Federal Reserve Forum on Consumer Research & Testing: Tools for Evidence-based Policymaking in Financial Services, November 9, 2010
Panel Two: Exploring research methodologies for consumer testing and studying consumer behavior

Hal Ersner-Hershfield:
Okay, great. So thanks again Jeanne and Ellen for organizing today. So in the spirit of talking about methods, I'll be presenting some new methods to be used in consumer research and in doing so, highlight one research stream that I've been engaged in the last four or five years, which is motivated by the problem. It starts sort of all--a lot of the research stream in here today which is that, here's the start of the problem. Life expectancy in America has been rising precipitously. So in the last hundred years--that's not the problem quite yet. We'll get to the problem a second. In the last century, more years have been added on to life expectancy than of all the previous millennia combined. So that's good. The problem is that saving rates have not kept up with this increased amount of time that's been tacked onto the end of life. So if we look at the personal saving rate, it's been going down over time. The National Retirement Risk Index from Alicia Munnell's group has been increasing over time and significantly so after the fiscal crash of 2008. And then, maybe this is most compelling, the McKinsey Global Institute estimates of fully two-thirds of baby-boomer households won't be able to maintain their preretirement standard of living in retirement.

So we've started talking today about a lot of the reasons for why this sense of under saving occurs, and Dan and John talked a lot about of them. I'm gonna zoom in on one theoretical foundation for why some economists and philosophers and psychologist have started to think about the reason why this under saving occurs. And this is in large part based on philosophy by a philosopher named Derek Parfit. And so the basic idea, and Dan touched on this earlier, is that we think about our future selves in different ways. So we feel potentially more connected to our future self in five years than we would to our future self in, say, 10 years or 15 years or 20 years' time. And at the very extreme--and this is the important point -- at the very extreme, our future selves might feel like or might seem like different people all together to us. They might feel like strangers. Okay, so this is a philosophical notion. It might seem a little bit arcane. This idea is actually slipped into pop-culture a lot recently. So here is the one of my favorite examples of this. This is Seinfeld on this idea. Let's see if that works.

[Background Music]

“I never get enough sleep. I stay up late at night, 'cause I'm Night Guy. [Laughter] Night Guy wants to stay up late. What about getting up after five hours sleep? Oh, that's Morning Guy's problem. That's not my problem, I'm Night Guy. I stay up as late as I want. So you get up in the morning, you're you yawn, you're exhausted, groggy, ooh I hate that Night Guy! See, Night Guy always screws Morning Guy. There's nothing Morning Guy can do. The only Morning

Guy can do is try and oversleep often enough so that Day Guy loses his job and Night Guy has no money to go out anymore."

Okay, so Seinfeld brings up this basic idea. I have the quote here just in case my technology didn't work, but the basic idea here is that there is this real disconnect between selves over time. Seinfeld is talking about it on a night to day basis. We're talking about it in more of a long term sense here. But this disconnect can actually impact the feelings that we feel and really alter the decisions that we make.

So here's the basic hypothesis that we wanted to test. The way that we think about our future selves on an individual level can actually be an important predictor of saving behavior. So if I think of my future self as being very different from who I am now, if it's a stranger, my motivation to save will be low. If on the other hand my future self seems very similar to me, if I feel very connected to it, my motivation to save will be much higher.

So we first wanted to test this theory using some of the more standard methods that we've talked about just using paper and pencil methods. And Dan and his group have used some of these methods as well. And what we did is we borrowed from some previous psychological research. We know one of the best ways to get people to answer how connected or similar they feel to other people in their lives is to ask them to just pick a group of overlapping circles. So we apply this to the self and we said, "Pick one of these groups here and, you know, think about it in terms of how this describes how similar you think about yourself in 10 years. So down here we have a case in which there's almost complete overlap between the self now and the self in 10 years and up here is the case in which there's no overlap. The self feels very different from who I am now.

So we gave this to a sample of adults across the Bay area and across the lifespan— so ages 18 to 86 – and we wanted to map their scores on that scale to how much money they've saved over time. And to examine that, we assess their assets that they'd accrued over their lifetimes, and we just gave them a categorical scale here. Zero meant zero dollars all the way up to 16, which was one and a half million dollars. Assets here were things like saving in savings accounts, the portion of their homes that they owned, et cetera. And what we're looking for is a relationship. So here on the X-axis is similarity to the future self in 10 years' time and on the Y-axis is the amount of assets that people had accrued. What we get is this significant positive relationship. So the more similar people felt to their selves in 10 years' time, the more assets they'd accrued over time. And this holds even when we control for other predictors of saving and assets, such as age and education.

And we have a couple of other studies, and as I said there're other lab groups that have other studies, that converge on the same finding here. But one thing you'll notice is that this used a measure of self report. And so we're starting to talk a little bit about some of the problems inherent in self report. We have social desirability which might not play as much into this, but it is a subjective measure. And so we wanted to try to draw on other measures to figure out if we

could measure this general idea here in a more objective, nonobtrusive way. And here we drew on biological measures.

So there's been sort of an emerging field known as social neuroscience and neuroeconomics which looks at how people make decision in the brain. And so here's a previous finding that really informed a lot of our thinking. There's been a spate of research over the last 10 years that shown that when we think about ourselves, there's a lot more activation in one portion of the brain compared to when we think about another person. So the brain can tell the difference between me and not me. I'm gonna go with Adair here just 'cause I think this is what we're doing today. So the brain can distinguish between the self and another person. So we wanted to take the same finding and reason that if the future self is thought of as another person, we should see those same differences in the brain – the same me, not me differences – and it should be in that same region, that region that codes for the self versus other. But most importantly, on an individual difference level, this is getting to what John was talking about before, these differences between current self and future self should be predictive of saving behavior.

How do we measure this? We had our participants come into the lab, we put them into a scanner, a brain scanner that is, and we had them make judgments about a number of trait words. And we had them judge whether these trait words apply to themselves now, themselves in 10 years' time, another person now and another person in 10 years' time. And we had them come back to the lab two weeks later after going through this whole task and we gave them a decision making exercise. And this is known as a temporal discounting task where we just ask people to make a number of choices between smaller monetary rewards that they can obtain immediately versus larger monetary rewards that they have to wait for. So the first thing that we saw was that we replicated all that previous research. We see greater activation or blood flow to this one portion of the brain when people thought about themselves compared to when they thought about another person. So we know that our equipment is working. This is good.

What we want to know is whether or not in this same region there was a difference between the current self and the future self, and that's exactly what we found. So you can see within there there's an area of the brain that codes for current self versus future self and it's helpful to look at this time course graph here. We don't have to go into the details here but the important thing to note is that there's a real difference between activation that's elicited by the current self -- that's the blue line there -- and activation that's elicited by thinking about the future self. And most importantly, that future self activation is on par with activation that's elicited by another person. So on an average level in the brain, people think about the future self as if it's another person.

Okay. Again, this is across all of our subjects. So we wanted to look on an individual level whether or not differences in this difference here were predictive of saving on that discounting task. So we created a different score. We just subtracted the future self activation from the current self activation and plotted along this X-axis here. So the further out we get to the right, the bigger the difference there is between the current self and the future self -- in other words, the

more the future self looks like another person to an individual on a neuro level. On the Y-axis is our measure of discounting behavior, and what we see is this positive correlation. So on this neuro level, the more the future self looks like another person, the worse these people are at saving. The more they'd opt for that immediate reward that they could get right now rather than waiting for that larger reward that they could get at a later point in time. So we have a couple of studies here that sort of triangulated on this same general finding that the way we conceive of the future self is a predictor of the types of decisions that we make.

So we wanted to move beyond this and think about, "Well, how can we actually design interventions based on this knowledge? How can we actually get in there and bridge that gap between the current self and the future self?" So here is--oops, sorry, well, there was an ad that flashed up there very quickly. Lincoln Financial -- you may have seen this ad before and they have these ad campaigns where they say, "If your future self could talk to you, what would it say?" There is one where there's a guy on a plane and his older self sits next to him and says, "I'm going to sit in first class because I've been saving now. I can afford that." Never mind whether that's actually a good decision to buy a first class ticket.

So this idea of presenting people with an image of their future selves is actually, you know, potentially a good one and based on some theory. So, Derek Parfit and economist named George Loewenstein have put up this idea that one of the reasons that we have these failures of intertemporal choices is because we can't vividly imagine what we might be like, what we might want, what we might desire in the future. And Loewenstein has pointed to how vivid examples can really affect behavior. So there's very few pulmonologists who smoke because every day they are exposed to the blackened lungs, the effect of smoking on their cancer victims. So what we wanted to know is could we show vivid examples of people's future self to them and could that then cause them to save more? Could it show them what they'll actually be like, what they'll actually look like, how they'll exist in the future?

So to do this we use the combination of what's known as age progression software – it's a software package called FaceGen – as well as virtual reality. And I'll explain this in a second. But I'll first give you an example of the age progression software. I've to use my face 'cause everyone seems to get a little offended if I just randomly pick them out of the crowd. So we first take a neutral image of our subject. This is my neutral image which I've now decided doesn't look so neutral. [Laughter] So that's me without any emotion. We feed this into the computer morphing software, the age progression software, and it creates what's known as avatar or a digital version of a person. And so here is the digital version of me. And then we run this avatar though the age progression algorithm. This is a standard algorithm in the software package that makes me look older. So this is me when I'm about 68 years old. I've shown this to my fiancé, and she is still going through with the marriage plan. So that's good. What the software does is essentially what a lot of--is what happens with aging, sags the cheeks a bit, drops the jowls, grays the hair, we add some sunspots in.

We then have our participants go into a virtual reality environment and interact with either their current selves or their future selves. And so this is how virtual reality works. The person is in this lab. They've got this mask on. They've got little lights on their shoulders. They walk around. There's cameras that tract their movements and display it back to the computer. The computer then displays the movements back to the goggles. If I move forward a couple of steps the world that I see moves back a couple of steps and this is what it looks like if you're an actual participant in the study. So the guy is in here. When he walks around, he sees this room and in the--and he sees a virtual mirror. And in that mirror he'll see his avatar. So this is what it actually looks like. So he walks around. It looks very real. And at a certain point, he gets to this mirror and what he sees in that mirror is the avatar – either himself now or himself in the future – and that avatar moves with him. So if he turns his head the room tilts. If he sort of turns to this side, the room turns to this side. It's just like looking into an actual mirror. So--and he can disappear if he wants. That's another example.

So in this study our participants either see themselves in the future or themselves now, and they spend about five minutes interacting with that self. They answer a series of questions to get them to identify with that avatar that they see in the mirror. They come out. They do a number of questionnaires, one of which is a monetary allocation task where we say if you're given a certain amount of money, how would you allocate that money, what are the bins that you would put that in? What we see is that those people who are exposed to their future self allocate about twice as much money to a hypothetical retirement account. And we've run some other studies that I won't get into now controlling for a lot of alternative explanations.

But what we wanted to do next, what I wanna show you now, is drawing on the idea that this is an exciting finding. However, it's expensive to run. It's not necessarily that practical. It's not every investment company or large employer that can take a virtual reality lab and just plop it into their HR office and get people to make these decisions based on this. So we've been working now with some computer scientists. We've created an online version of this tool. And what we do here is we--And our first version of this, we had undergrads, again, and we asked them to make a retirement allocation decision. This is just like the way it would be on your second week of the job when you have to make your decision in terms of where will my money go to my 401(k) plan. And they get a slider bar and they can drag this bar to change their allocations and they see some numbers. So on the one hand they see the percent of their paycheck that they'll be able to take home annually in their working life. On the other hand they see the percent of their final years working income that they'll be able to take home in their retirement period. And we've got a financial model here. We've actually worked with a real economist -- not psychologists doing this sort of work, who's created the model for us and we have a number of different parameters in addition to these ones. We set these parameters to be fixed for the lab experiments so that we can sort of put everyone on a level playing field.

But what happens in this study is we put our participants into this sort of computer design here and we have them make decisions but they see an image of either their current self or their future

self. And importantly, that avatar changes emotional expression as a function of the decisions that they make. So if I decide for not a lot of money, I'll be sad. If I opt for a lot of money for myself now, I'll be happy. And then we flip this for the condition in which people see their future selves. If I take too much money for myself now, my future self will be sad. If I opt for a lot of money, my future self will be happy. And we're running a lot of different versions of this where we're putting in a lot of different controls here. But the idea being this acts as sort of a nudge towards allocation behavior. So here's how it looks for a subject. They drag the slider and that face actually changes in real time as a function of the allocations that they make. So it smiles over there. And so, again, we wanted to see whether or not exposure to that future self could actually cause--give rise to a bump in allocation behavior. And in this first round -- this is, again, a hypothetical allocation -- we see a significant increase in allocation behavior. It's on the order of about 1.75 percent. It's not huge in magnitude, but that actually does have some serious effects over time once we think about obviously compound interest and what not.

So what comes next? Right now there's some manual work that's required for these aged images. We're now working with some other computer scientist to actually automate this age progression procedures so people could take an image of themselves with their iPhone or their webcam and then automatically age their faces. This isn't that far off from the future. There's other apps now that can change your face. There's a terribly named app on the iPhone called FatBooth which takes an input picture and makes you look heavier. So, my friend – we'll see if he'll still say he's my friend – Dan. This is last night before Ethiopian food and then this is an image of him that's been fat morphed, we'll call it. So there's--he still looks good, he still looks good. So we wanna try to automate these procedures, the idea being that financial advisors can use this image generation process with clients when making allocation decisions. Employers can implement this type of technology into retirement allocation sessions. And potentially in the future, consumers could customize this type of software for use with their online banking tool. So I might make a decision in advance and say if I spend more than 400 dollars a month on, you know, category X, then when I log on to my Bank of America account my avatar is not gonna be happy with me or something like that. There's a lot of different tweaking we can do here with this software.

And then of course the bigger questions that arise from this are whether such interventions will work for real people – rather than just, you know, smaller samples of college undergrads – on real decisions. What are the boundary conditions? So for what decision does this work and for what decisions doesn't this work? And then sort of in a related vein, for whom do such interventions work? So what are the individual differences that might make it so that people are more susceptible to these types of interventions than other people?

So thanks so much for your time.