendliness or confidence as an individual personality trait, but in a virtual environment, friendliness and confidence can be architectural traits of the social system. This suggests that avatar design can be thought of as a form of social engineering in virtual communities.

Every day, millions of users interact with each other via graphical avatars in real time in online games (Chan & Vorderer, 2006). All of them are using an avatar that differs from their physical appearance to some degree. In fact, most of them are using avatars that are attractive, powerful, youthful, and athletic. Theoretical frameworks of understanding our digital self-representations are important because choosing who we are is a fundamental part of being in a virtual environment. As my studies showed, who we choose to be in turn shapes how we behave. While avatars are usually construed as something of our own choosing - a one-way process - the fact is that our avatars come to change how we behave.

REFERENCES

Allport, G. (1954). *The nature of prejudice*. Reading, MA: Addison-Wesley.

Anderson, C. A., & Bushman, B. J. (2001). Effects of violent video games on aggressive behavior, aggressive cognition, aggressive affect, physiological arousal, and prosocial behavior: A meta-analytic review of the scientific literature. *Psychological Science, 12*(5), 353-359.

Anonymous. (1998). To reveal of not to reveal: A theoretical model of anonymous communication. *Communication Theory, 8*, 381-407.

Ansfield, M., & Wegner, D. (1996). The feeling of doing. In P. Gollwitzer & J. Bargh (Eds.), *The Psychology of Action*. New York: Guildford Press.

Bailenson, J., & Blascovich, J. (2004). Avatars. In W. S. Bainbridge (Ed.), *Encyclopedia of Human-Computer Interaction*: Berkshire Publishing Group.

Bailenson, J., Blascovich, J., Beall, A., & Loomis, J. (2003). Interpersonal distance in immersive virtual environments. *Personality and Social Psychology Bulletin, 29*, 1-15.

Bailenson, J., Garland, P., Iyengar, S., & Yee, N. (2006). Transformed Facial Similarity as a Political Cue: A Preliminary Investigation. *Political Psychology, 27*, 373-386.

Bailenson, J., Iyengar, S., Yee, N., & Collins, N. (2006). *Facial Similarity as a Voting Heuristic: Some Experimental Evidence.*Unpublished manuscript.

Bailenson, J., Shum, M. S., Atran, S., Medin, D. L., & Coley, J. C. (2002). A bird's eye view: Triangulating biological categorization and reasoning within and across cultures. *Cognition, 84,* 1-53.

Bailenson, J., & Yee, N. (2005). Digital Chameleons: Automatic Assimilation of Nonverbal Gestures in Immersive Virtual Environments. *Psychological Science, 16*, 814-819.

Bailenson, J., & Yee, N. (in press). Virtual interpersonal touch and digital chameleons. *Journal of Nonverbal Behavior*.

Bailenson, J., Yee, N., Blascovich, J., & Guadagno, R. (in press). Transformed Social Interaction in Mediated Interpersonal Communication. In E. Konijn, M. Tanis,

S. Utz & A. Linden (Eds.), *Mediated Interpersonal Communication*: Lawrence Erlbaum Associates.

Bandura, A. (1977). *Social Learning Theory*. Englewood Cliffs, New Jersey: Prentice Hall.

Bandura, A. (1986). *Social Foundations of Thought and Action: A Social Cognitive Theory*. Englewood Cliffs, New Jersey: Prentice-Hall.

Bandura, A., Ross, D., & Ross, S. (1961). Transmission of aggressions through imitation of aggressive models. *Journal of Abnormal and Social Psychology, 63*, 575-582.

Bargh, J. (1996). Automaticity in social psychology. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 169-183). New York: Guilford.

Bargh, J., Chen, M., & Burrows, L. (1996). The Automaticity of Social Behaviour: Direct effects of trait concept and stereotype activation on action. *Journal of Personality and Social Psychology, 71*, 230-244.

Bargh, J., & Pietromonaco, P. (1982). Automatic information processing and social perception: The influence of trait information presented outside of conscious awareness on impression formation. *Journal of Personality and Social Psychology, 43*, 437-449.

Bem, D. (1972). Self perception theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology (Vol. 6)*. New York: Academic Press.

Berscheid, E., & Walster, E. (1979). *Interpersonal Attraction*. Menlo Park: Addison-Wesley.

Bolton, G. (1991). A Comparative Model of Bargaining: Theory and Evidence. *The American Economic Review, 81*, 1096-1136.

Brewer, M., Dull, V., & Lui, L. (1981). Perceptions of the elderly: Stereotypes as prototypes. *Journal of Personality and Social Psychology, 41*, 656-670.

Brock, T. (1965). Communicator-recipient similarity and decision change. *Journal of Personality and Social Psychology, 1*, 650-654.

Bruner, J. (1951). Personality dynamics and the process of perceiving. In R. Blake & G. Ramsey (Eds.), *Perception: An approach to personality* (pp. 121-147). New York: Ronald Press.

Burger, J., Messian, N., Patel, S., del Prado, A., & Anderson, C. (2004). What a coincidence! The effects of incidental similarity on compliance. *Personality and Social Psychology Bulletin, 30*, 35-43.

Burgoon, J. (1978). A communication model of personal space violation: Explication and an initial test. *Human Communication Research, 4*, 129-142.

Burgoon, J., & Walther, J. (1990). Nonverbal expectancies and the evaluative consequences of violations. *Human Communication Research, 17*, 232-265.

Burgoon, J., Walther, J., & Baesler, E. (1992). Interpretations, evaluations, and consequences of interpersonal touch. *Human Communication Research, 19*(237-263).

Busemeyer, J., Dewey, G., & Medin, D. (1984). Evaluation of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 638-648.

Bushman, B. J., & Anderson, C. A. (2002). Violent video games and hostile expectations: A test of the general aggression model. *Personality & Social Psychology Bulletin, 28*(12), 1679-1686.

Byrne, D. (1971). *The Attraction Paradigm*. New York: Academic Press.

Cappella, J., & Panalp, A. (1981). Talk and Silence Sequences in Informal Conversations. Interspeaker Influence. *Human Communication Research, 7,* 117-132.

Carver, C., Ganellen, R., Froming, W., & Chambers, W. (1983). Modeling: An analysis in terms of category accessibility. *Journal of Experimental Social Psychology, 19*, 403-421.

Chan, E., & Vorderer, P. (2006). Massively Multiplayer Online Games. In P. Vorderer & J. Bryant (Eds.), *Playing computer games - Motives, responses, and consequences*. Mahwah, NJ: Lawrence Erlbaum.

Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: The perception-behavior link and social interaction. *Journal of Personality & Social Psychology, 76*(6), 893-910l.

Christenson, P., & Roberts, D. (1998). *It's Not Only Rock and Roll: Popular Music in the Lives of Adolescents*. Cresskill, NJ: Hampton Press.

Collins, A., & Loftus, E. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 6*(407-428).

Culnan, M. J., & Markus, M. L. (1987). Information technologies. In F. M. Jablin & L. L. Putnam (Eds.), *Handbook of organizational communication: An interdisciplinary perspective* (pp. 420-443). Thousand Oaks, CA: Sage Publications, Inc, 1987, 781.

Davis, M. H., Conklin, L., Smith, A., & Luce, C. (1996). Effect of perspective taking on the cognitive representation of persons: A merging of self and other. *Journal of Personality and Social Psychology, 70*, 713-726.

Devine, P. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*, 5-18.

Devine, P., & Baker, S. (1991). Measurement of racial stereotype subtyping. *Personality and Social Psychology Bulletin*, *17*, 44-50.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5-18.

Dijksterhuis, A., & van Knippenberg, A. (1998). The relation between perception and behavior, or how to win a game of Trivial Pursuit. *Journal of Personality and Social Psychology, 74*, 865-877.

Dijksterhuis, A., & van Knippenberg, A. (1999). On the paramters of associative strength: Central tendency and variability as determinants of stereotype accessibility. *Personality and Social Psychology Bulletin, 25*, 529-536.

Dion, K., Berscheid, E., & Walster, E. (1972). What is beautiful is good. *Journal of Personality and Social Psychology*, *24*, 285-290.

Ekman, P. (1992). An argument for basic emotions. *Cognition and Emotion, 6*, 169-200.

Elio, R., & Anderson, J. (1981). The effects of category generalizations and instance similarity on schema abstraction. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 397-417.

Flanagin, A. J., Tiyaamornwong, V., O'Connor, J., & Seibold, D. R. (2002). Computer-mediated group work: The interaction of member sex and anonymity. *Communication Research, 29*(66-93).

Forsythe, R., Horowitz, J., Savin, N., & Sefton, M. (1994). Fairness in simple bargaining experiments. *Games and Economic Behavior, 6*, 347-369.

Frank, M., & Gilovich, T. (1988). The dark side of self and social perception: Black uniforms and aggression in professional sports. *Journal of Personality and Social Psychology, 54*, 74-85.

Friend, R. M., & Vinson, M. (1974). Leaning over backward: Jurors responses to defendants' attractiveness. *Journal of Communication, 24*, 124-129.

Galinsky, A. D., & Moskowitz, G. B. (2000). Perspective-taking: Decreasing stereotype expression, stereotype accessibility, and in-group favoritism. *Journal of Personality and Social Psychology, 78*, 708-724.

Geis, F. L. (1993). Self-fulfilling prophecies: A social psychological view of gender. In A. E. Beall & R. J. Sternberg (Eds.), *The psychology of gender* (pp. 9-54). New York, NY: Guilford Press, 1993, xxiv, 278.

Gerbner, G., & Gross, L. (1976). Living with television: The violence profile. *Journal of Communication, 26*, 172-199.

Gergen, K., Gergen, M., & Barton, W. H. (1973). Deviance in the dark. *Psychology Today, 11*, 129-130.

Giles, H., & Claire, R. (1979). *Language and Social Psychology*. Oxford: Blackwell.

Goffman, E. (1959). *The presentation of self in everyday life.* New York: Anchor Books.

Golby, A., Gabrieli, J., Chiao, J., & Ebenhardt, L. (2001). Differential responses in the fusiform region to same-race and other-race faces. *Nature Neuorscience, 4*, 845-850.

Hamilton, D., & Sherman, J. (1994). Stereotypes. In R. Wyer & T. Srull (Eds.), *Handbook of social cognition*. Hillsdale, NJ: Erlbaum.

Hancock, J., & Dunham, P. (2001). Impression Formation in Computer-Mediated Communication Revisited: An Analysis of the Breadth and Intensity of Impressions. *Communication Research, 28*, 325-347.

Hancock, J., Toma, C., & Ellison, N. (2007). *The truth about lying in online dating profiles.* Paper presented at the CHI 2007, San Jose, CA.

Higgins, E., & King, G. (1981). Accessibility of social constructs: Information - processing consequences of individual and contextual variability. In N. Cantor & J. Kihlstrom (Eds.), *Personality and social interaction* (pp. 69-121). HIllsdale, NJ: Erlbaum.

Higgins, E., King, G., & Mavin, G. (1982). Individual construct accessibility and subjective impressions and recall. *Journal of Personality and Social Psychology, 43*, 35-47.

Higgins, E., Rholes, W., & Jones, C. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology, 13*(141-154).

Hiltz, S., Johnson, K., & Turoff, M. (1986). Experiments in group decision making: Communication processes and outcome in face-to-face versus computerized conferences. *Human Communication Research, 13*, 225-252.

Hitzman, D. (1986). Schema abstraction in a multiple-trace memory model. *Psychological Review, 93*, 411-428.

Homa, D., Sterling, S., & Trepel, L. (1981). Limitation of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 418-439.

Homa, D., & Vosburgh, R. (1976). Category breadth and the abstraction of prototypical information. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 322-330.

iTinySoft. (2002). Magic Morph. Retrieved March 21st, 2004, from http://www.effectmatrix.com/morphing/index.htm

Iyengar, S. (2002). *Experimental Designs for Political Communication: From Shopping Malls to the Internet*. Paper presented at the Workshop in Mass Media Economics.

Jacobson, D. (1999). Impression formation in cyberspace: Online expectations and offline experiences in text-based virtual communities. *Journal of Computer-Mediated Communication, 5*.

Jacoby, L., & Witherspoon, D. (1982). Remembering without awareness. *Canadian Journal of Psychology, 32,* 300-324.

James, W. (1890). *Principles of psychology*. New York: Holt.

Jarvenpaa, S., & Leidner, D. (1998). Communication and trust in global virtual teams. *Journal of Computer-Mediated Communication, 3*.

Johnson, R., & Downing, L. (1979). Deindividuation and valence of cues: Effects on prosocial and antisocial behavior. *Journal of Personality and Social Psychology, 37*, 1532-1538.

Johnston, L. (2002). Behavioral mimicry and stigmatization. *Social Cognition, 20*(1), 18-35.

Jones, E. E. (1964). *Ingratiation*. New York: Irvington.

Jones, E. E., & Nisbett, R. E. (1971). *The actor and the observer: Divergent perceptions of the causes of behavior.* Morristown, NJ: General Learning Press.

Jones, E. E., & Pitman, T. S. (1982). Toward a general theory of strategic self-presentation. In J. Suls (Ed.), *Psychological perspectives on the self* (Vol. 1, pp. 231-262). HIllsdale, NJ: General Learning Press.

Katz, J., & Postal, P. (1964). *An integrated theory of linguistic descriptions*. Cambridge, MA.: MIT Press.

Kelley, H. (1950). The warm-cold variable in first impressions of persons. *Journal of Personality, 18,* 431-439.

Kelly, G. (1955). *The psychology of personal constructs*. New York: Norton.

Kendon, A. (1970). Movement Coordination in social interactions. *Acta Psychologica, 32*(2), 101-125.

Kiesler, S., Siegel, J., & McGuire, T. W. (1984). Social psychological aspects of computer-mediated communication. *American Psychologist, 39*(10), 1123-1134.

Krcmar, M., & Vieria, E. (2005). Imitating Life, Imitating Television: The Effects of Family and Television Models on Children's Moral Reasoning. *Communication Research, 32*, 267-294.

LaFrance, M. (1982). Posture Mirroring and Rapport. In M. Davis (Ed.), *Interaction Rhythms: Periodicity in Communicative Behavior* (pp. 279-298).

Lakin, J. L., & Chartrand, T. L. (2003). Using nonconscious behavioral mimicry to create affiliation and rapport. *Psychological Science, 14*(4), 334-339.

Lakin, J. L., Jefferis, V. E., Cheng, C. M., & Chartrand, T. L. (2003). The chameleon effect as social glue: Evidence for the evolutionary significance of nonconscious mimicry. *Journal of Nonverbal Behavior, 27*(3), 145-162.

Langlois, J., Kalakanis, L., Rubenstein, A., Larson, A., Hallam, M., & Smoot, M. (2000). Maxims or myths of beauty?: A meta-analytic and theoretical review. *Psychological Bulletin, 126*, 390-423.

Lea, M., O'Shea, T., & Spears, R. (1992). "Flaming" in computer-mediated communication. In M. Lea (Ed.), *Contexts of Computer-Mediated Communication* (pp. 89-112). New York: Harvester Wheatsheaf.

Levelt, W., & Kelter, S. (1982). Surface form and memory in question answering. *Cognitive Psychology, 14*(78-106).

Levy, B., & Langer, E. (1994). Aging free from negative stereotypes: Successful memory in China and among the American deaf. *Journal of Personality and Social Psychology, 66*, 989-997.

Lippman, W. (1922). *Public opinion*. New York: Harcourt Brace Jovanovich.

Liu, Y., & Shrum, L. (2003). What is Interactivity and Is It Always a Good Thing? Implications of Definition, Person, and Situation for the Influence of Interactivity on Advertising Effectiveness. *Journal of Advertising, 31*, 53-64.

Macrae, C. N., Bodenhausen, G. V., Milne, A. B., & Jetten, J. (1994). Out of mind but back in sight: Stereotypes on the rebound. *Journal of Personality and Social Psychology, 67*, 808-817.

McCauley, C., & Stitt, C. (1978). An individual and quantitative measure of stereotypes. *Journal of Personality and Social Psychology, 36*, 929-940.

Medin, D., & Schaffer, M. (1978). Context theory of classification learning. *Psychological Review*(85).

Medin, D., & Wattenmaker, W. (1987). Category cohesiveness, theories, and cognitive archaeology. In U. Neisser (Ed.), *Concepts and conceptual development* (pp. 25-62). Cambridge, England: Cambridge University Press.

Merola, N., Penas, J., & Hancock, J. (2006). *Avatar color and social identity effects: On attitudes and group dynamics in virtual realities.* Paper presented at the ICA 2006, Dresden, Germany.

Meyer, D., & Schvaneveldt, R. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90*, 227-234.

Miller, N., & Dollard, J. (1941). *Social Learning and Imitation*. New Haven, CT: Yale University Press.

Minda, J., & Smith, J. (2002). Comparing protoype-based and exemplar-based accounts of category learning and attentional allocation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 275-292.

Nass, C., & Lee, K. (2001). Does computer-synthesized speech manifest personality? Experimental tests of recognition, similarity-attraction, and consistency-attraction. *Journal of Experimental Psychology, 7*, 171-181.

Nass, C., Moon, Y., Fogg, B., Reeves, B., & Dryer, C. (1995). Can computer personalities be human personalities? *Proceedings of CHI 1995*, 228-229.

Nass, C., Robles, E., Bienenstock, H., Treinen, M., & Heenan, C. (2003). Voice-based disclosure systems: Effects of modality, gender of prompt, and gender of user. *International Journal of Speech Technology, 6*, 113-121.

NCHS. (2004). National health and nutrition examination survey 2003-2004. Retrieved 5th September, 2006, from http://www.cdc.gov/nchs/about/major/nhanes/nhanes2003-2004/nhanes03_04.htm

Neely, J. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology, 106*, 226-254.

Nelson, C. (2001). The development and neural bases of face recognition. *Infant and Child Development*, *10,* 3-18.

Neumann, R., & Strack, F. (2000). Mood Contagion: The automatic transfer of mood between persons. *Journal of Personality & Social Psychology*, *79*, 211-223.

Nosofsky, R. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39-57.

NYCLU. (2006). Who's watching? Video camera surveillance in New York City and the need for pubic oversight. Retrieved 18th April, 2007, from http://www.nyclu.org/surveillancecams/

Osherson, D., Smith, E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review, 97*, 185-200.

PARC. (2005). Leveling Time. from http://blogs.parc.com/playon/archives/2005/07/leveling_time.html

Park, J., & Schaller, M. (2005). Does attitude similarity serve as a heuristic cue for kinship? evidence of an implicit cognitive association. *Evolution and Human Behavior*, *26*, 158-170.

Peterson, R. (2001). On the Use of College Students in Social Science Research: Insights from a Second-Order Meta-analysis. *Journal of Consumer Research, 28*, 450-461.

Posner, M., & Keele, S. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology: Human Learning and Memory, 77*, 353-363.

Posner, M., & Snyder, C. (1975a). Attention and cognitive control. In R. Solso (Ed.), *Information processing and cognition: The Loyola symposium*. Hillsdale, N.J.: Erlbaum.

Posner, M., & Snyder, C. (1975b). Facilitation and inhibition in the processing of signals. In P. Rabbitt & S. Dornic (Eds.), *Attention and performance V*. New York: Academic Press.

Postmes, T., & Spears, R. (2002). Behavior online: Does anonymous computer communication reduce gender inequality? *Personality & Social Psychology Bulletin*, *28*, 1073-1083.

Postmes, T., Spears, R., & Lea, M. (1998). Breaching or building social boundaries? *Communication Research, 25*, 689-699.

Postmes, T., Spears, R., & Lea, M. (2000). The formation of group norms in computer-mediated communication. *Human Communication Research, 26*, 341-371.

Provine, R. (1986). Yawning as a stereotyped action pattern and releasing stimulus. *Ethology, 72*, 109-122.

Provine, R. (1992). Contagious Laughter: Laughter is sufficient stimulus for laughs and smiles. *Bulletin of the Psychonomic Society, 30*(1-4).

Quillian, M. (1962). A revised design for an understanding machine. *Mechanical Translation, 7*, 17-29.

Quillian, M. (1967). Word concepts: A theory and simulation of some basic semantic capabilities. *Behavioral Science, 12*, 410-430.

Reed, S. (1972). Pattern recognition and categorization. *Cognitive Psychology, 3*, 382-407.

Reeves, B., & Nass, C. (1996). *The Media Equation: How People Treat Computers, Television, and New Media like Real People and Places*: University of Chicago Press.

Regan, D. T., & Totten, J. (1975). Empathy and attribution: Turning observers into actors. *Journal of Personality and Social Psychology, 32*, 850-856.

Rheingold, H. (1993). *The Virtual Community*. Reading: Addison-Wesley Publishing Company.

Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382-439.

Schiano, D., & White, S. (1998). *The first noble truth of cyberspace: People are people (even when they MOO).* Paper presented at the Proceedings of CHI, Los Angeles.

Schroeder, R. (Ed.). (2002). *The social life of avatars*. London: Springer-Verlag.

Schwartz, D., Blair, K., Biswas, G., Leelawong, K., & Davis, J. (in press). Animations of Thought: Interactivity in the Teachable Agent Paradigm. In R. Lowe & W.

Schnotz (Eds.), *Learning with Animation: Research and Implications for Design*. UK: Cambridge University Press.

Shanteau, J., & Nagy, G. (1979). Probability of acceptance in dating choice. *Journal of Personality and Social Psychology*, *37*, 522-533.

Sherif, M., Harvey, O. J., White, B. J., Hood, W. R., & Sherif, C. W. (1961). *Intergroup cooperation and competition: The Robbers Cave experiment*. Norman, OK: University Book Exchange.

Sherman, J. (1996). Development of mental representations of stereotypes. *Journal of Personality and Social Psychology*, *70*, 1126-1141.

Short, J., Williams, E., & Christie, B. (1976). *The social psychology of telecommunications.* London: Wiley.

Shriffin, R., & Schneider, W. (1977). Controlled and automatic human information processing: Vol 2. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84,* 127-190.

Shruptine, K. (1975). On the Validity of Using Students as Subjects in Consumer Behavior Investigations. *Journal of Business*, *48,* 383-390.

Smith, E. (1990). Content and process specificity in the effects of prior experiences. In T. Srull & R. Wyer (Eds.), *Advances in social cognition* (pp. 1-59). Hillsdale, NJ: Erlbaum.

Smith, E., Shoben, E., & Rips, L. (1974). Structure and process in semantic memory: A featural model for semantic devisions. *Psychological Review*, *81*, 214-241.

Smith, E., & Zarate, M. (1992). Exemplar-based models of social judgment. *Psychological Review, 99,* 3-21.

Snyder, M., & Swann, W. B. (1978). Behavioral confirmation in social interaction: From social perception to social reality. *Journal of Experimental Social Psychology, 14*(2), 148-162.

Snyder, M., Tanke, E. D., & Berscheid, E. (1977). Social perception and interpersonal behavior: On the self-fulfilling nature of social stereotypes. *Journal of Personality & Social Psychology, 35*(9), 656-666.

Spears, R., & Lea, M. (1994). Panacea or panopticon? The hidden power in computer-mediated communication. *Communication Research, 21*, 427-459.

Srull, T., & Wyer, R. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology, 37*, 1669-1672.

Srull, T., & Wyer, R. (1980). Category accessibility and social perception: Some implications for the study of person memory and interpersonal judgments. *Journal of Personality and Social Psychology, 38*, 841-856.

Steele, C. M. (1997). A threat in the air: How stereotypes shape the intellectual identities and performance of women and African Americans. *American Psychologist*, *52*, 613-629.

Steele, C. M., & Aronson, J. (1995). Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology, 69*, 797-811.

Steuer, J. (1993). Defining Virtual Reality: Dimensions Determining Telepresence. *Journal of Communication, 42*, 73-93.

Storms, M. D. (1973). Videotape in the attribution process: Reversing actors' and observers' points of view. *Journal of Personality and Social Psychology, 27*, 165-175.

Stroop, J. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18*(643-662).

Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behavior. *European Journal of Social Psychology*, *1*, 149-178.

Taylor, S., & Crocker, J. (1981). Schematic bases of social information processing. In E. Higgins, C. HErman & M. Zanna (Eds.), *Social cognition: The Ontario Symposium* (pp. 89-134). Hillsdale, NJ: Erlbaum.

Towson, S. M. J., Zanna, M. P., & MacDonald, G. (1989). Self-fulfilling prophecies: Sex role stereotypes as expectations for behavior. In R. K. Unger (Ed.), *Representations: Social constructions of gender* (pp. 97-107). Amityville, NY: Baywood Publishing Co, Inc.

Trevino, L., & Webster, J. (1992). Flow in Computer-Mediated Communication: Electronic Mail and Voice Mail Evaluation and Impacts. *Communication Research, 19*, 539-573.

Tulving, E., Schacter, D., & Stark, H. (1982). Priming effects in word-fragment completion are independent of recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8*, 336-342.

Valins, S. (1966). Cognitive Effects of False Heart-Rate Feedback. *Journal of Personality and Social Psychology, 4*, 400-408.

van Baaren, R. B., Holland, R. W., Kawakami, K., & van Knippenberg, A. (2004). Mimicry and Prosocial Behavior. *Psychological Science, 15*(1), 71-74.

van Baaren, R. B., Holland, R. W., Steenaert, B., & van Knippenberg, A. (2003). Mimicry for money: Behavioral consequences of imitation. *Journal of Experimental Social Psychology, 39*(4), 393-398.

Walther, J. (1996). Computer-mediated communication: Impersonal, interpersonal, and hyperpersonal interaction. *Communication Research, 23*(1), 3-43.

Walther, J., Anderson, J., & Park, D. (1994). Interpersonal effects in computer-mediated interaction: A meta-analysis of social and anti-social communication. *Communication Research, 21*, 460-487.

Walther, J., Slovacek, C., & Tidwell, L. (2001). Is a Picture Worth a Thousand Words?: Photographic Images in Long-Term and Short-Term Computer-Mediated Communication. *Communication Research, 28*, 105-134.

Wegner, D. (1994). Ironic processes of mental control. *Psychological Review, 101*, 34-52.

Wegner, D., Schneider, D., Carter, S., & White, L. (1987). Paradoxical effects of thought suppression. *Journal of Personality and Social Psychology, 58*, 409-418.

Woodcock, B. (2006). MMOG Chart. from http://www.mmogchart.com/

Word, C. O., Zanna, M. P., & Cooper, J. (1974). The nonverbal mediation of self-fulfilling prophecies in interracial interaction. *Journal of Experimental Social Psychology Vol. 10(2)*, 109-120.

Wyer, R., & Carlston, D. (1979). *Social cognition, inference, and attributions*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Yee, N. (2006). The Demographics, Motivations, and Derived Experiences of Users of Massively Multi-User Online Graphical Environments. *Presence: Teleoperators and Virtual Environments, 15*, 309-329.

Yee, N., & Bailenson, J. (2006, August 24 – 26). *Walk A Mile in Digital Shoes: The Impact of Embodied Perspective-Taking on The Reduction of Negative Stereotyping in Immersive Virtual Environments.* Paper presented at the PRESENCE 2006: The 9th Annual International Workshop on Presence, Cleveland, Ohio, USA.

Yee, N., & Bailenson, J. (in press). The Proteus Effect: The Effect of Transformed Self-Representation on Behavior. *Human Communication Research*.

Young, T. J., & French, L. A. (1996). Height and perceived competence of U. S. Presidents. *Perceptual and Motor Skills*, *82*, 1002.

Zajonc, R. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, *9*, 1-27.

Zajonc, R. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, *35*, 151-175.

Zajonc, R. (2001). Mere exposure: A gateway to the subliminal. *Current Directions in Psychological Science*, *10*, 224-228.

Zajonc, R., & Markus, H. (1984). Affect and cognition: The hard interface. In C. Izard, J. Kagan & R. Zajonc (Eds.), *Emotions*, *Cognition, and Behavior* (pp. 73-102). Cambridge: Cambridge University Press.

Zaki, S., Nosofky, R., Stanton, R., & Cohen, A. (2003). Prototype and exemplar accounts of category learning and attentional allocation: A reassessment. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 6*, 1160-1173.

Zimbardo, P. (1969). The human choice: Individuation, reason, and order vs. deindividuation, impulse and chaos. In W. J. Arnold & D. Levine (Eds.), (Vol. 17, pp. 237-307). In W. Arnold & D. Levine (Eds.), *Nebraska Symposium on Motivation* (Vol. 17, pp. 237-307). Lincoln, NE: University of Nebraska Press.