JMR

INCREASING SAVING BEHAVIOR THROUGH

AGE-PROGRESSED RENDERINGS OF THE FUTURE SELF

HAL E. HERSHFIELD

DANIEL G. GOLDSTEIN

WILLIAM F. SHARPE

JESSE FOX

LEO YEYKELIS

LAURA L. CARSTENSEN

JEREMY N. BAILENSON

IN PRESS, JOURNAL OF MARKETING RESEARCH

Hal E. Hershfield is Post-doctoral Fellow and Visiting Assistant Professor at the Kellogg School of Management, Northwestern University, 2001 Sheridan Rd, Evanston, IL 60208, USA; Email: h-ersnerhershfield@kellogg.northwestern.edu. Daniel G. Goldstein is Assistant Professor of Marketing at London Business School and Principal Research Scientist at Yahoo Research, 110 West 40th Street, 17th Floor, New York, NY 10028, USA; Email: dgg@yahoo-inc.com. William F. Sharpe is the STANCO 25 Professor of Finance, Emeritus at Stanford University Graduate School of Business, 518 Memorial Way, Stanford, CA 94305-5015, USA; Email: wfsharpe@stanford.edu. Jesse Fox is an Assistant Professor of Communication at The Ohio State University, 3016 Derby Hall, 154 North Oval Mall, Columbus, OH, 43210-1339; email: fox.775@osu.edu. Leo Yeykelis is a Ph.D. student in the Communication Department at Stanford University, 420 Serra Mall, Stanford, CA 94305, USA; Email: yeyleo@stanford.edu. Laura L. Carstensen is the Fairleigh Dickinson, Jr. Professor of Public Policy and Professor of Psychology at Stanford University, Department of Psychology, 450 Serra Mall, Stanford, CA 94305, USA; Email: laura.carstensen@stanford.edu. Jeremy Bailenson is Associate Professor of Communication at Stanford University Department of Communication, 420 Serra Mall, Stanford, CA 94305, USA; Email: jbailenson@stanford.edu.

INCREASING SAVING BEHAVIOR THROUGH

AGE-PROGRESSED RENDERINGS OF THE FUTURE SELF

ABSTRACT

Many people fail to save what they need to for retirement (Munnell, Webb, and Golub-Sass 2009). Research on excessive discounting of the future suggests that removing the lure of immediate rewards by pre-committing to decisions, or elaborating the value of future rewards can both make decisions more future-oriented. In this article, we explore a third and complementary route, one that deals not with present and future rewards, but with present and future selves. In line with thinkers who have suggested that people may fail, through a lack of belief or imagination, to identify with their future selves (Parfit 1971; Schelling 1984), we propose that allowing people to interact with age-progressed renderings of themselves will cause them to allocate more resources toward the future. In four studies, participants interacted with realistic computer renderings of their future selves using immersive virtual reality hardware and interactive decision aids. In all cases, those who interacted with virtual future selves exhibited an increased tendency to accept later monetary rewards over immediate ones.

Keywords: Retirement saving, temporal discounting, future self-continuity, immersive virtual reality, intertemporal choice

Although in the United States the age associated with retirement is 65, life expectancy at that age is 18 years (Arias 2007), which reflects a gradual increase in retirement length that has been observed since the 19th century (Lee 2001). Unfortunately, with this prolonged golden age comes the risk of outliving one's money or undergoing a sudden decrease in quality of life. Munnell, Webb, and Golub-Sass (2007), for example, used data from the 2004 Survey of Consumer Finances and calculated a benchmark replacement rate for each household and found that 43% of households fell at least 10% short of reaching target replacement rates. When taking into account the recent financial crisis, the retirement picture grows even bleaker. Using the same benchmark replacement rate, Munnell, Webb, and Golub-Sass (2009) found that a few years later, the percentage of households who would fall short of reaching their retirement goals had grown to 51 percent. Moreover, the McKinsey Global Institute (2008) observed that fully two-thirds of Early Baby Boomers (born between 1945 and 1955) do not have the resources to maintain their pre-retirement standard of living in retirement, perhaps resulting from individuals not being able to forecast the consequences of their investment decisions (Goldstein, Johnson, and Sharpe 2008). Such a lack of preparedness is particularly problematic in the United States given the status of the already weakened social security system. In the words of Federal Reserve chairman Ben Bernanke, "The arithmetic is, unfortunately, quite clear" (Chan and Hernandez 2010): to avoid overwhelming budget cuts, saving behavior needs to change.

Discounting of the future was long described as exponential in economic models (e.g., Samuelson 1937), but recent behavioral research suggests that it is better fit by a hyperbolic or quasi-hyperbolic function in certain situations (Kirby and Marakovic 1996; Laibson 1997; Laibson, Repetto, and Tobacman 1998; Strotz 1956; Zauberman, Kim, Malkoc, and Bettman 2009). In particular, rather extreme discounting is thought to occur in situations in which

immediate gains (as opposed to short-term gains) are traded off against long-term ones (Chapman 1996; Frederick, Loewenstein, and O'Donoghue 2003; Laibson 1997; O'Donoghue and Rabin 1999). Such steep discounting is thought to lead to decisions that would be rejected with enough planning (Hoch and Loewenstein 1991; Lynch, Nettemeyer, Spiller, and Zammit 2010), and failure to save for retirement is held up as a prototypical example of discounting to excess (Diamond and Koszegi 2003; Laibson, et al. 1998).

Two broad types of remedies aim to increase saving and reign in excessive discounting (see Ho, Lim, and Camerer (2006) for a review of other interventions in multiple domains). Remedies of the first type reduce the lure of immediacy, and the amount that one is able to consume in the present, through pre-commitment towards starting to save at a later date (e.g., Thaler and Benartzi 2004). Pre-commitment involves any device in which consumers enact constraints in the present to limit undesirable (or promote desirable) behaviors in the future (Ariely and Wertenbroch 2002). Elster (1977) points to the example of Ulysses tying himself to the ship mast so that he could listen to the Sirens' songs without jumping overboard, and Schelling (1983) notes that alcoholics take Antabuse to ensure that they will not engage in future drinking. To this end, several companies have successfully implemented programs that use pre-commitment strategies that allow employees to commit future income to retirement plans (e.g., Choi, Laibson, Madrian, and Metrick 2006; Thaler and Benartzi 2004).

Remedies of the second type increase the appeal of waiting to spend, and the expected enjoyment of future spending, by directing people's imagination towards future uses for money. If consumers are often insensitive to the ways in which present consumption necessarily limits the ability to spend money later on (Frederick, Novemsky, Wang, Dhar, and Nowlis 2009), it may be due to people failing to elaborate on the positive future consequences of waiting

(Nenkov, Inman, and Hulland 2008; Nenkov, Inman, Hulland, and Morrin 2009). Indeed, interventions that encourage people to elaborate on future outcomes and consider future uses for money have been shown to increase patience on intertemporal choice tasks (e.g., Weber, et al. 2007).

In this article, we explore a third and complementary route, one that deals not with present and future rewards, but with the connection between present and future selves. In line with thinkers who have suggested that people may fail to identify with their future selves through a lack of belief or imagination (Parfit 1971; Schelling 1984), we propose that enabling people to interact with photo-realistic age-progressed renderings of themselves will cause them to allocate more resources towards the future. In four studies, participants interacted with computer-generated representations of their future selves using immersive virtual reality hardware and interactive decision aids before making decisions about whether to consume in the present or future.

FUTURE SELF-CONTINUITY

Theoretical Foundations and Empirical Evidence

Philosophers, psychologists, and economists have argued that one important determinant of intertemporal choice is an individual's sense – or lack thereof – of psychological connection to his or her future self (Ainslie 1975; Elster 1977; Parfit 1971, 1987; Schelling 1984; Strotz 1956; Thaler and Shefrin 1981). Psychological connectedness is conceived of as varying with respect to the age difference between different "selves," such that one might feel more connection to one's future self in one year than one's future self in 40 years, and accordingly

might care less about a distant self (Parfit 1971). At the extreme, with a total lack of psychological connectedness, one's future self may seem like a different person altogether. As Butler (1736) first pointed out, "…if the self or person of today, and that of tomorrow, are not the same, but only like persons, the person of today is really no more interested in what will befall the person of tomorrow, than in what will befall any other person" (p. 344). Parfit (1987) expands on this concept:

> If we now care little about ourselves in the further future, our future selves are like future generations. We can affect them for the worse, and, because they do not now exist, they cannot defend themselves. Like future generations, future selves have no vote, so their interests need to be specially protected. Reconsider a boy who starts to smoke, knowing and hardly caring that this may cause him to suffer greatly fifty years later. This boy does not identify with his future self. His attitude towards this future self is in some ways like his attitude to other people. (p. 319-320).

To those estranged from their future selves, saving is like a choice between spending money today or giving it to a stranger years from now. Presumably, the degree to which individuals feel connected to their future selves should make them appreciate that they are the future recipients and thus affect their willingness to save. We hold the view that it is not important whether in various senses, a person *actually* does change over time. Although trait-level personality characteristics (Roberts and DelVecchio 2000) and general interests (Low, Yoon, Roberts, and Rounds 2005) remain relatively consistent over the course of a life time, the cells in our bodies turn over, and attributes as personal as our names, noses and reputations can be willfully altered beyond recognition. What matters, however, in the sense of the law, is that one person has but one identity, and with this essential link the assets of the present and future selves are yoked together. Like Parfit, we hypothesize that individuals who *feel* as though the future self is a different person fail to acknowledge this connection, that is, fail to *identify* with themselves in the future.

Do people think of the future self as a different person? Research has demonstrated that people make attributions about the future self in the same manner that they do for others, by attributing the future self's behavior to dispositional factors rather than situational ones (Pronin and Ross 2006; Wakslak, Nussbaum, Lieberman and Trope 2008), and make decisions for the future self using a similar process that they use to make decisions for other individuals (Pronin, Olivola, and Kennedy 2008). On an implicit level, Ersner-Hershfield, Wimmer, and Knutson (2009) showed that thinking about the future self elicits neural activation patterns that are similar to neural activation patterns elicited by thinking about a stranger. More importantly, the degree to which people report feeling connected to their future selves is related to their intertemporal decision-making. Ersner-Hershfield, Garton, Ballard, Samanez-Larkin, and Knutson (2009), for example, found that higher levels of future self-continuity were positively associated with financial assets, and Bartels and colleagues showed that perceived connectedness to the future self was predictive of choices on several different temporal discounting tasks (Bartels and Rips 2010) and distinct from other constructs related to temporal preference, such as present bias (Bartels and Urminsky, in press). Moreover, in their neuroimaging study, Ersner-Hershfield and co-authors (2009) demonstrated that participants who showed the greatest neural activation differences between thoughts about the current self and thoughts about the future self also showed the steepest discount rates.

Increasing Connectedness to the Future Self

Parfit (1971) remarked that neglecting the future self might be "caused by some failure of imagination, or some false belief". The notion of a future self presents many challenges to the imagination. With the countless directions a life and appearance might take, one is unsure with which among this infinity of future selves to identify. The multiplicity of possible future selves is

acknowledged by Parfit, who cites Proust (1927/1981) to illustrate this point: "…we are incapable, while we are in love, of acting as fit predecessors of the person whom we shall presently have become and who will be in love no longer…" (p. 631). Similarly, a composite representation of all future selves would necessarily be impoverished and the quest to identify a singular image may result in indecision or lack of confidence about what the future holds.

Parfit (1987) also observes that "when we imagine pains in the further future, we imagine them less vividly, or believe confusedly that they will somehow be less real, or less painful" (p. 161). In line with this sentiment, recent research has argued that consumers suffer "empathy gaps" and may misunderstand how they will feel in the future about decisions that they make in the present (Wilson and Gilbert 2005; Loewenstein, O'Donoghue, and Rabin 2003). Loewenstein (1996) theorized that a more vivid impression of oneself engaging in some action in the future might intensify the emotions that are linked to thinking about that scenario. These intensified emotions might, in turn, allow an individual to be better informed regarding the future consequences of a present decision. For example, pulmonologists tend to smoke less than other doctors, perhaps because seeing blackened and withered lungs on a daily basis increases the negative emotions that are associated with smoking (Loewenstein 1996).

AGE-PROGRESSED EMBODIMENTS OF THE FUTURE SELF

To the extent that people can feel more connected to a vividly imagined future self, they should be motivated to save more money for the future. Accordingly, in what follows, we examine the association between seeing an embodiment of one's self in the future and the propensity to save for retirement or accept later monetary rewards over immediate ones. Whereas

philosophers and psychologists have used the term "embodied" to signify that organisms are influenced by their own physical, incarnate bodies (Lakoff and Johnson 1999), we borrow the term from the virtual reality literature, where it is associated with a particular visual or otherwise perceivable representation of a body. As pointed out by Lanier (1992), a perceivable embodiment helps forge an association with the self in a digital environment.

In the current research, we present people with renderings of their future selves made using age-progression algorithms that forecast how physical appearances will change over time. People are of course capable of imagining their future selves at any time, so why should the presentation of renderings lead to different behavior than that resulting from day-to-day imagination? We have several reasons to suspect a difference. First, though capable, people may not ordinarily elect to imagine their future selves. Second, even if people regularly imagine their future selves, they may not be imagining themselves at a distant retirement age. Third, imagination of the future may be propositional ("I will have enough money to leave to my children") rather than visual; vivid visual imagery is thought to exert strong influences on preferences and memory (Loewenstein 1996; Standing 1973). Fourth, as discussed, an imagined self may be uncertain, vague, and probabilistic, whereas a computerized rendering is in contrast definite and specific. Fifth, because people may have variable confidence in their abilities to imagine their future selves, renderings created by an objective forecasting model may be viewed as more authoritative. Lastly, as Parfit (1971) notes, neglect of the future self can arise from a failure of the imagination. Because imagination from the starting point of a graphical rendering may require less effort and attention than imagination from a blank slate, people may more easily imagine the future self when seeded with an image that is based on their present appearance (as the renderings in the experiment are). These studies were motivated, in part, by others who have

used virtual reality as a tool (Blascovich and Bailenson 2011) for influencing: consumer

behavior (Ahn and Bailenson, in press), health behavior (Fox and Bailenson 2009), financial

decision making (Yee and Bailenson 2007), and memory (Segovia and Bailenson 2009).

Our work adds to a growing body of literature that examines effective interventions for

increasing saving behavior. Whereas previous work has explored how both changing one's

environment in advance (Ariely and Wertenbroch 2002) and elaborating on potential outcomes

can lead to more patient intertemporal choices (Nenkov et al. 2008, 2009), our manipulations

take an earlier starting point. Instead of altering the way consumers think about future rewards,

the manipulations presented here seek to aid consumers in imagining the future self who will

benefit (or suffer) from the outcomes of decisions made today.

OVERVIEW

A variety of novel technologies were incorporated in the current studies, including

immersive virtual reality (VR). Blascovich and Bailenson (2011) have demonstrated the utility of

this methodology for social science, in that VR can recreate situations that are difficult to

simulate in the physical world while maximizing mundane realism and experimental control. In

Study 1, we used immersive VR to actually put participants inside of a visual representation of

their body and face as they will approximately look in the future (Yee and Bailenson 2006).

Study 2 extends the results of Study 1 to more implicit dependent variables and rules out demand

effects. Study 3A tests whether these interventions can work in field conditions, namely delivery

over the Internet, without special virtual reality hardware, and only a few user photographs as

input. Finally, Study 3B assesses the generalizability of these results using a community sample,

and also examines the extent to which the manipulations actually enhance future self-continuity.

STUDY 1: IMMERSIVE VIRTUAL REALITY AND THE FUTURE SELF

In Study 1, we used collaborative virtual environments (CVEs; see Blascovich and Bailenson 2011) to study the effects of exposure to a digital representation of one's future self. CVEs are communication systems in which multiple participants share the same three-dimensional digital space despite occupying remote physical locations. In a CVE, immersive virtual environment technology monitors the movements and behaviors of individual participants and renders those behaviors within the CVE via avatars. These digital representations are tracked naturalistically by optical sensors, mechanical devices, and cameras. In Study 1, one group of participants saw a digital representation of their current selves in a virtual mirror, and the other group of participants saw an age-morphed version of their future selves in the virtual mirror. To ensure identification with their avatars and an equal amount of time spent in the virtual reality environment, all participants had a brief conversation with a confederate. We hypothesized that those participants who saw the age-morphed version of their future selves would be more likely to allocate money toward a hypothetical retirement savings account after exposure than those who saw the current version of themselves.

Materials

Money allocation task. The money allocation task was a novel task created for the purposes of this experiment. In it, participants were told to imagine that they had just unexpectedly received $1,000 and were asked to allocate it among four options: "Use it to buy something nice for someone special," "Invest it in a retirement fund," "Plan a fun and

extravagant occasion," and "Put it into a checking account."

Age progression. We first used preset algorithms from a computer software package (Facegen Modeler; Singular Inversions 2004) which a) locates key points on the face from a front-on and profile photograph, b) builds a three-dimensional model of the face, and c) morphs the shape and texture of the model to simulate the aging process to create a persuasive visual analog of a 70 year-old version of a current college student (See Figure 1 for an example of the age-progression procedure). To create the aged photos, the age-progression algorithm of the FaceGen Modeler software package was applied with identical settings to each photo. Because the software is specifically designed to manipulate facial features and not hair, an artist next digitally retouched each image using Adobe Photoshop to change the original hair color from the front-on photograph of the participant to gray. We used an identical procedure (i.e., creating the three dimensional model) for the non-aged avatars, except that we did not use the aging algorithm nor did we gray the hair color.

Virtual reality. In the immersive virtual reality environment, the participant can enter an immersive virtual reality system and see their future self avatar in a virtual mirror (Bailenson, Beall, and Blascovich 2002). In the simulation, cameras monitor the movements and behaviors of the participant and display those behaviors via avatars (digital representations of people). Figure 2 demonstrates the CVE apparatus that is in place in Stanford University's Communication Department, with the critical hardware labeled with letters. The virtual setting was a white room that had the same exact dimensions as the physical room participants were in (see Figure 2). 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 2). Thus, the mirror image tracked and

reflected six degrees of freedom such that when the participant moved in physical space, his or her avatar moved in perfect synchrony in the mirror. The confederate's avatar was located 5 meters in front of the participant, facing the participant, and remained invisible until the conversational portion of the experiment began. Additionally, the confederate physically remained behind a curtain until the conversational portion of the experimental task (see below).

Apparatus. Perspectively correct stereoscopic images were rendered at an average frame rate of 60 Hz. The simulated viewpoint was continually updated as a function of the participants' head movements, which were tracked by a three-axis orientation sensing system. The position of the participant along the X, Y, and Z planes were tracked via an optical tracking system. Participants wore an nVisor SX head-mounted display (HMD) that featured dual 1,280 horizontal by 1,024 vertical pixel resolution panels that refreshed at 60 Hz. See Figure 2 for the equipment setup.

Emotion questionnaire. Participants were asked to rate the extent to which they felt each of 15 different emotions (Positive: accomplishment, amusement, contentment, excitement, happiness, interest, joy, and pride; Negative: anger, anxiety, disgust, fear, frustration, irritation, sadness) on a 7-point scale (From 1 = "Not at all" to 7 = "Extremely").

Method

Participants. Fifty participants (33 women; M = 20.13 years) took part in the study. Participants received course credit or $10.

Experimental procedure. Participants were randomly assigned to one of two conditions (current self or future self). Participants were told that they were going to enter the virtual reality environment and see their own (or aged) face on a digital avatar, and that they would be answering a series of personal interview questions. The experimenter then showed participants

two images of their avatar before they entered the virtual reality environment. So that

participants in both conditions saw an equal number of images, participants in the control

condition were shown a front and side view of their digital avatar, whereas participants in the

experimental condition were shown the young version of their avatar and then the age-morphed

version.

Next, participants were outfitted in the head-mounted display (HMD) and entered the

virtual environment, where they could see their digital avatar in the virtual mirror in front of

them. Participants saw themselves in a room that was exactly the same dimensions as the

physical lab room depicted in Figure 2. Following Yee and Bailenson's (2007) procedure, to

enhance identification with the avatar, the experimenter asked the participant to turn around 180

degrees and verify that they saw a mirror in front of them. After verbal affirmation, 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 60 and 75 seconds.

Participants were then asked to turn back around to face the front of the room (i.e., their

original orientation). Slightly ahead of time, the experimenter had triggered the program to make

the confederate's avatar visible to the participant in the virtual world. 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 (e.g., "What is your name," "Where are you from?," "What

is your passion in life?"). These interview questions were included to enhance identification

between the participant and his or her avatar (Yee and Bailenson 2007). Their behaviors were not

scripted, and they were instructed to use natural head movements when interacting with the

participant.

After the interview questions were asked, the experimenter removed the participant's HMD and then the participant completed the money allocation task, the emotion questionnaire, and filler tasks for another experiment. The participant was then paid and debriefed.

Results and Discussion

We had hypothesized that participants in the future self condition would allocate more money toward retirement than would participants in the current self condition. In line with this prediction, participants who were exposed to their future selves in virtual reality allocated more than twice as much money toward the retirement account (M = $172; SD = $214) than did participants who were exposed to their current selves (M = $80; SD = $130), t(48) = 1.83, *p* = .035 (one-tailed), d = .519.

It is possible, though, that seeing one's aged face may have depressed one's current mood. If this was in fact the case, the appeal of the short-term options (e.g., planning a fun and extravagant occasion) in the monetary allocation questionnaire may have been diminished not because of enhanced connectedness to the future self, but because these amusing options are less appealing when experiencing negative moods (though see work by Cryder, Lerner, Gross, and Dahl (2008) which demonstrates that sadness leads to *more* spending). Such a scenario would leave a larger pot of money to be devoted to the long-term saving option. However, results indicate that participants in the future self condition (M = 2.52; SD = .95) did not exhibit significantly stronger negative emotions (M = 2.62; SD = 1.27) than participants in the current self condition, t(48) = .33, *p* = .75. Further, when controlling for the average of all negative emotions, condition is still a predictor of amount of money allocated toward retirement, F(1, 47) = 3.19, *p* = .04 (one-tailed).

Thus, results suggest that interaction with a vivid version of one's future self causes individuals to give modestly greater weight to long-term saving. It is possible, however, that demand characteristics of the study played a role in participants' behavior (Orne 1962). Although informal post-experiment interviews suggested that this was not the case, in Study 2, we attempted to formally rule out these demand characteristics.

STUDY 2: RULING OUT DEMAND CHARACTERISTICS

In Study 2, we sought to rule out two alternative explanations for our findings from Study 1. First, because the monetary allocation task occurred directly after the virtual reality paradigm, it is possible that participants in the experimental condition felt pressure to allocate more to the long-term account. Thus, in Study 2, we separated the virtual reality portion of the study from the decision-making portion, and provided a cover story that masked our research purposes. Secondly, it is possible that experimental condition participants in Study 1 were merely primed with the concept of aging, and this prime prompted them to save more for retirement (i.e., Bargh and Chartrand 1999). As such, in Study 2, we exposed participants to either their own aged avatar or another research participant's aged avatar. Finally, in Study 2, we sought to extend our findings to three different dependent variables. Namely, we used a short-term temporal discounting task, a long-term temporal discounting task, and a retirement spending questionnaire. Including all three measures allowed us to test the degree to which our manipulations increase a focus on the future self for all types of intertemporal choices rather than just long-term ones. Even though short-term and long-term discounting may operate on different processes (McClure, Laibson, Loewenstein and Cohen 2004), opting for larger later rewards in both cases might

nonetheless be linked to better saving behavior. That is, by repeatedly foregoing short-term rewards one could inevitably end up saving more over the long run (Thaler and Shefrin 1981); money not spent on pleasurable experiences in the present can be used for other purposes in the future. We thus hypothesized that participants who interacted with avatars depicting their future selves, as opposed to future others, would demonstrate a greater willingness to choose larger, future rewards on all three tasks.

Materials

Short-term temporal discounting task. The temporal discounting task consisted of 21 choice trials (Kirby and Marakovic 1996). Each trial included one smaller immediate reward paired with one larger delayed reward. Immediate values ranged from $15 to 83, whereas the delayed values ranged from $30 to $85 over delays of 10 to 75 days. The task was incentive compatible: participants were instructed that at the end of the experimental session, one response would be chosen at random and participants would be paid 33% of that choice at the appropriate delay (e.g., If on a given question a participant chose $30 in 27 days, they would receive $10 in 27 days). The number of delayed choices was counted to index discount rate (Magen, Dweck, and Gross 2008). Although the Kirby and Marakovic (1996) procedure excludes data from individuals who either chose all of the immediate or all of the delayed options, the Magen et al. (2008) procedure has the advantage of retaining these individuals for subsequent analysis.

Long-term temporal discounting task. The long-term temporal discounting task was created for the purposes of this experiment and was based on the Kirby and Marakovic (1996) short-term temporal discounting task. Again, each trial included one smaller immediate reward paired with one larger delayed reward. As in Kirby and Marakovic (1996), we used a hyperbolic discounting function:

(1) $V = A / (1 + kD)$

where *V* is the value of a present (or immediately available) gain, *A* is the amount of a future

gain, *k* is a discount parameter that varies across individuals, and *D* is the amount of time that

individuals must wait for the future gain (Mazur 1987). Discount rates (*k*) ranged from .07 to .86.

Immediate values ranged from $2,575 to $25,840, while the delayed values ranged from $76,965

to $98,191 over delays of 35 to 40 years. Again, the number of delayed choices was counted to

index the discount rate.

      Retirement Spending Questionnaire. The Retirement Spending Questionnaire

(Binswanger and Carman 2009) consisted of one question that asked participants to choose

between six options of how they would like to spend the money they earned during their lifetime.

Each option detailed the amount of money one would want to spend monthly during one's

working life and during retirement. For example, Option 1 detailed a case in which one could

spend $2,950 per month during one's working life and $1,900 per month during retirement. At

the other end of the spectrum, Option 6 detailed a case in which one could spend $2,600 per

month during one's working life and $3,600 per month during retirement. Thus, higher scores on

this scale were associated with greater allocation toward retirement. Participants were explicitly

told to answer the question as if prices remained constant (i.e., as if there was no inflation).

Method

      Participants. Twenty-one participants (15 women; M = 20.08 years) took part in the

study. Participants received $10 pay and were told that the experiment would take approximately

an hour to complete.

      Experimental procedure. Participants were assigned to be in one of two conditions

(future self or future other), and were then given the following instructions. Participants were

told that we were experimenting with technology that would put their own aged face (or someone else's aged face) on a digital avatar. Next, participants were brought into the same immersive virtual reality environment from Study 1, where they could see either their digital avatar in a virtual mirror in front of them (future self condition) or the aged digital avatar of another same-sex and same-race participant (future other condition). As in Study 1, Yee and Bailenson's (2007) procedure to enhance identification with the avatar was conducted.

Participants were then allotted three minutes to talk about the ways that they were similar (in terms of personality, appearance, temperament, major preferences, beliefs, values, ambitions, etc.) to the person depicted by the avatar in the mirror. After three minutes had elapsed, the experimenter took the HMD off of the participant and then the participant completed several short filler questionnaires asking them about their experience in the virtual reality environment. To eliminate demand characteristics, we ostensibly separated the manipulation from our main dependent variables. Namely, after completing these questionnaires, a computer message appeared notifying participants that they had completed the experimental session. At this point in the study, however, approximately 45 minutes had elapsed, and participants had been led to believe that they were participating in an hour-long experiment at a pay rate of $10 per hour. The experimenter feigned surprise at the early completion time and told participants that they were free to go but in order to receive the full $10 they needed to be present for a full hour of research participation. Participants were then told that there were a few other short surveys that were being conducted in the lab by other researchers in which they could participate to earn the remaining pay. All participants agreed to participate in the additional survey. To further mask our research purposes, participants were ostensibly given a choice of three different surveys in which to participate. All three surveys, however, linked to the same set of questionnaires: the

short-term temporal discounting task, the long-term temporal discounting task, and the

retirement spending questionnaire. Upon completing these tasks, participants were debriefed and

paid.

Results and Discussion

To measure the unique impact of interacting with one's future self on saving behavior,

we conducted a repeated-measures ANOVA on standardized scores from the three dependent

variables (Short-term temporal discounting task: $M_{\text{Future Self}} = 0.36$, $SD_{\text{Future Self}} = 1.04$ and M_{Future}

$_{\text{Other}} = -0.32$, $SD_{\text{Future Other}} = 0.89$.; Long-term temporal discounting task: $M_{\text{Future Self}} = 0.27$, SD

$_{\text{Future Self}} = 0.83$ and $M_{\text{Future Other}} = -0.25$, $SD_{\text{Future Other}} = 1.11$; Retirement Spending Questionnaire:

$M_{\text{Future Self}} = 0.28$, $SD_{\text{Future Self}} = 0.99$ and $M_{\text{Future Other}} = -0.25$, $SD_{\text{Future Other}} = 0.99$). Results

indicated that participants in the future self condition exhibited more saving behavior across the

three tasks than did participants in the future other condition, $F(1, 19) = 4.14$, $p = .056$.

Thus, Study 2 employed a well-masked paradigm in which we directly compared the

effects of interacting with one's imagined future self to interacting with another imagined older

person. In line with our prediction, vivid exposure to one's future self in an immersive virtual

reality environment – compared to exposure to another person's older avatar – led to increased

saving in both short- and long-term decision-making tasks.

In both Studies 1 and 2, we presented age-progressed renderings of the future self, but did

not highlight the consequences of decisions that are made in the present. Furthermore, although

these findings are encouraging conceptually, because of high costs and time efficiency, most

companies will not be able to use immersive virtual reality to convince their employees to

contribute additional money to retirement accounts. Accordingly, in Study 3A we sought to

highlight future consequences of present decisions and translate our methods and findings to

more accessible formats for widespread use as an intervention.

STUDY 3A: GENERALIZATION TO FIELD CONDITIONS

Whereas the previous study aimed to rule out demand effects, this study takes an applied perspective and gives the virtual renderings emotional qualities in order to intentionally exert a demand or "nudge" in the context of an online decision aid. The decision aid provides users with accurate estimates of present and future spending levels that will be attainable given various saving rates. Its estimates, described in the Appendix, take inflation, earnings growth, employer 401k contribution match, age, social security, and many other factors into account. The information given is so useful that it alone could serve as a basis for prudent savings decisions, and providing participants with it should greatly reduce their uncertainly about how saving will affect their material wealth in the short and long term. Owing to this, the present design allows us to mitigate the effect of income uncertainty when estimating the persuasive effect of the virtual renderings. The second objective of this study is to test whether the kind of intervention proposed in this article can be practically adapted for widespread Internet participation from home, office, or elsewhere. To conduct this study, no virtual reality hardware was necessary, and people needed to only submit three digital photos to take part.

The ability to vividly imagine not just the face of the future self, but also the emotional reactions of the future self, may affect the willingness to save for the future. Virtual reality research has demonstrated that exposure to virtual cause-and-effect actions can change actual behavior. For example, participants shown virtual versions of themselves losing weight with exercise and gaining weight with inactivity were more likely than controls to exercise afterwards

(Fox and Bailenson 2009). To test whether exposure to an emotional future self can affect the propensity to save for retirement, participants in two conditions interacted with current or age-progressed renderings of themselves that appeared to respond emotionally to changes in income.

Materials

Age progression. In the first phase of the study, participants came in person to the laboratory where they were photographed making happy, sad, and neutral facial expressions, after which they were instructed that they would carry out the remainder of the study online.

Using the three photos as a starting point, we prepared three age-progressed photos (for participants in the future self condition) or three unaged photos (for participants in the current self condition). Each set of three photos served as the basis for 11 emotional variant photos in two sets. A "future" set of 11 photos showed the participants as they might look as older adults (approximately age 65), with levels of facial expression ranging from sad to happy. The "present" set had the same 11 levels of facial expression but showed the participants at their current age.

To create the three aged photos, we used the same procedure from Study 1, employing the age-progression algorithm of the FaceGen Modeler software package and Adobe Photoshop to change the hair color to gray. The 11 emotional variant photos within the "present" and "future" sets were created using Fantamorph software to make "morphs" of the sad, neutral, and happy input images. A morph is akin to a weighted average of two images and looks quite natural (Bailenson, Iyengar, Yee, and Collins 2008). The weighting scheme in Table 1 was applied to create the morphs. This process was the same for making the present and future photos, the only difference being the three starting images used.

Retirement allocation slider bar. In both conditions, participants made a decision about

how much they would like to contribute toward their retirement fund. To help in making this decision, they were given a slider that controls their level of contribution. Instead of showing participants the percentage of salary saved at each slider position (which is not particularly informative for making such decisions), as they explored various levels, the interface indicated how much it would impact their current and future income. In Figure 3, the left-hand percentage displays current income as a percentage of salary remaining after contributions to Social Security and Medicare. The right-hand percentage value displays income in retirement expressed as a percentage of income in the year immediately before retirement. Both percentages updated dynamically as the cursor was moved over the slider.

The values represented on the slider were calculated in the same manner in both conditions, and are based on a detailed financial model summarized in the Appendix. The inputs given to the model were the same for all participants and chosen to be representative for students at the university where the studies were conducted. In particular, the model made its calculations based on an individual aged 22, starting work at age 23, retiring at age 68, and earning a salary of $64,000 (in today's dollars) by age 40. Allowable percentages of salary that could be saved toward retirement ranged from 0% to 10%. These parameters were used to generate realistic tradeoffs (in percentage terms) for typical American college students, and were never seen by the participants.

Figure 3 depicts the interface in the "present" and "future" conditions, and Figure 4 gives an indication of how the movement of the slider bar affected the emotions of the future face; the behavior of the *present* face can be inferred from what follows. When the present face is displayed, it becomes happier as the slider is moved left (and present income increases) and sadder as the slider is moved right (and present income decreases). However, in the future self

condition, the future self face becomes sadder as the slider is moved left (and retirement income decreases) and happier as the slider is moved right (and retirement income increases). The two percentages are both always visible to participants in both conditions. Because these percentages provide detailed information about the tradeoffs of saving, and should reduce participants' uncertainty about the material consequences of saving, estimates of treatment effects should be conservative and relatively isolated from the effect of uncertainty. In practice, savers are not able to estimate their future income levels as easily as with the decision aid, so the persuasive effect of the faces may also be conservative relative to what would be obtained in the field. In addition, to control for other factors which might bias the results, we asked participants to complete two scales:

MacArthur Scale of Subjective Social Status. The MacArthur Scale of Subjective Social Status (Adler, Epel, Castellazzo, and Ickovics 2000) was used to capture participants' socioeconomic status (SES). The scale is in a pictorial format and presents a 10-runged "social ladder" and asks participants to click on the rung on which they feel they stand in relation to other people in the U.S. Scores on the ladder have been predicted by employment grade, education, household and personal income, household wealth, satisfaction with standard of living, and feelings of financial security (Adler et al. 2008) (M = 7.08, SD = 1.70).

Perceived income stability. To assess participants' beliefs about their future careers, we asked them to report how stable they thought the income flow would be in their future careers on a 5-point scale (M = 3.83, SD = .90).

Method

Participants. Participants were 42 students (22 women; Age: M = 21.00 years) who were recruited from a university subject pool who agreed to participate in a study on decision-making

for a $15 Amazon.com gift certificate. Participants were randomly assigned to one of two groups: Future Self (n = 21) and Current Self (n = 21).

   Procedure. In the second phase of the study, three to four weeks after the participants were photographed, all participants received an email inviting them to complete the study online. The URLs in the emails were customized so that participants were assured to see their own photos. Participants were given instructions on how the slider bar would affect current and future income and how to interpret the numbers at each end of the slider, after which they read a variant of the following sentence modified to reflect whether they would see their present face or future face: "Additionally, you will see a virtual image of yourself *now* (at retirement age), to clarify how the movements of the slider affect your *income while you are working* (future income)." The next web page contained a four-alternative multiple-choice question testing comprehension of the relationship between the slider bar movement and the amounts of present and future income. Those failing to answer the question correctly were directed back to the instruction pages until they passed the test.

   At this point, participants were presented with the decision aid. To avoid anchoring effects, the slider did not have a visible handle set to any point. Instead, a vertical line intersecting the slider appeared when the mouse hovered over the slider bar. The hovering point of the mouse determined the continuously-updated current and retirement income percentages as well as the facial expression (from happy to sad), according to condition. When participants clicked on the slider bar, a gray line appeared to mark their last active choice and subjects were given the option to submit this choice or to continue exploration. Upon submission, the slider bar position and sequence of clicks was recorded to the database.

   After the retirement allocation task, participants completed a demographics questionnaire

(in which they answered questions about sex, age, and race/ethnicity), the MacArthur Scale of Subjective Social Status, and the perceived income stability item. Participants were debriefed via email.

Results and Discussion

Retirement allocation. We compared the effect of condition (current self or future self) on retirement allocation[1]. In line with our prediction, participants in the future self condition allocated a significantly higher percentage of pay toward retirement (M = 6.76%; SD = 1.68%) than did participants in the current self condition (M = 5.20%; SD = 2.35%), t(38) = 2.38, *p* = .023, d = .771. Although the effect size for this result is medium, the difference between conditions of 1.56 percentage points is practically quite significant.

Potential moderators. Because previous research has demonstrated that both SES (Grable 2000; Henry 2005) and income stability (Bajtelsmit and Bernasek 2001) affect retirement allocation decisions, we independently tested the moderating role of these variables. Following Judd and Kenny's (2010) recommendation, we first centered all predictor variables, and created interaction terms by taking the product of condition and SES and the product of condition and perceived stability. Results indicated that neither SES (Condition X SES interaction, β = .24, t(36) = .99, *p* = .33) nor perceived income stability (β = -.15, t(36) = -.83, *p* = .41) were significant moderators, and Condition remained a significant positive predictor of retirement allocation when controlling for SES, β = .36, t(36) = 2.33, *p* = .025, and perceived income stability, β = .36, t(36) = 2.27, *p* = .029.

STUDY 3B: STRONGER TEST OF GENERALIZATION TO FIELD CONDITIONS

In Study 3B, we sought to rule out an alternative explanation for the findings from Study 3A. Namely, it is possible that instead of results being due to exposure to the future self, they are due to participants reacting to the valence the different faces present. Rather than being motivated to save more by the presence of the future self, participants could have merely been moving the retirement slider bar toward whichever face was smiling. To correct for this potential confound, we conducted an additional experiment that was identical to Study 3A except that instead of being presented with emotional images of their current or future selves, participants were exposed to neutral versions of these selves (i.e., faces that neither frowned nor smiled).

Study 3B had two additional purposes. First, to ensure that the results from the first three studies were generalizable, participants in Study 3B were adults from a national online pool (as opposed to undergraduates from a university). Second, we examined a potential mechanism that could underlie the relationship between exposure to the future self and enhanced saving behavior. Because impoverished intertemporal decision-making has been linked to a lack of continuity with one's future self (Ersner-Hershfield, Garton et al. 2009; Bartels and Rips 2010), we assessed whether the present manipulation would actually boost continuity with that self. To do so, we gave participants the Future Self-Continuity Scale developed by Ersner-Hershfield, Garton et al. (2009). We hypothesized that participants who were exposed to images of their future selves would allocate more current income towards retirement and report a greater sense of continuity with their future selves. We further hypothesized that future self-continuity would act as a mediator between experimental condition and retirement allocation.

Method

Future Self-Continuity Scale. The Future Self-Continuity Scale, which was based on

Aron, Aron, and Smollan's (1992) Inclusion of the Other in the Self Scale, measures the degree

to which participants feel similar to their future selves. Participants pick a pair of Euler circles

(out of a possible 7 pairs) that best represents how similar they feel to their future selves in ten-

years time (See Ersner-Hershfield, Garton et al. (2009) for a more detailed description). Higher

scores indicate more continuity with one's future self.

Participants. Participants were 40 adult of US nationality aged 18 to 35 (22 women; Age:

M = 26.27 years) who were recruited from the Mechanical Turk online community and agreed to

participate in a study on virtual reality in exchange for $11. The study comprised two parts. In

the first, potential participants submitted photographs of themselves holding up a sign that

displayed the current date and time. This was done to help verify that the photos were indeed

those of the participants and not stock photos. A few weeks later, in the second part, the 56

participants who submitted acceptable photos were randomly assigned either the Future Self or

Current Self groups. The response rates did not differ significantly by condition (68% in the

future self condition vs 75% in the current self condition), $\chi^2(1, N = 40) = .35$, $p = .55$. During

the experiment, we administered a question to ensure that participants understood the task

instructions and one participant from each condition failed this test. Our final sample comprised

18 participants in the Future Self condition and 20 in the Current Self condition.

Procedure. The experimental procedure was identical to that of Study 3A except for three

key differences. First, participants uploaded their photos instead of being photographed in

person. This should help assure that a wide range of photos can be used for this type of treatment

in the field. Second, and most importantly, instead of seeing an image of their current or future

self that changed emotional expression as a function of their retirement allocation decision,

participants in this study saw a fixed image of their current or future self that maintained a

neutral expression. That is, as participants altered their retirement allocations, the facial expression on the image of their current or future self (depending on the condition) did <u>not</u> change. Finally, participants also completed the Future Self-Continuity Scale. Would the result of Experiment 3A replicate 1) without the faces taking on emotional expressions, 2) with photos taken under a variety of conditions, and 3) with a national sample of participants instead of a sample of undergraduates from one university?

Results

We compared the retirement allocations of the present- and future-self conditions. In line with our prediction, participants in the future self condition allocated a significantly higher percentage of pay toward retirement (M = 6.17%; SD = 2.30%) than did participants in the current self condition (M = 4.41%; SD = 2.44%), t(36) = 2.28, p = .029, d = .742. As in Studies 1, 2, and 3A, exposure to an age-progressed rendering of the future self led to an increase in saving behavior.

Participants in the future self condition also reported a higher level of continuity with their future selves (M = 4.39, SD = 1.61) than did participants in the current self condition (M = 3.25, SD = 1.92), t(36) = 1.97, p = .057, d = .64. Further, we find evidence for mediation: when regressing allocation percentage on condition and future self-continuity, the direct effect from condition to retirement becomes non-significant (β = .235, t(35) = 1.54, p = .134) while future self-continuity remains a significant predictor of retirement allocation (β = .384, t(35) = 2.51, p = .017). Preacher and Hayes' (2008) bootstrapping procedures established that this mediation was indeed significant—the bias-corrected confidence interval (CI) of the bootstrapping mediation test did not include zero ($CI_{95\%}$ = .0757, 1.5731; N = 38; 10,000 re-samples).

GENERAL DISCUSSION

The present work represents the first demonstration of a new kind of intervention in which people can be encouraged to make more future-oriented choices by having them interact with age-progressed renderings of their own likenesses. In four separate studies, using both undergraduates and community members as participants, we provide evidence that manipulating exposure to visual representations of one's future self leads to lower discounting of future rewards and higher contributions to saving accounts. Moreover, Study 2 shows that our effects are not due to thinking about aging *per se*, nor due to demand effects, but rather, can arise simply from direct exposure to renderings of the future self. Studies 3A and 3B demonstrated that effects of this type may be able to translate to the field at low expense. The findings were modest with medium to large effect sizes (Cohen 1988) and were consistent across the four studies in the predicted direction. We set out to explore and test a novel methodology, and results should be interpreted as preliminary, but just the same, given the American situation of low saving rates and increasing life expectancy, our results are encouraging and have a managerial implication: when making important long-term decisions, vivid representations of one's future self should increase the future-orientation of saving decisions.

This visual intervention may in some sense address the failures of belief or imagination of which Parfit (1971) wrote. Study 3B demonstrated that exposure to an image of one's future self does increase continuity with that future self. However, we still do not fully understand all of the underlying psychological processes that are affected by our manipulations, and there is much research to be done on this topic. It is possible that exposure to the future self causes more parity in emotional experiences between the current self and the future self. As noted earlier, research

on the hot-cold empathy gap (Loewenstein, O'Donoghue, and Rabin 2003) has suggested that the emotions an individual feels in the present are much stronger than the emotions that the same individual *expects* to feel in the future. Such empathy gaps are, of course, not borne out by experience: most similarly valenced emotional occurrences take on a similar profile whether they occur today or in five years. One reason, then, that people make unwise intertemporal choices is because they erroneously allocate too much weight to the feelings of the seemingly more emotional current self and not enough to the future self. Thus, our manipulation could work by helping consumers to recognize that the future self will be just as emotional as the current self is. Alternatively, the effectiveness of our manipulation may rest in its ability to force consumers to temper the emotions they feel in the present to make them more on par with those that they expect to feel in the future.

It is also possible that our manipulations may be successful in part because they are considerably more engaging than traditional interventions. To the extent that consumers are more likely to follow through on self-control tasks when they are fun and engaging (Laran and Janiszewski in press), manipulations such as ours that rely on age-progression visualizations may, in addition to enhancing connectedness to the future self, also make thinking about saving for retirement a more entertaining and compelling endeavor. A better understanding of the underlying process will no doubt inform future interventions aimed at helping people save more for retirement.

Limitations and Future Directions

Hypothetical exercises. We used novel technology – immersive virtual reality and realistic age-progression software – to assist the imagination. The theoretical foundation upon which we based our hypotheses, however, does not necessarily suggest that the people must *see*

representations of their future selves. Parfit (1987), Schelling (1984), and Loewenstein (1996) simply claimed that future focus will increase to the extent one is aware that one's future self is a living, breathing, individual who is dependent on the choices of the current self. We predicted that using immersive virtual reality and realistic age-morphing software would be a particularly compelling and effective method because it does not rely on differences in imagination abilities; age-progression software and immersive virtual reality can make the future self more realistic regardless of one's ability to envision the future. Nevertheless, it is possible that hypothetical exercises that bring the future self into one's mind (e.g., composing a letter to the future self) might also be effective at increasing long-term saving behavior. We have discussed how imagination differs from interacting with renderings, and future research could test whether such differences matter.

Temporal duration of effects. In all four studies, participants made intertemporal choices either during or shortly after exposure to their future selves (or current selves). As such, the temporal duration of our manipulations remains unknown. However, we believe that the most important practical application of future self saliency manipulations occurs in the actual decision-making moment. Indeed, when an employee decides how much of her paycheck she would like to allocate to a 401(k), this is a decision that is normally made once and rarely revisited (Benartzi and Thaler 2007; Madrian and Shea 2001; Samuelson and Zeckhauser 1988). Thus, the moment when one makes an intertemporal choice is arguably the most important time to consider the interests of the future self. To this end, in ongoing work, we are exploring the effectiveness of using the slider bar paradigm from Studies 3A and 3B to aid actual employees when they are making their retirement decisions. This endeavor will help us in determining whether decisions made on the hypothetical retirement allocation task from Studies 3A and 3B

translate to incentive-compatible rewards (though previous research has demonstrated that discounting rates of real versus hypothetical rewards show similar patterns (Coller and Williams 1999)).

Age differences in effectiveness of manipulation. Similarly, in the present studies, we used a relatively young sample of research participants. This choice was deliberate, as the optimal time to start saving for retirement is early in life (Bernheim, Skinner, and Weinberg 2001). It is unclear, however, whether vivid exposure to the future self would hold such powerful effects for older consumers (Yoon, Cole, and Lee 2009).

Generalizability of novel technology. One final limitation of the current work is that we are relying on technology that is not in everyone's living room or displayed in financial institutions across the world. However, the use of avatars and virtual worlds is accelerating greatly. Game platforms such as the Nintendo *Wii* and Microsoft's new *Kinect* are transforming the realm of home entertainment into an immersive experience, and due to the recent success of three-dimensional movies such as *Up* and *Avatar*, new stereo displays will be available for consumers to purchase this year (Blascovich and Bailenson 2011). Technology's development is accelerating, and the notion of people interacting via avatars with their bank accounts or financial planners is not an outlandish one. The intervention in Studies 3A and 3B requires only a camera and a connection to the internet, both of which can be found on many smartphones and other portable web-enabled devices. Similar face-rendering applications are already in existence and require the user to upload a photo and wait approximately 10 seconds for a result. On the iPhone platform, for example, one application allows a user to upload a photo taken with his or her phone and then instantly see what he or she would look like after gaining roughly 100 pounds. Nonetheless, given the general difficulty that most employers have in encouraging

employees to sign up for benefits programs, implementing the interventions from Studies 3A and 3B on a mass-scale would by no means be a simple endeavor.

Conclusion

Due to the confluence of increasing life expectancy, a weakened economy, and a large number of underprepared households, getting consumers to save more for retirement is at this time an issue of great economic importance. To complement interventions that emphasize or deemphasize the appeal of present and future rewards, the present work operates on another distinction, that between present and future selves, to increase saving for the future.

REFERENCES

Adler, Nancy E., Elissa S. Epel, Grace Castellazzo, and Jeanette R. Ickovics (2000), Relationship of subjective and objective social status with psychological and physiological functioning: Preliminary data in healthy white women. *Health Psychology, 19*(6), 586-592.

––– , Archana Singh-Manoux, Joseph Schwartz, Judith Stewart, Karen Matthews, and Michael G. Marmot (2008), Social status and health: A comparison of British civil servants in Whitehall-II with European- and African-Americans in CARDIA. *Social Science Medicine, 66*(5), 1034-1045.

Ahn, Sun J., and Jeremy N. Bailenson, (in press), Self-endorsing versus other-endorsing in virtual environments: The effect of brand attitude and purchase intention. *Journal of Advertising*.

Ainslie, George (1975), Specious reward: Behavioral theory of impulsiveness and impulse control. *Psychological Bulletin, 82*(4), 463-496.

Arias, Elizabeth (2007), United States Life Tables, 2004. *National Vital Statistics Reports, 56*(9), 1-40.

Ariely, Dan, and Klaus Wertenbroch (2002), Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science, 13*(3), 219-224.

Bailenson, Jeremy N., Andrew C. Beall, and Jim Blascovich (2002), Mutual gaze and task performance in shared virtual environments. *Journal of Visualization and Computer Animation, 13*, 1-8.

––– , Shanto Iyengar, Nick Yee, and Nathan A. Collins (2008), Facial similarity between voters and candidates causes social influence. *Public Opinion Quarterly, 72*(5), 935-961.

––– and Kathryn Y. Segovia (2010), Virtual doppelgangers: Psychological effects of avatars who ignore their owners. In W. S. Bainbridge (Ed.), *Online worlds: Convergence of the real and the virtual* (pp. 175-186), New York, NY: Springer.

Bajtelsmit, Vickie L., and Alexandra Bernasek (2001), *Risk preferences and the investment decisions of older Americans, AARP Public Policy Institute report no. 2001-11*. Washington, DC: American Association of Retired Persons.

Bargh, John A., and Tanya L. Chartrand (1999), The unbearable automaticity of being. *American Psychologist, 54*(7), 462-479.

Bartels, Daniel M., and Lance J. Rips (2010), Psychological connectedness and intertemporal choice. *Journal of Experimental Psychology-General, 139*(1), 49-69.

––– and Oleg Urminsky (in press), On intertemporal selfishness: The perceived instability of identity underlies impatient consumption. *Journal of Consumer Research*.

Benartzi, Shlomo, and Richard Thaler (2007), Heuristics and biases in retirement savings behavior. *Journal of Economic Perspectives, 21*(3), 81-104.

Bernheim, B. Douglas, Jonathon Skinner, and Steven Weinberg (2001), What accounts for variation in retirement wealth among U.S. households? *The American Economic Review, 91*(4), 832-857.

Binswanger, Johannes, and Katherine G. Carman (2009), How real people make long-term decisions: The case of retirement preparation. *Discussion Paper 2009-73*.

Blascovich, Jim, and Jeremy N. Bailenson (2011), *Infinite Reality*. New York, NY: William Morrow.

Butler, Joseph (1736), *The analogy of religion, natural and revealed, to the constitution and course of nature : To which are added two brief dissertations: I. Of personal identity. II. Of the nature of virtue*. London: Printed for James John and Paul Knapton.

Chan, Sewell, and Javier C. Hernandez (2010, April 7), Bernanke says nation must take action soon to shape fiscal future. *The New York Times,* p. B1.

Chapman, Gretchen B. (1996), Temporal discounting and utility for health and money. *Journal of Experimental Psychology-Learning Memory and Cognition, 22*(3), 771-791.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick (2006), Saving for retirement on the path of least resistance. In E. J. McCaffrey and J. Slemrod (Eds.), *Behavioral Public Finance: Toward a New Agenda* (pp. 304-351), New York: Russell Sage Foundation.

Cohen, Jacob (1988), *Statistical power analysis for the behavioral sciences (2nd ed),* New Jersey: Lawrence Erlbaum.

Coller, Maribeth, and Melonie B. Williams (1999), Eliciting individual discount rates. *Experimental Economics, 2(2)*, 107-127.

Cryder, Cynthia E., Jennifer S. Lerner, James J. Gross, and Ronald E. Dahl (2008), Misery is not miserly: Sad and self-focused individuals spend more. *Psychological Science, 19*(6), 585-530.

Diamond, Peter, and Botond Köszegi (2003), Quasi-hyperbolic discounting and retirement. *Journal of Public Economics, 87(9-10)*, 1839-1872.

Elster, Jon (1977), Ulysses and Sirens: A theory of imperfect rationality. *Social Science Information, 16*(5), 469-526.

Ersner-Hershfield, Hal, M. Tess Garton, Kacey Ballard, Gregory R. Samanez-Larkin, and Brian Knutson (2009), Don't stop thinking about tomorrow: Individual differences in future self-continuity account for saving. *Judgment and Decision Making, 4*(4), 280-286.

—— , G. Elliott Wimmer, and Brian Knutson (2009), Saving for the future self: Neural measures of future self-continuity predict temporal discounting. *Social Cognitive and Affective*

Neuroscience, 4(1), 85-92.

Fox, Jesse, and Jeremy N. Bailenson (2009), Virtual self-modeling: The effects of vicarious reinforcement and identification on exercise behaviors. *Media Psychology, 12*(1), 195-209.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue (2003), Time discounting and time preference: A critical review. In G. Loewenstein, D. Read and R. F. Baumeister (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (pp. xiii, 569), New York: Russell Sage Foundation.

——— , Nathan Novemsky, Jing Wang, Ravi Dhar, and Stephen Nowlis (2009), Opportunity cost neglect. *Journal of Consumer Research, 36*(4), 553-561.

Goldstein, Dan G., Eric J. Johnson, and William F. Sharpe (2008), Choosing outcomes versus choosing products: Consumer-focused retirement investment advice. *Journal of Consumer Research, 35*(10), 440-456.

Grable, John E. (2000), Financial risk tolerance and additional factors that affect risk taking in everyday money matters. *Journal of Business and Psychology, 14*(4), 625-630.

Henry, Paul C. (2005), Social class, market situation, and consumers' metaphors of (dis)empowerment. *Journal of Consumer Research, 31*(4), 766-778.

Ho, Teck H., Noah Lim, and Colin F. Camerer (2006), Modeling the psychology of consumer and firm behavior with behavioral economics. *Journal of Marketing Research, 43(3)*, 307-311.

Hoch, Stephen J., and George Loewenstein (1991), Time-inconsistent preferences and consumer self-control. *Journal of Consumer Research, 17*(4), 492-507.

Judd, Charles M., and David A. Kenny (2010), Data analysis in social psychology: Recent and recurring issues. In S. T. Fiske, D. T. Gilbert and G. Lindzey (Eds.), *The handbook of social psychology* (5th ed., Vol. 1, pp. 115-139), Boston, MA: McGraw-Hill.

Kirby, Kris N., and Nino Marakovic (1996), Modeling myopic decisions: Evidence for hyperbolic delay-discounting within subjects and amounts. *Organizational Behavior and Human Decision Processes, 64*(1), 22-30.

Laibson, David (1997), Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics, 112*(2), 443-477.

——— , Andrea Repetto, and Jeremy Tobacman (1998), Self-control and saving for retirement. *Brookings Papers on Economic Activity(1),* 91-196.

Lakoff, George, and Mark Johnson (1999), *Philosophy in the flesh: The embodied mind and its challenge to Western thought*. New York City: Basic Books.

Lanier, Jaron (1992), Virtual Reality: The promise of the future. *Interactive Learning International, 8*(4), 275-279.

Laran, Juliano, and Chris Janiszewski (in press), Work or fun?: How task construal and completion influence regulatory behavior. *Journal of Consumer Research*.

Lee, Chulhee (2001), The expected length of male retirement in the United States, 1850-1990. *Journal of Population Economics, 14*(4), 641-650.

Loewenstein, George (1996), Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes, 65*(3), 272-292.

——— , Ted O'Donoghue, and Matthew Rabin (2003), Projection bias in predicting future utility. *Quarterly Journal of Economics, 118*(4), 1209-1248.

Low, Douglas K., Mijung Yoon, Brent W. Roberts, and James Rounds (2005), The stability of interests from early adolescence to middle adulthood: A quantitative review of longitudinal studies. *Psychological Bulletin, 131(5)*, 713-737.

Lynch, John G., Richard Nettemeyer, Stephen A. Spiller, and Alessandra Zammit (2010), A generalizable scale of propensity to plan: The long and the short of planning for time and money. *Journal of Consumer Research, 37*(6), 108-128.

Madrian, Brigitte C., and Dennis F. Shea (2001), The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics, 116*(4), 1149-1187.

Magen, Eran, Carol S. Dweck, and James J. Gross (2008), The hidden zero effect: Representing a single choice as an extended sequence reduces impulsive choice. *Psychological Science, 19*(7), 648-649.

Mazur, James E. (1987), An adjusting procedure for studying delayed reinforcement. In J. Commons, J. E. Mazur, J. A. Nevin and H. Rachlin (Eds.), *The effect of delay and of intervening events on reinforcement value* (Vol. 5, pp. 55-73), Hillsdale, NJ: Lawrence Erlbaum.

McClure, Samuel M., David I. Laibson, George Loewenstein, and Jonathan D. Cohen (2004). Separate neural systems value immediate and delayed monetary rewards. *Science, 306*(5695), 503-507.

McKinsey Global Institute. (2008), *Talkin' 'bout my generation: The economic impact of aging U.S. baby boomers*. San Francisco, CA: McKinsey and Co.

Munnell, Alicia H., Anthony Webb and Francesca Golub-Sass (2007), *Is there really a retirement savings crisis?: An NRRI analysis* (Vol. 7-11), Chestnut Hill, MA: Center for Retirement Research, Boston College.

——— , ——— , ——— (2009), *The national retirement risk index: After the crash*. Chestnut Hill, MA: Center for Retirement Research, Boston College.

Nenkov, Gergana Y., Jeffrey J. Inman, and John Hulland (2008), Considering the future: The conceptualization and measurement of elaboration on potential outcomes. *Journal of Consumer Research, 35*(1), 126-141.

——— , ——— , ——— , and Maureen Morrin (2009), The impact of outcome elaboration on susceptibility to contextual and presentation biases. *Journal of Marketing Research, 46*(6), 764-776.

O'Donoghue, Ted, and Matthew Rabin (1999), Doing it now or later. *American Economic Review, 89(1)*, 103-124.

Orne, Martin T. (1962), On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist, 17*, 776-783.

Parfit, Derek (1971), Personal Identity. *Philosophical Review, 80*(1), 3-27.

——— (1987), *Reasons and persons* (Repr. with corrections. ed.). Oxford: Clarendon Press.

Preacher, Kristopher J., & Andrew F. Hayes (2008). Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods, 40*, 879-891.

Pronin, Emily, Christopher Y. Olivola, and Kathleen A. Kennedy (2008), Doing unto future selves as you would do unto others: psychological distance and decision making. *Personality and Social Psychology Bulletin, 34*(2), 224-236.

——— and Lee Ross (2006), Temporal differences in trait self-ascription: When the self is seen as an other. *Journal of Personality and Social Psychology, 90*(2), 197-209.

Proust, Marcel (1927/1981), *A la recherche du temps perdu (A remembrance of things past)* (A. Mayor, Trans. Vol. 1). New York, NY: Random House, Inc.

Roberts, Brent W., and Wendy F. DelVecchio (2000), The rank-order consistency of personality from childhood to old age: A quantitative review of longitudinal studies. *Psychological Bulletin, 126(1)*, 3-125.

Samuelson, Paul A. (1937), A note on the measurement of utility. *The Review of Economic Studies, 4*(2), 155-161.

Samuelson, William and Richard Zeckhauser (1988), Status quo bias in decision-making. *Journal of Risk and Uncertainty, 1*(1), 7-59.

Schelling, Thomas C. (1983), *The mind as consuming organ.* Boston, MA: Harvard Institute of Economic Research.

——— (1984), Self-Command in practice, in policy, and in a theory of rational choice. *The American Economic Review, 74*(2), 1-11.

Segovia, Kathryn Y., and Jeremy N. Bailenson (2009), Virtually true: Children's acquisition of false memories in virtual reality. *Media Psychology, 12(4)*, 371-393.

Singular Inversions (2004), Facegen Modeller 3.1. Vancouver, B.C.

Standing, Lionel (1973), Learning 10,000 pictures. *The Quarterly Journal of Experimental Psychology, 25*(2), 207-222.

Strotz, Robert H. (1956), Myopia and Inconsistency in Dynamic Utility Maximization. *The Review of Economic Studies, 23*(3), 165-180.

Thaler, Richard H., and Shlomo Benartzi (2004), Save More Tomorrow (TM): Using behavioral economics to increase employee saving. *Journal of Political Economy, 112*(1), S164-S187.

——— and Hersh M. Shefrin (1981), An economic theory of self-control. *Journal of Political Economy, 89*(2), 392-406.

Wakslak, Cheryl J., Shiri Nussbaum, Nira Liberman, and Yaacov Trope (2008), Representations of the self in the near and distant future. *Journal of Personality and Social Psychology, 95*(4), 757-773.

Weber, Elke U., Eric J. Johnson, Kerry F. Milch, Hannah Chang, Jeffrey C. Brodscholl, and Daniel G. Goldstein (2007), Asymmetric discounting in intertemporal choice: A Query Theory account. *Psychological Science, 18*(6), 516-523.

Wilson, Timothy D., and Daniel T. Gilbert (2005), Affective forecasting: Knowing what to want. *Current Directions in Psychological Science, 14*(3), 131-134.

Yee, Nick, and Jeremy N. Bailenson (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, OH.

——— and ——— (2007), The Proteus effect: Self transformations in virtual reality. *Human Communication Research, 33*, 271-290.

Yoon, Carolyn, Catherine A. Cole, and Michelle P. Lee (2009), Consumer decision making and aging: Current knowledge and future directions. *Journal of Consumer Psychology, 19*(1), 2-16.

Zauberman, Gal, B. Kyu Kim, Selin A. Malkoc, and James R. Bettman (2009), Discounting time and time discounting: Subjective time perception and intertemporal preferences. *Journal of Marketing Research, 46*(4), 543-556.

FOOTNOTES

1. Two participants from the future self condition were outliers in that they made a

   retirement allocation that was more than 1.5 times the interquartile range of the scale, and

   were thus excluded from further analysis.

TABLE 1
WEIGHTING OF INPUT PHOTOS TO CREATE MORPHS

Expression level	Input Photo 1	Weight 1	Input Photo 2	Weight 2
1 (Saddest)	Sad	99%	Neutral	1%
2	Sad	80%	Neutral	20%
3	Sad	60%	Neutral	40%
4	Sad	40%	Neutral	60%
5	Sad	20%	Neutral	80%
6	Happy	1%	Neutral	99%
7	Happy	20%	Neutral	80%
8	Happy	40%	Neutral	60%
9	Happy	60%	Neutral	40%
10	Happy	80%	Neutral	20%
11 (Happiest)	Happy	99%	Neutral	1%

Morphs are composite photos made from two input photos. Expression levels 1, 6, and 11 use 1% weighting to ensure that all photos will indeed be morphs and not original images.

A)

B)

C)

Figure 1. Example of morphing procedure. A) Actual photo of first author. B) Non-aged digital avatar. C) Aged digital avatar.
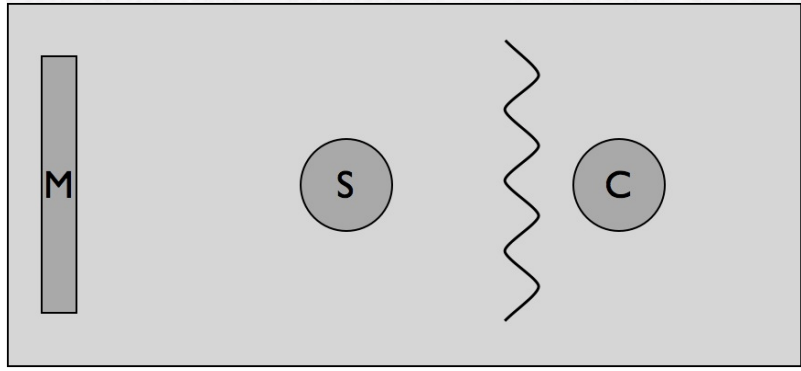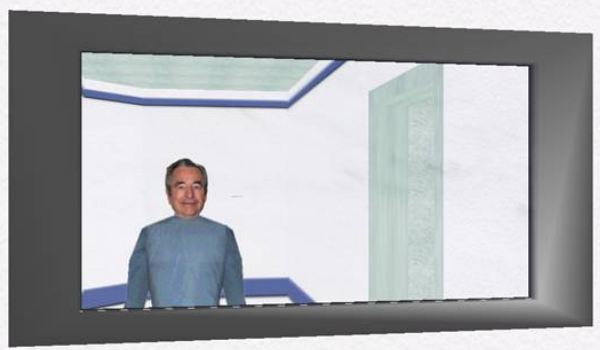
Figure 2. The top panel demonstrates the virtual reality technology: A) head-mounted display and orientation tracking device, B) behavior tracking cameras, and C) image generator. The middle panel demonstrates the image seen by the user depicted in the top panel. As he moves and gestures in the physical world, his aged virtual self image in the mirror moves with him. In the bottom panel is a diagram showing the layout of the room, with the position of the virtual mirror (M), the position of the Subject (S), the curtain, and the position of the Confederate (C).

Please use the scale below to indicate your preferred retirement allocation

Current Annual
Income
XX%

Retirement
Income
XX%

Next

Please use the scale below to indicate your preferred retirement allocation

Current Annual
Income
XX%

Retirement
Income
XX%

Next

Figure 3. The "current self" (top) and "future self" (bottom) conditions of Study 3A.

Please use the scale below to indicate your preferred retirement allocation
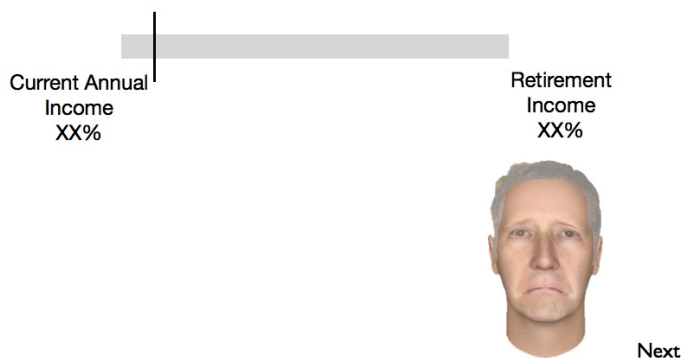
Current Annual Income XX%   Retirement Income XX%

Next

Please use the scale below to indicate your preferred retirement allocation

Current Annual Income XX%   Retirement Income XX%

Next

Please use the scale below to indicate your preferred retirement allocation

Current Annual Income XX%   Retirement Income XX%

Next

Figure 4. The position of the slider affects facial expressions in the photos displayed in Study 3A. This figure depicts the "future self" condition. As the slider moves to the left (top third), the future face becomes sadder, consistent with having less money in retirement. When the slider is in the middle (middle third), the future face has a neutral expression. As the slider moves to the right (bottom third), the future face becomes happier.

APPENDIX

The User Interface

The user is asked to give his or her sex. All other parameters are set by the software. The setting is one in which the subject has one decision variable – the percentage of gross salary to be contributed to a 401(k) plan each year. This is chosen by moving a slider. The leftmost position involves no such contribution. The effect is shown by a number or picture representing the average percentage of net income to gross income during the working years. The net income is gross income less employee contributions for social security and medicare and contributions to the 401(k) plan.

The output of the financial model is the ratio of net retirement income in the first year of retirement to the gross income in the year prior to retirement. This is shown by a number or picture on the right. As the slider is moved right the ratio for the working years will fall and the ratio for the retirement years will rise.

Pre-retirement Assumptions

The user is assumed to be 22 years old, to start work at 23, to have saved $5,000 already, and to retire at 68. All calculations are in real terms. The user is assumed to have a salary history that follows a path consistent with the cross-section of average earnings for year-round workers in the U.S. and an annual progression of average real wages of 1.1% per year (consistent with the intermediate projections of the U.S. Social Security Administration). The assumed path is calibrated to reach a real value of $64,000 per year when the worker is age 40.

The position of the slider determines the user's desired contribution to a 401(k) plan, expressed as a percentage of salary. His or her employer is assumed to provide a contribution

equal to 50% of the user's contribution, up to 6% of salary – the most common formula used by employers in the U.S. The employee's contribution is capped at a real amount equal to $16,500 in real terms, increased by 1.1% per year, on the assumption that the law is changed to keep up with real wages.

During the working years, salary is reduced by deductions for social security and medicare. The former is equal to 6.2% of employee pay up to a cap that starts at $106,800 and is assumed to increase in real terms by 1.1% per year. The latter is assumed to equal 2.9% of salary in every year.

Post-retirement Assumptions

When the worker retires, income is assumed to come from social security payments less the medicare deduction plus money received from an annuity purchased at retirement with the proceeds from the 401(k) fund.

The social security benefits are based on the average indexed social security earnings for the 35 highest years of contributions. The formula includes two breakpoints that are assumed to increase in real terms during the user's working life by the assumed real wage growth rate of 1.1%. From this amount a constant real value of $1,152 is deducted each year to cover Medicare costs after retirement.

Contributions to the 401(k) account during the working years are assumed to earn an annual riskless real rate of interest of 2.5%. This falls between the real rate assumed in Social Security projections (2.9% in 2009) and typical rates on long-term Treasury Inflation Protected Securities (slightly under 2% in mid-2010). Using these assumptions, the real value of the account when the user reaches retirement age is calculated. The proceeds are assumed to be used to purchase a single-person real annuity, priced using the Society of Actuaries' Table RP-2000

mortality assumptions, a real rate of interest of 2.5% and a 15% surcharge for costs and compensation for mortality table risk.

Finally, real retirement income for the first post-retirement year is computed by adding to the 401(k) annuity income the social security benefit and subtracting the Medicare deduction. This is divided by the real salary for the last working year to provide the value for the right-hand side of the slider.