

Neuroeconomics: The Secret Life of Homo Economicus
March 17, 2007
The Philoctetes Center

Nersessian: I can't fill David Kirkpatrick's shoes because they're both literally and symbolically much too big for me, but I'll try my best and see what I can do. So I'll start with introducing the panelists. Elizabeth Phelps is Professor of Psychology and Neuroscience at New York University. She is currently the President of the Society for Neuroeconomics and serves on the Boards of the Association for Psychological Science and the Society of Neuroethics. Most of her work has been focused on learning, memory, and emotion. Paul Glimcher is Associate Professor of Neural Science and Psychology at New York University's Center for Neural Science. He's an Investigator for the National Eye Institute and Founding President of the Society for Neuroeconomics. Bill Brown has long been in the economic forecasting business, formerly as Chief Economist at JP Morgan and currently as an Economic Advisor to Barclays Capital and as an independent macroeconomic trader. Alberto Bisin is Associate Professor of Economics at New York University. He's an Associate Editor of *Research in Economics* and *Economic Theory* and has written numerous articles for economics journals.

I think the best way to start today is by asking Elizabeth Phelps, who is the current President of the Society of Neuroeconomics, to tell us what neuroeconomics is.

Phelps: I'm going to defer to the Founding President because I think Paul has been instrumental in pushing this field forward and I'm sort of riding on his coattails a little bit.

Glimcher: I'm going to defer to Danny, who's not here. No, I'll tell you my personal answer and it's not everyone's answer. I was trained as a neurophysiologist. I learned how to study the activity of single nerve cells in humans and animals making decisions. I learned how to use brain scanners on people and animals. What I was interested in was decision-making—how you pick a wife, how you pick an ice cream flavor. That seemed to be a reasonable thing for a neurobiologist to study, but it became clear that, first off, neurobiologists were ill-equipped theoretically to think about that problem; they really were engineers who built wiring diagrams and did not think deeply about mental life, psychological processes, or prediction of optimality theory. And this led me to psychology, to think about mental life and mental experience, and I realized over the course of many years—and eventually became a professor in that department, as well—that you had to combine the insights of psychology—experimental insights of how people actually behaved—with neuroscientific methods if you were going to develop any comprehensive description of human behavior. Now that's not too surprising.

As I learned more and became driven to understand everyone who had thought about decision-making, it became clear to me that at NYU, there's actually a department of decision-making, and that department was called the Economics Department. That drove me over many years to try and synthesize into my own view the economic approach, which is really, I think, finally about specifying what people *ought* to do. Psychology is sort of empirically what they *do* do and neuroscience is how the hardware achieves what they *do* do, but can it explain the gaps between what they *should* do and what they *do* do? So for me, neuroeconomics is the synthesis of at least these three disciplines. Its goal really is to adjudicate the gap between economic theory, which

tells us so clearly now what an efficient chooser, an efficient person, would do, and psychological observation, which tells us that people don't do that. There's a hole there, because economics is this spectacular theoretical device that works very well in large groups and works, with apologies, quite poorly with regard to large chunks of individual behavior. Psychology is a good descriptive tool, but it really, at this moment, lacks any coherent, sweeping theoretical view.

So you've got this problem: psychologists can tell us what's wrong with economic theory; economists can tell us what people ought to do; but you don't have a theory of what people do. And it's my hope that neuroscience will fill that gap and that these three disciplines together will create an algorithmic or mechanistic model of human behavior, which will be able to unite all three levels. I think that the grand, unifying excitement of that and the reason why there's so much excitement around neuroeconomics today—at least in my mind—is because that would be a synthesis very different from the one that occurred between molecular biology and biology in the '50s, or between physics and chemistry at the turn of the 1900s. It would actually be a fusion across the border of the social and natural sciences. It would suggest that the kinds of grand unities we see throughout the natural sciences will progress ultimately to the social sciences, as well.

Nersessian: Do you want to add something?

Phelps: No, I mean if we're going to go through how we got to this, I'll give my little story and then talk about how I became involved in neuroeconomics. I was interested in understanding human behavior broadly, so I trained initially in psychology and cognition, but then realized that looking at the brain was going to be a very powerful way to help us understand cognition and I re-trained in neuroscience. So for the past several years, I've been specializing in how emotion changes how we learn and remember and, of course, one of the consequences of learning that something is emotional is that we're going to make different choices. And so I got to this field because I became interested in how what I'm studying about emotion and learning relates to the choices we make. We assume, of course, that when we find out something is fearful or rewarding or something like that, we're going to make a different choice.

When I got into looking at choices I found, much like Paul, the psychology of decision-making has merged quite a bit with economics. It got me into thinking about economics and collaborating with economists. So I came at this more from just the psychological perspective of understanding how it is when we learn things, emotion changes them, and how that changes the choices that we make. In doing this, I'm just starting to really learn about economics and collaborating with economists and understanding this perspective. The theoretical and computational aspects of economics is now adding something to my psychological neuroscience work that wasn't there previously, so I see the merging of the two enhancing both fields. The Society of Neuroeconomics, which Paul was instrumental in starting—I went to the second meeting. The first one was invitation-only and I didn't get invited. At the second meeting it was decided to form a society, which came about, and there's been about four meetings so far. One of the things they've been trying to do is alternate the Presidents to be both a neuroscientist or a psychologist, and an economist. So Paul was the first President, the second was Colin Camerer,

who is an economist, and I'm the third, bringing in a little more of the psychological aspects of the field.

Glimcher: Can I add it was not I who did not invite you.

Phelps: It had nothing to do with Paul. I think one of the interesting things for me, as somebody who has gone between fields within the study of behavior, neuroscience, and psychology, and within psychology, different subdisciplines—emotion spanning social psychology, clinical psychology, and cognitive psychology—is how the culture of these two fields merged. So that psychology and neuroscience has become a rather broad tent because we study all of human behavior. The field of economics, now that I'm learning about it, doesn't seem to be such a broad tent. I think it's going to be more of a challenge for the neuroscience and psychological research. I mean, psychology's been doing this for quite a while, but for the neuroscience research to influence how the economists are thinking, I think there's already been quite an influence from economics in how neuroscientists are thinking these days.

Nersessian: Now we have two non-neuroeconomics people, so what's your reaction to what is going on in terms of neuroeconomics and all the excitement about it?

Bisin: I come to this from economics and economics is the way I was taught and what most economists still nowadays are taught is rational choice. So that's what we do. This is the basic assumption and then there is a lot of analysis of what this implies for all sorts of behavior. Economists don't think of themselves as just thinking about economics. I myself have done a lot of work, which many would think of as sociology, maybe. So that's the basic idea. Now this has been criticized by sociologists and anthropologists—social sciences in general—for a long time. By this I mean rational choice. The reaction of economists usually being “well, whatever, we don't care,” and that's because the criticism was, to our understanding, in part because economists are arrogant but also because that criticism, from my point of view, was nested in not really understanding what an economic model is.

Then the psychologists started criticizing and that sounded much more interesting to us. Many of us, including myself, thought there were a lot of very interesting ideas there and, you know, from the 1970s on, there's been an enormous amount of research, usually experiments in the lab, which have shown that in many very consistent, coherent ways, people were not playing rational choice, or not acting as rational actors, the way we would think they did. That still doesn't mean the rational-choice model is wrong, but it makes you think that there's something interesting to do. So that's my reaction. I think there's a lot of interesting stuff in here.

I would say I completely disagree with Paul when he said economics is about what people ought to do—we think of it as a positive theory, not as a normative theory, including rational choice. So the whole point here is that economists are interested in behavior and think about rational-choice behavior not as interesting behavior per se, but in terms of the implication behavior has on economic or social choices and equilibrium. So the fact that you see people not really quite acting as rational actors in the lab, if our rational-choice models or financial markets work well, we think that's good. It means that the deviation from rational choice in some sense gets

corrected by the large numbers, by the market, or whatever you want to think about, and they're good enough.

So my reaction to this is I think rational choice has been incredibly successful in many fields and I think the evidence in the lab and the evidence that is coming out much more recently from the other lab—I mean in the brain—is incredibly interesting in giving us possibly some new general view about how people act. I actually completely agree with Paul when he says in psychology we just don't have a basic model that we can apply. That's why I think rational choice has been so successful because first, it's incredibly simple and second, incredibly deep, much deeper than just what the words mean, and very unifying. And yes, in some sense, wrong, and we are looking for something else of the same order methodologically, with something else that looks like rational choice, in the sense of being a very simple, or a couple or a few, very basic principles that we can push from. So in some sense, I look at this literature and I'm very excited because it looks like we're going there. We're not there, obviously, but we're going there.

I must say also, personally, when I got into this I started reading psychology and I was very disappointed. I think we—psychologists and economists—have a very different way of thinking in terms of what a model is, what a theory is. When I started reading neuroscience—I mean, not real neuroscience because I don't understand, but the part of neuroscience Paul spoke of, the part of neuroscience that is closer to economics—to me that seemed very exciting because we were speaking a relatively similar language. We agree on what a model is, we agree on what explaining means, we agree on many things, on many levels, and that seems very fruitful.

Nersessian: Can I just ask him to say two words on what rational choice is?

Bisin: Yes. Rational choice: the way economists think about this is they set a set of rules—we call them axioms—that an agent, which we call rational, should satisfy when he chooses. So these are things of this sort: if on this big table suppose these are glasses of different liquids and this is a Coke, a beer, and something else. If on this table I take the Coke—there's a beer here and orange juice over there—and so if I take the Coke, fine, it must mean somehow that I like Coke. Now, let's go to a different table: there still is Coke, there still is orange juice over there, but we took away beer. If now, because beer is not there, I choose orange juice, then this is not rational. This is a very simple kind of argument. From this and, of course, other stuff of this sort, a lot of logical reasoning has come up. It tells you that an agent that satisfies in his choice this kind of behavior will end up maximizing some utility function. And that's where most economists start from. So this background is called decision theory. People have looked at, in some sense, what does rationality really mean. And in a nutshell, it really means this, what I just said about the glasses. What it implies is the maximization of some utility function, so most economics don't do the glasses; most economics is about utility function and things that an agent maximizes, chooses an element of some set, which makes this utility function of the agent as high as possible. That's what rational choice is.

Nersessian: Thank you. Can I get Bill to say something?

Brown: Well, I come from a little bit of a different background in the sense that I do no research or investigation. I'm a practical economist, having been in that business a long time. But in

recent years, when I've read some of this literature, it seems very natural to me because one of the things I did at JP Morgan was participate in their risk-management effort, which was basically how to control traders to make them make fewer mistakes and do a better job. In that process, it was very clear that humans are not good at doing certain things, and it's the exact same things I read in the papers on neuroeconomics. They don't deal well with probability, distributions, particularly the tails and small probabilities. They're very associative; they associate one thing with another for reasons that aren't rational, you can say. They're very scenario-oriented; they like to think in terms of a pattern that plays out that is like the real world. They can't separate themselves from that. And all of these things are not very helpful in terms of pricing a derivative security, say, or positioning in the market. You need to be rational if you want to make money in that business.

So back in the '70s we used to try to use ways to train people to be better. For example, we'd ask them to give us probability distributions—this is our best traders—and we would collect them and talk about them and make people think about them and show them the historical ones. And then we'd come back six months later and say, "Look, you've given me this probability distribution on where the ten-year bond yield's going to be, and fifty-five percent of the time, the result is in a box over here that you put zero probability of happening." There's something wrong here. You know, you're just not quite right. We thought this would make them do a better job. I would say all of this effort in the '70s was hopeless. You'd tell them that and six months later they'd come back and do the same thing. And these were the best traders—these guys weren't awful, they were actually doing a pretty good job out there.

It really goes to show you what the research says: people are hard-wired to work in certain ways and they do certain things well and they do not do other things well and it's very difficult to get them to do it. Now since then, the financial world has moved in ways of just taking these decisions out of individuals' intuitive grasp. You've got big computer models—algorithms—that are driving these things. And an investment bank would never let some individual by the seat of his pants say, "I think that's what that derivative security ought to be worth." It's just not how it's done. You haven't made the people better; you've taken the control away from the people, or you've structured it in a way that you've protected yourself from their shortcomings, at least to a large degree. That's how you got these things working better—not improving the people, but by putting up a separate system or structure. So I sort of come from a world that uses some of these things and experiences them.

Glimcher: Can I correct for a moment, Alberto? Let me put it this way: I was raised as a biologist, and the theory of natural selection is, if anything, the centerpiece of all biological thought. The theory of natural selection says that if there are two animals trying to live in the same niche, the one who maximizes his reproductive fitness better wins. That is actually economics, where the utility function is replaced by this defined notion of fitness and success of generation—the number of your genes that repeat. That is to say, if there were two animals, one of which was a better utility-function maximizer and more rational with regard to his fitness than the other, the animal that is the better maximizer wins. So we would expect to find after 600 million years of evolution on this planet that animals ought to be pretty close to good at everything they have to do to survive, and we believe that the rational-choice model would be a good predictor of their behavior. Now, in fact, that's true: every time this has been taken to the

field with animals, it just works perfectly and the animals' behavior is bang-on. We're talking about individual animals. There are beautiful studies of birds trying to decide whether or not to grab a mealworm. Based on the size of the mealworm, how long it takes to uncover the mealworm from its shell before they eat it, you can adjust the weight of the mealworm, and the birds know exactly where the cut off is. If a mealworm gets one iota too small, the bird knows to turn it down because the bird is perfectly optimizing its use of time and its gathering of energy.

To me, this stands in stark contrast to the way we understand the decisions of individual humans. I make a very clear distinction here between microeconomics, the study of individual human decision-makers, and macroeconomics, which is kind of what Alberto is leaning on as being so good. I agree with that. We don't understand the decisions of individual human decision-makers using economic models anything as well as we understand these birds. You ask people to make simple decisions like losses and gains and they are petrified about losing money. If I say to someone, "Here's a lottery ticket. It carries a fifty percent chance of losing a hundred bucks. How much do I have to pay you to buy that lottery ticket from you?" or something like that, and they overvalue that lottery ticket. It's like, they're so petrified of that loss, and they'll pay anything to get rid of it, which makes no sense, if what they're trying to do is maximize what they have in their pockets.

So there is, I think Alberto has to admit, a huge gap in the explanation of individual human behavior at the economic level. Now this washes out at the market level. That is to say, the rational-choice models do great at the market level because you put a couple billion agents together and you put a couple of giant investment houses that are using trading algorithms and accounting for twenty percent of the market in any given instant, and now the markets look really good and rational-choice theory predicts really well. So what I see in economics is that really important gap between the quality of prediction in microeconomics and the quality of prediction in macroeconomics. The good news for economists is it doesn't matter. The bad news for us, as scholars trying to understand human behavior, is the first theory doesn't work nearly as well as the second.

Bisin: Should I answer? Not answer, react? I think I agree with most of what Paul said. I have a different view about it. So we agree that macroeconomics is perfect.

Nersessian: Only for the moment.

Bisin: Only for the moment. But there's one point that I want to stress here: it's not macroeconomics or finance or labor economics—the way that economists think of it does not come out of the large number of agents only. It relies crucially on rational choice. So it's not that any model with a million agents or many models with a million agents would work. All the predictions that economists think about in finance, that economists think about in macro, that economists think about in labor, are predictions that crucially depend on rational choice. I was, but most economists are not born this way. This came out of a big fight in the '50s and '60s in economics, in macroeconomics, between rational-choice people and non-rational-choice people. Actually, the negative reaction that most economists have against behavior economics comes from here because it somehow throws away—how do you say it in English?—throws away with the bathtub—

Glimcher: The bathwater.

Bisin: In the sense that a lot of the stuff that the economists understood in the '60s, which also goes under the heading of rational choice, is not really about rational choice. It's about what we call rational expectation—how we believe agents think the world behaves, so that part is crucial and has been thrown away with behavior economics, and that's something that the economists react very badly about. The first point I'm trying to make is it's not just with a million agents—rational-choice theory tends to work reasonably well, and reasonably well in economics as such. Now it's true that when you try to measure individual behavior, Paul is exactly right. When we look into the lab we see people not doing what we as rational-choice people expect them to do, it's true. There is no discussion. That's what I said at the beginning. That's where people like me and many other people got interested in psychology and neuroeconomics, because there's something really systematic—the fear of losses or the way people discount the future. There are many dimensions in which it's clear there's something going on. And the fact that we don't see that too much in macro and finance is important, but we will not know, in a sense, until we have the right models that deal with this issue at the micro level. We will have to take this model to the macro level and see if they improve. So in the '60s, when economists were thinking about models without rational expectation—there was the big fight, Chicago versus MIT, or whatever—at that point, macroeconomists there were very happy with their macro models. And then these people from Chicago came and said, "You shouldn't be happy; your models are wrong." So now we are happy and maybe we're wrong.

Nersessian: If you're happy, you're wrong, but that's okay.

Bisin: So it's useful to go and look into this. That's the point I wanted to make. I think in some sense there's also a difference in objectives here. Economists are interested in decision theory as a tool and decisions as a tool. Paul and Liz, neuroeconomists, I think, are more interested in decision theory, per se, which is fine. How people make decisions—that's very important. For economists this is the basic assumption: we're very interested, but there are some details that we don't care about. I mean, we call them details—they might be very important issues. The point, I think, is to try to understand if we can develop a model that can substitute for rational choice and which encompasses and takes care of all these things that we see in the lab, which seem to point to people having all sorts of other dimensions that are not rational.

Nersessian: So the question is, if I understand it, why do we need to understand neuroscientifically why people make the wrong choices? Why not recognize they make the wrong choice and account for it?

Phelps: I think this gets to an issue of why do you need neuroscience to understand any behavior, right?

Nersessian: Well, no, because neuroscience is different from neuroeconomics.

Phelps: I understand that, but I'm trying to make the point of what can neuroscience bring to our understanding of behavior, and economic behavior in particular. I think we can have

psychological models about how things work, but when we look to see how they're instantiated in the brain, we may find subtleties that help us understand our psychological and economic model. This has been a debate not just in neuroeconomics, but in cognitive neuroscience and social neuroscience and any field where you are interested in some aspect of human behavior and trying to see what you gain from understanding brain function. So we know that understanding brain function gives us another model of how behaviors are organized and put together; it's where behavior comes from. By understanding things at this level, we're going to get a much more sophisticated and nuanced understanding of the behavior itself. I think that's true of economic behaviors and other behaviors.

Glimcher: Yes, I agree, completely. In the 1950s, the way we pharmacologically treated mental illness was we took drugs—we ground up plants, basically, separated them into sub-components, and fed them to people who were psychotic to see if they got better. One could say that works. You have enough people, you feed them enough plants, and sooner or later you'll find something out. We don't do that anymore. Now what we do is try to understand the neural basis of psychoses; we try to understand whether there's a transmitter defect where it lives; what reuptake device doesn't work properly; we try and develop pharmacological strategies that are direct, pointed interventions. I'm going to bet that we've made more progress at this iterated level, where you do some psychology, you do some neuroscience, you do some molecular biology, you go back and forth, you build links between the levels, that that's a more productive way to treat mental illness than just grinding up plants and feeding them to people who are psychotic. Although, I admit, both work. So that's my answer to Alberto.

There are lots of big puzzles in economics that are unresolved. As Alberto knows, one of my favorites is that people are more impulsive than they should be. They seem to value immediate gains too highly. This leads to all kinds of macro-scale problems. People under-save for retirement. I mean, think about this problem: people don't save enough for retirement. How do we solve that in this country?

Bisin: That's an implication of those models, as you know, and that's one possible interpretation of—

Glimcher: Let me say it a different way, then. Our government, fearing—based on these observations—that people don't save enough for retirement, pays people to save for retirement. That's how we solve this problem. Incentivized—in order to make sure that you don't starve to death when you are old—

Bisin: The government does it because otherwise they would have to feed people who are starving to death. And if people know the government would have to feed people who are starving to death, they won't save for retirement. Rational expectation, there you go.

Brown: Just on this point of how you can have people producing irrational economy—there's reasons for this. It's the big numbers. It's a feedback system, like in evolution, that if you don't do it right, you lose. And there's a clear feedback through money. There's arbitrage. There's all kinds of reasons why an economic system, particularly a finance system, could produce rational results with irrational people underlying it. I think economists shouldn't be criticized too much

for those rational assumptions. But there are a lot of other areas of behavior where these conditions don't hold and there isn't any arbitrage, and there aren't large numbers. Take designing a system for cars that would make them safer—warning systems, lane things, all kinds of things. It's very critical how the mind works and how humans behave. It's very useful to know these things when you're thinking about what system in a car would improve it. Or the big problem in medical care: the error rate. You know, flying an airplane, you'd never accept the error rate you're getting in hospitals now. Why is that and how do you reduce it? Do the people behave wrongly? Doctors underestimate low probability events; they don't wash their hands unless you absolutely force them to because it's a very low probability that'll hurt. But thousands of doctors not washing hands thousands of times—it gets big. So I think there are all kinds of areas outside of finance or sort of traditional economics where knowledge of how the human mind works can be very useful.

Glimcher: I think I hear you two guys—speaking of the washing of hands—like Pontius Pilot saying, “We only really care about the markets. Economics doesn't reach down to individual choice. That may be crazy, but that's not our business.” And I want to push you on that.

Bisin: Can I try and answer the question that you asked, Paul, which is why neuroeconomics is useful or why neuroscientists should do their job and we do ours? I don't want to come across as the guy who says they shouldn't. Let's take some of the examples that Paul had. People fear losses. People might be too impulsive. All right, so what's been the typical way of dealing with this for an economist, and I want to claim in part also psychologist. As I said, we have utility function; we have people who say they maximize utility function. Most of the things that Paul said, he implied that what's inside utility function, what we care about, is money or consumption. But economists are very good at expanding this stuff, so they will say, “Oh, you know, it's not just that we care about money. We also care negatively in the sense that we care about losses.” There you go, the puzzle is solved. Of course, it's not solved, it's—

Glimcher: That's totally unfair, Alberto. Because in order for that to be true, there would have to be an existing axiom set for reference to preferences.

Bisin: No, but I'm getting there. I'm saying economists have done that, psychologists have done that, and it's crazy. They've said, “Oh, you know, it's not that people only care about consumption and money, they also care about not losing it.” Or they said, “Oh, it's not really that they're impulsive or anything, there's something in their utility function that makes them like in a weird way the present more than the future.” Things of this sort. So, you know, these are patches. You can go along with patches. Some people believe—and, as I already said, I don't—that people under-consume. We can write down a model where we have that sort of a patch. People care too much about the present then in patches, because the model will predict that people under-consume or under-save in some sense. That's not an explanation. I think this is why neuroeconomics is important. All these little patches—you can do it once, you can do it twice, you can do it three times—at the end, you've—

Nersessian: You've now said that the reason people don't save enough is because in their temporal lobe, there is an area that works in such a way as to not make them realize that it's important to save. Let's say you discover that. Then what do you do?

Bisin: No, what I want is a unifying model of many or most of these little things, because I believe there's something underneath. There's hopefully a unifying model. I don't care about any one specifically. I know Paul dislikes this, that I don't have a particular attachment to this, but there are all these two brains models out there, yes? So this is a unifying mechanism. It can be the right one or the wrong one, but it's a unifying mechanism, which will try to explain both the fear of losses and the present bias. So what I'm looking for in neuroeconomics as an economist is really this: some underlying model of what goes wrong with my rationality model at the individual level. Once I have that, I can try to fit it to the macro model and see if I was wrong before.

Phelps: I just want to give you an example to sort of follow up why it would matter if we knew about this part of the temporal lobe. We should get back to the two-people-inside-your-head debate at some point. But we did a study where we had people learn through trial and error to trust somebody or not trust somebody. So, somebody's going to reward you or they're not going to reward you if you trust them in an economic exchange. Prior to them doing this, we gave them a little vignette: This person is basically a good guy—nothing to do with economic behavior—he helps his friends out and that kind of thing; he works for Teach for America. This other guy is kind of sleazy, and this guy is kind of average. What we see when we do this is when you're working with a good guy, you don't use the part of your brain that seems to be responding to trial and error feedback that you're getting through these economic transactions. You use another learning and memory system when the guy is good to help guide your choices, and your choices don't change as much with a good guy, even though they're all rewarding you to the same degree. So you could say, well, maybe we would have seen that behavior, and we did. But we know something about how those brain systems work and what kind of information they encode and what kind of information they use when they're driving behavior versus when another memory system is driving behavior. This now tells us something about the psychology that we wouldn't have known otherwise. We can take everything we know about those brain systems now and use that to understand why it is that when it's a good guy, it doesn't bother you so much that he cheats you. That's going to now inform our psychological model of how this could work. So that's just an example of how it is that seeing something in the brain might then change your psychological model, because we're bringing everything we know at the brain science level to the psychological model in that case. Ultimately, we care about what people do. We care about that level of description. But knowing about *that* level of description adds information to *this* level of description.

Nersessian: Well I think the findings in neuroscience that are of the kind that Liz just mentioned are extremely important and I think whatever new findings are coming out of neuroscience are fascinating and exciting. But I also know that there is kind of over-excitement about neuroscience these days, so much so that I myself was involved some fifteen, twenty years ago, starting in neuropsychology, thinking that neuroscience was going to enlighten us and tell us a lot about psychoanalysis. I remember very well that in my excitement, I decided to write a paper about how neuroscience and its theory of emotions was going to help psychoanalysts with their understanding of Affect Theory. And I presented the paper in New Orleans and I was feeling as I was presenting that I am really talking a lot about nonsense, because I am confusing my hopes and my wishes about what it can give us with what it is actually giving us. I read Colin

Camerer's paper, which I have here, and I had this very intense reaction that it was exactly the way I was feeling. This is the paper "Neuroeconomics: Using Neuroscience to Make Economic Predictions." He seems to be trying very hard to convince you that neuroscience is going to do great things for us, but I still don't know what it's doing and how it's going to do it.

Glimcher: I wish Colin were here today. I have to say that speaking for myself—but I suspect Liz and Alberto feel the same way—we have exactly that sense reading that particular paper.

Nersessian: Oh, good.

Audience: He was the second president?

Glimcher: He was the second president, yes. Colin is a great guy—don't get us wrong. He's a major figure in economics today.

Phelps: He's wonderful.

Glimcher: He's really serious—he's learned a tremendous amount of neuroscience; he has a lab; he does scanning experiments. Colin is even more optimistic than the two of us and he has made promises about the future, which I'm always cautious about. You know, Liz and I have an expression between us that a lot of people use, which is when we pick up the journal, *Neuron*, or the *Journal of Neuroscience*, and there's a colorful brain on the cover and it's so beautiful and you read the paper and it's like they took some guy and they read him a disgusting story while he was in the brain scanner and then they made him watch a clip from *The Sound of Music* and then they took the difference in brain activation between those two. And this is the neural structure responsible for being disgusting.

Phelps: Or for Julie Andrews.

Glimcher: Julie Andrews, right. They do *The Sound of Music* with and without Julie Andrews and then they find the Julie Andrews Center. So Liz will say to me, "Did you see that paper on Julie Andrews on the cover of *Neuron*?" And I'll say, "You mean the spots-on-brains paper?" There is this tendency right now to produce these beautiful spots-on-brains maps, and for reasons that are a mystery—I know to the two of us—*Time* magazine, *Newsweek*, they love this stuff. All you have to do is produce a colorful brain. I find economists particularly credulous in this regard—if you will excuse me, Alberto—

Bisin: I've never done it, so you can do it.

Glimcher: David Laibson gets up with the spots-on-brain picture and says, "This economics theory is proven," and you're at the American Economic Association meeting, there are five hundred people out there who eat partial derivatives for breakfast, but you show them a colorful brain and they're like, "Oh, so that pretty much settles it."

Bisin: No, I agree with all Paul has said, but there are two things: one I mentioned, which is that economics is incredibly difficult and dry and it's become like that very recently. Part of it is this

rational expectation stuff, which makes it really hard mathematically. So people are not used to thinking of economists as—I mean we're not—but as physicists. People realize—by people I mean journalists, people in the street who don't do this as a profession—that when they talk to a physicist, they don't understand. That makes sense, they don't understand. Physics is hard. But when they do it with an economist, they say, “This guy must be saying a lot of things to confuse me. It's got to be simpler than that.” So I think these people know this and they just feed it to the public and the journalists. Now you finally understand they were actually just cheating you. You know, they said economics was hard. Instead it's very easy, let me tell you. I don't think the economists actually buy that much into the spots on the brain. I remember early on when you and Aldo Rustichini, a friend of ours—many people came to give talks in economics, to the people in the economics department, about neuroscience. After a while the people—the standard economists—would just say, “I'm not going to come; they're just going to show me some brain pictures. They're very cute and nice. The brain moves and colors up and it's beautiful and then they show me monkeys and that's it and I don't want to see it.” So I think economists react badly, too badly, to this. I mean there's a lot of interesting stuff that they don't see because they say, “Look. Look at the brain, which moves.” For some reason, by the way, they're very good at this picture of the brain which moves and flows and stuff. We're not good at doing—

Glimcher: You don't do graphics very well.

Bisin: Yes, we don't do graphics very well.

Glimcher: Well, certainly the economists that we like and trust are not easily taken by this stuff and the quality of the debate has improved in the last three years, just beyond words. The first meeting of the Study for Neuroeconomics was a series of presentations by noted economists and neurobiologists. Of the twenty talks, maybe four had data to present. The remaining sixteen were basically this paper you're talking about. Like, “If I could measure this, it would completely revolutionize this, and here's how I'm going to measure it.”

Nersessian: But one of the things in the paper is if you give somebody oxytocin, they become more trustful. So what are you going to do—if you want to cheat them, you give them oxytocin and have them trust you and take their money? So it seems like the finding doesn't lead you into a place where you have a solution to the issue. So oxytocin increases bonding of the mother and the child, increases trust, it does a lot of nice things and a lot of opiates do a lot of nice things as we've recently found out, but how does that help us in terms of determining or affecting the person's economic behavior, which is what neuroeconomics hopefully will one day do, no?

Glimcher: No, that's not my goal.

Audience: There's actually an answer to that—economic studies have shown that societies and cultures where the trust level is high do much better in terms of everybody else's welfare.

Nersessian: But what are you going to do, give everybody—

Glimcher: I don't think that's Ed's question.

Nersessian: What are you going to do, give them hormones? Give the society hormones?

Audience: That's a question for culture. But if you know the biological effects, culture will—

Nersessian: Yes, but economics has to do more with culture than the biology of an individual.

Audience: Well, that particular question, yes.

Phelps: I just want to make a couple of points about the persuasion of neuroscience and particularly what Paul mentioned—everyone seems to love these spots-on-brains. And this is not just happening in economics, you know, it's happening in law and social psychology and all sorts of other areas where now everyone sees a brain picture and they think somehow they've solved it.

I was part of a panel once at the AAAS and it was called the “Cognitive Paparazzi” and it was about just this phenomenon. One of the psychologists who spoke was a man named Frank Keil, who's a professor at Yale and who's a developmental psychologist. He's done this beautiful research showing that if you show people a picture, they over-interpret their understanding of the phenomenon. And so it doesn't matter if you're describing helicopters and there's a picture of how the propellers are connected or if it's a picture of a spot on the brain, saying, you know, this is the “Julie Andrews area.” It doesn't matter. But if you ask people specific questions about what they actually know, then they realize they don't actually understand it. They're pretty good if they only got a verbal description. Then their ability to estimate what they know is much better. So I think this has something to do with why all of a sudden people feel like they have some understanding when there's this beautiful picture, and why it's influenced fields more than perhaps the data at this point would merit.

But of course the other issue, as we all know, is really, you know, “this is where our behavior comes from.” I think people give some level of primacy to this level of description as opposed to other levels of description. I think both of those phenomena come into play and I think we have to be very careful about saying what we are actually learning from this. I think the jury's still out when people talk about even psychology or cognitive neuroscience. The study of cognition has looked at the brain for a very long time, and in the first five or ten years you could easily argue that we didn't learn a whole lot about the study of cognition from studying the brain. Certainly since then, I think almost everyone would disagree with you. So I think this field is very, very young and I think we'll start to see where this is going to benefit economics in ways I don't think we can predict right now. We can't argue this hasn't had a huge influence on things like psychiatry and other fields, and I think that's going to be true for economics in ways that may be hard for us to answer today.

Glimcher: There's this question of what's the goal. I'm finally a scholar of human behavior—I want to understand why people behave, why they make the choices they do and how they work at many different levels. You can really tell I'm not an economist when I'm hanging out with economists. I don't care about making the markets more efficient. I don't actually care about improving the way people save or really even knowing whether they save enough, obviously. So I do share your skepticism. The oxytocin results are really striking. To me what it means is a

behavior as complicated as trust can be reduced to a single hormone, at least to some degree. That's an amazing neuroscientific finding. By understanding where all the oxytocin receptors are, we can leverage that to understand a little bit more about the neural basis of trust. So all of those are good. If I can show you that those brain areas are active during some other kind of exchange, I can relate that new behavior to this underlying neurocircuit, and that allows me to improve my model of that behavior as well. Does this mean we should make everyone carry nasal inhalers and use them every hour? I don't know; you'd have to ask a welfare economist. And I'm definitely not a welfare economist.

Nersessian: People are restless to ask questions. Let me just do one other thing and then we'll go to the audience. As you may or may not know, I had a lot of communications back and forth with Professor Kahneman, who decided not to come. One of the reasons why he said he didn't want to come is because he thought one of his friends would be criticized and he didn't know enough about the work of the friend to be a good person to defend that friend. The friend is Jonathan Cohen from Princeton and, as you know, Kahneman is from Princeton. What is it about the work of Jonathan Cohen that the two of you were going to attack so harshly?

Phelps: So since he's not here, let's just jump all over Jonathan.

Glimcher: Turn off the web feed and we'll answer that question. Just kidding.

Nersessian: Scientifically, what is your criticism? What is it you dispute or disagree with?

Glimcher: I can't explain it. It has its roots in economics.

Bisin: Okay, so—

Phelps: We'll all chip in after you're done.

Bisin: Of course, I don't know enough. Jonathan Cohen is a neuroscientist and I don't know enough about his work. But let me say what I think is interesting about his work and what I don't like about his work, to give you an idea of what it's about. From the economics point of view—he'll take issue with the data. So the idea here, as I said, is that as economists, we are looking for a unifying model of choice that is close enough to rational choice, because we think rational choice works fine, but also able to explain a bunch of things that we see that do not agree with rational choice. One possible unifying model is this two brains thing. Again, I'm talking about this as an abstract model, not about anything that is really happening in the brain, because I don't know anything about what's happening in the brain. As a general abstract model, this is a model where there are two parts of choice, one that's more impulsive and the second that's more controlled or cognitive, or however you want to call it. So there are many examples that you can fit into this. The typical example that economists have in mind is people who get run over by a car when they go to England. So they look at the wrong side of the street and it's not that they really want to be run over by the car, they just make a mistake. So what's the mistake? You can rationalize this by saying, "Oh, you know, there's some part of the brain which is somewhat automatic, which makes them do something—look to one side of the street—and if you point

their attention that this is London, you should look at the other side, they'll do it." But you have to point it to their attention because otherwise, they don't do it.

Jonathan Cohen looked at a bunch of models of behavior of this kind. I think one of the things he had in mind—you'll know much better than I do—is this typical psychological experiment called "Stroop," where you feed people words like the word "red" written in blue ink and people automatically read the word, but if you ask them to name the ink color, they'll name the ink color, but something else goes on in the brain. So he has a model and a theory about whatever goes on in the brain.

Some people in economics have tried to think about this general abstract model—"Stroop" and the way you look when you cross the street in London—as models trying to understand this interaction between impulsive, or automatic, and more controlled behavior. So you can think of something like that when you think about choice over time, this issue of present bias. You can think there's the automatic part of your brain, whatever that is, that makes you grab something very quickly because that's what you're used to doing, but then there's something else that says "hold on." And the interaction between these two works in some way and you have a theory—some people have used this to think about tails and distribution. People don't assume automatically that there's nothing in the tail but when you make them think, they'll say, "Yes, there's something in there." So this seems like a reasonable, possible general model of behavior.

Jonathan Cohen has done some of his work in psychology and then people in economics—David Laibson, who we mentioned before—have thought of this as explaining, in this particular case, choice over time. So they have some experiments, published in *Science*, where they look at differential activity in the brain in choice environments where people either have a present immediate choice and environments where they don't, and they try and look at differential activation and they see something that lights up when you have the immediate reward, and things that don't light up in the other case. They call them "beta areas" and that magnificently fits with the model. So this second part, as you can see, I don't think it makes any sense—there is no model to fit. I think of that paper a lot along the lines of the paper that you mentioned of Colin's; it just claims way too much. But from my point of view, this vague model of the two brains—it's an interesting model. It's an interesting way to go. I think the debate that had to do with Kahneman—Paul and Liz will go on to this—they actually do not believe there's anything like these two components in the brain, and so in some sense they're even more against that paper.

Glimcher: For me, this is an awkward place to lay out this story, amongst all of you. Because what I see this as is something that's very closely related to traditional Freudian theory. The notion of two different processes inside your mind or skull—presumably the same thing—which compete with each other for control of your behavior is an ancient idea, of course. We could go to the Greeks. Of course, this would be a good place to go to the Greeks, also, through Freud. In the 1950s and 60s, as many of you know, there was a sort of grand battle in neurobiology over whether there would be a very simple mapping between a chunk of the brain—at the time called the limbic system and the Freudian notion of the Id—and a chunk of the brain known as the neocortex and Freudian notions of the Ego and Superego. That debate sort of ended it basically, and the decision was that that simple mapping was wrong because you could remove most pieces of the limbic system and show no effect on very emotional, automatic, impulsive behaviors. So it

seemed very hard to believe that the structures that were supposed to be responsible for the Id could be removed and Id-like behaviors would remain.

These ideas, I think, are an important part of our culture and they have surfaced in economics recently in the form of two processes. How do you explain the fact that a person who is a drug user will tell you when he's sitting in a room like this that he doesn't want to use drugs and he'll act in ways—if you say, “Well, I can set it up. Tomorrow you'll have drugs,” he'll say, “I don't want it, I'm trying to stay off them.” So you say, okay, his preference is not to take drugs. Then you take him off to his pusher's house and he looks at the door of the pusher's house and you say, “But you don't want any drugs, right?” And he says, “Oh no, I want drugs. Please give me drugs.” So you can't reconcile that because he seems to have switched his preferences, right? Well, so here's the solution: there are actually two people in his head. One of those guys is in control when he's sitting in this room and that guy doesn't want to take drugs. And there's the second person in his head; that person wants to take drugs and something about standing in front of his pusher's house gives that second person control over the behavior of the individual. I actually have no problem with that as an explanatory framework.

There has been a rush to use neurobiological evidence to validate those claims and in a number of economic models, people have said, “Look, the two competing models are those that involve multiple agents, multiple selves, multiple actors hidden inside your head, and those which do not. So what we will do is design a scanning experiment that will take a month or two to run and that experiment will yield neurobiological evidence to the following type. Under conditions where Agent 1 is active, where the impulsive guy is in control, I'll scan you. When the non-impulsive guy is in control, I'll scan you. I'll take the difference in brain activation patterns between the two and if there's any difference at all, that proves there are two people in your head.” Now, as a cautious neuroscientist, that is not a clever experiment and it doesn't prove anything.

The experiments that would actually validate those hypotheses are very complicated. Liz and I have both worked on experiments like that—they take years and years to complete and the very first of them are all now being published. They're about three years behind the leading edge papers on these same questions. They are actually all coming down the opposite way from the first-time papers. They're not showing clear evidence of two systems—one switching on, the other switching off. Now that is not to say there are not many things that contribute to our preferences and decisions; it's not to say our preferences are always stable. It's just to say it's not going to be as easy as the limbic brain—now, I don't even know what that means anymore, but I read that word all the time in *Time* magazine.

Phelps: We're trying to ban it.

Glimcher: Liz and I are actively trying to ban it. We do not allow our graduate students to use that word anymore. They can name any brain structures they want, but we don't know what the limbic system is, so they can't use that word.

Nersessian: I should say that unless there is a misunderstanding, in psychoanalysis, you wouldn't say there are two brains. Leaving the drug addict aside, because he could sit in front of you and lie a hundred times and then go out and take his drug—that's a whole other problem. But a lot of

people who don't have drug addiction can have totally competing ideas. Some guy breaks up with his girlfriend and he thinks this is the best thing; he wanted to break up with her for two years and every time he broke up, she came back and said, "Oh, I want to have the relationship again." So the moment she dates somebody else, he's running after her. Now he knows he didn't want to be with her and he was breaking up with her, but now that she's dating somebody else... So psychoanalysis says, well, let me understand why there's this behavior, because the psychoanalytic conviction is that there's an understandable connection between what on the surface may appear as opposing types of behavior. So, just to clarify the psychoanalytic position.

Glimcher: But in *On the Ego and the Id*, Freud tells a story that sounds much more like—when you read it through the first time, it really seems like there are two agents. And I think that's the view—

Nersessian: Freud is trying to do the same thing as an anatomist does with the brain—there's the frontal lobe, there's this, there's that. If you look at the diagram in that paper, he has it completely connected and interconnected—in fact, the Superego he connects to the Id. So it's complicated and you can't look at it in that kind of—

Glimcher: But as you know there's been a long history in neuroscience of trying to do that and I just think this is the most recent example of that.

Phelps: All I wanted to add to this was that I think the work Danny's done, which unfortunately he isn't here to talk about today, really started to go away from the rational-choice model and come up with another way that people make choices, and it's been a very useful heuristic to describe this as two systems. I think that's pushed the field forward quite a bit. There are properties of one system and properties of another system and they work in somewhat predictable ways, and that's been a great psychological level of description.

But if you start to look at the neural level of description, there are a lot of subtleties and variations that come into play. The interactions I'm most interested in are emotion and cognition because of my area of expertise, and there's been this argument that you make a multitude of choices based on your emotional system versus your cognitive system, and there are different choices and there are two systems that interact. I think that neuroscience evidence doesn't support that; it suggests a much more subtle and nuanced view. Emotion is not one thing; it's not represented in one way. These systems that underlie emotion and cognition aren't independent in the brain. There are brain structures that are more or less involved in emotional behaviors, but they interact extensively at every level of analysis with those that are considered cognitive behaviors. So it just seems to say that this dichotomy that we've used to understand psychological behavior and has been very helpful in driving the field forward, and now we have in these two parts of the brain, completely takes away what neuroscience has brought to the table. Neuroscience is going to give us a much more subtle and nuanced view of how these things work and that's going to go back and now refine our psychological models. It doesn't mean they're not good or important psychological models. Right now, they're just somewhat incomplete in understanding all of human behavior, which is expected as the field progresses.

Nersessian: One more comment and then—

Brown: One reason why this side of the room is more relaxed on this topic than the other side—to be honest, economics, macroeconomics, all the things we use to explain how markets work, how GDP grows, is not solidly grounded on how an individual makes decisions. It is not. At the macro level it has a life of its own. Therefore, economists tend to be willing to say, “Let’s say there are two minds and it could be interesting.” So a lot of the economists joke, “Let’s assume if you can’t open it. It’ll be all right.” The economists are quite relaxed about it because it really doesn’t build up and change your understanding of how individuals work. Economics doesn’t rise and fall on this. Macroeconomics, most of economics, is not really grounded on those individual behaviors.

Audience: In physics you look at the science using this interplay between theory and instrument and experiment. So you’ve alluded to this—I’d like to know what the experiments agenda for neuroeconomics are. You talk about subtlety and nuance. What are the three or four experiments—the Michelson-Morley, the electronic experiment—what other kind of experiments do you think are really going to make these people sit up and pay attention, not unlike the way that, as Danny commented, the Freeming experiment, the Herb Simon stuff, subjective expected utility, affected and influenced quote-unquote rational choice in economics.

Phelps: Part of the problem is we all do our own experiments, so I’m going to tell you my experiment. I don’t know the ones that are going to make the economists sit up and notice; I feel that’s a little bit yet to be determined.

Glimcher: I have a more precise answer. You know, I’m much more aggressive. Here’s the real answer: the goal is the generalized neural theory of decision. That is, a piece-wise explanation of the architecture of the brain and how those elements interact computationally to yield free human choice behavior.

Audience: But why do you want to know?

Glimcher: Good question. I don’t know. I’m a kid who took stuff apart.

Audience: It’s a simple question; you’re avoiding the answer.

Glimcher: Well, let me answer the second half of his question before we go any further. My answer is I want a generalized neural theory of decision-making in the same way that we have a standard model of the atomic structure of the universe which again, we don’t need—I grant you that—but I love that there’s a standard model of matter. I don’t know.

Nersessian: You’re a scientist.

Glimcher: Yes, I’m a neuroscientist. Now here’s what I think are going to be the qualities of the Michelson-Morley class experiments: they are going to have to start with a problem economists are interested in and have not resolved at the microeconomic level. Alberto knows exactly what I think is the candidate experiment and it’s one to explain this impulsivity. There are going to have to be a series of neuroscientific experiments which yield novel behavioral predictions. This is

really critical because what you're saying when you do that is that I can use the mapping and a study at the neuroscientific level to make a behavioral prediction that has never before been made by an economist and that an economist would not have thought to make. That behavioral prediction has to be validated, and then the last piece to make this a homerun is it has to undermine a major axiomatic theory.

Bisin: That's important.

Glimcher: And suggest an alteration to that axiomatic theory. As Alberto knows, there are two groups in my lab who are working on experiments that have this flavor and they may or may not succeed. These are the experiments that are really going to, I think, make it impossible to look away.

Audience: I would like to know more about economics. You mention about this lottery gentleman who wouldn't take the hundred dollars, which is irrational, but why don't you think about the forces that are getting him and so many others to go to the lottery. They're not going for a hundred dollars, they're going for a million dollars or whatever they think they might get. I don't know what that has to do with what part of the brain.

When you talk about economics, I usually think about totality, and this is mainly talk about consumerism, rather than who are you. In other words, being determines consciousness. I won't go into that too much, but you know, what about the people who are making this lottery thing? You discussed the automobiles not being developed well enough. There are a lot of mistakes there. Why does that happen? Certainly there are mistakes in the brain, but there are also such reasons as maybe more of a profit wants to be made, so therefore sometimes they may cut short some of the things that they should do. Also, if you're going to look at prices or cost, if you really want to know the value of things, you're not going to be able to see that, so what does that have to do with this?

Some people learn more than others because nobody's the same, we're all a little different. Some learn more than others and maybe that has something to do with what's in the brain, but the macroeconomics is so much more important. Let's look at these forces that are making people want that way. Or even talk about retirement savings—what are the motivations, why does it happen the way it happens? I mean, it's not about something in the brain; there are other reasons for it. Who gets more, who gets less?

Bisin: We've not been discussing this too much because economists do this and we all agree—Paul agrees—that there are complicated explanations about savings, which have nothing to do with the brain. Economists do that well, but there are aspects that are very important. When I reacted to Paul, saying I don't think it's true that people under-save, I don't think it's true, but you see actually a lot of evidence, which you can interpret in terms of under-saving, and if that's the case, let's discuss this and try to understand it. Now all the other stuff that you're mentioning, I completely agree. It's very important. I think we haven't discussed this because, as you said, it has less to do with neuroscience and much more to do with standard economics.

Phelps: I was simply going to say it's always happening in the brain; everything is happening in the brain. Now you're talking about factors that are not sort of individual factors or cognitive factors, and certainly adding information about culture and social factors is where the study of the human brain is going. So I don't think because you're talking about things that are not just individual choices that have influence from a range of factors that understanding things at the neural level won't help us understand those decisions, as well.

Audience: I would like to change the frame of reference a bit. I think it would be much clearer. You see, I think all people make rational choices all the time. Human beings make choices all the time. In fact, all animals make choices all the time. Is something good or bad, should they go towards it or away from it—they make choices all the time. It's based upon biological criteria, such as not getting killed, being able to reproduce, being comfortable instead of in pain, on the basic biological level. The microeconomists got all confused because the people they're studying didn't seem to be rational. It's because their definition of rational is much too narrow, their definition of rational being what the person is consciously aware of and is consciously reporting. If you want to understand the individual, it's much more complex, and we psychoanalysts do it better than the economists because we try to find what they're really afraid of that they don't realize, what they really desperately want, but can't admit to themselves. If you factor in all these things, then the choices are rational.

For example, there's a very famous economic game—the prisoner's dilemma. A guy in the next booth, who you don't know, tells you he has some money, a hundred dollars or whatever. He's going to split it with you and if you agree to the split, you'll get some and he'll get some. So if the guy in the booth says fifty-fifty, you're overjoyed and you say yes. If the guy in the booth says sixty-forty, you're happy, fine. If the guy in the booth says he's going to keep ninety and you're going to get ten dollars, very often the person becomes indignant and says the hell with it, no one gets any money. Yet that same person, if you would say to him, "Imagine this scenario: you found a wallet with a hundred dollars. You give it back to the person and he gives you a ten dollar reward. Would you do it?" And they say, "Oh, yes." What's the difference—emotional things about narcissistic injury, fairness, whatever. So if we take these things into account, we can understand the microeconomics. Macroeconomics is easier because we have people of different ages, with different requirements and different unconscious needs, and they sort of average out—the irrationals average out. So the macroeconomics people are dealing more with rational people.

Glimcher: So let me ask you one question before we go any further because I want to understand better. Would you say people never make irrational decisions or that people predominantly make rational decisions?

Audience: I'm reluctant to say never, but certainly predominantly.

Nersessian: I think there's a difference between explainable and rational. I think your example is more that it can be explained.

Glimcher: I don't agree. If we had had Danny here, this would have been a different discussion because Liz and I would have been beating up on him and we would have looked—

Phelps: I would not beat up on Danny.

Glimcher: —much more pro-rational choice. Here's my favorite puzzle like this. I used to get told this—it's a question, as Alberto said, of defining what's in the utility function. That's the whole point. George Lowenstein, a noted economist and psychologist, once gave me this puzzle: one of your graduate students goes off to a conference at some remote city. She's engaged, and she meets some nice guy and they go to a bar. She's sitting at the bar, trying to decide whether or not to go back to the hotel room with this guy. She does. When she wakes up the next morning and she's completely distraught and miserable and crying. How do you explain that, because it was an irrational act? And George has argued to us that that's because something emotional takes over control of her behavior. Now as a biologist, here's my read on that: we evolved to maximize fitness. She is just doing exactly what she was designed by evolution to do.

Nersessian: How come half of them may do it and half of them may not do it? Evolution?

Glimcher: So then I have to explain something to you from game theory called the Mixed Strategy Equilibrium that requires unpredictability in behavior in order to maximize welfare of the agent. So there's a very good answer, but it would take a little while to work through.

Nersessian: I don't think so, but I'll take your word for it because I don't think you have the answer. I think the only way you can have that answer is if you have the patient in your office and then you can draw the line that will explain to you why this person did what she did and another person, very similar in many ways, did not.

Glimcher: To me that is a parallel, perhaps even equivalent, explanation to the game theoretic one.

Audience: We can't do it forwards, only backwards. So if the person comes in and says, "I was so tempted, but I didn't"—

Glimcher: We could do it forward in the aggregate.

Audience: In the aggregate, yes. In mammalian models and in bird models of monogamous animals, you find this, what you just described. You find it in human societies, too. So that means that there's a biological possibility, a propensity, but it doesn't mean every single person would succumb without—

Glimcher: This may be helpful. There's a very, very famous puzzle by the economist and physicist and everything-else-ist, John Van Neumann. Sherlock Holmes is in Victoria Station in London and he jumps on a train, which is going to stop at Canterbury on its way to Dover. He knows that Dr. Moriarty, his arch enemy, has got to the station just after the train has pulled out. So he's okay on the train. Moriarty has a gun and he will kill Holmes if he can find him. So here's the decision Holmes has to make: if Holmes gets off at Canterbury—oh, sorry I left out one other fact: Moriarty's rich, which you're supposed to know, and he can charter a train, but he has to know where to stop, because if Holmes gets off at Canterbury and he shoots to Dover,

he'll miss Holmes at Dover. But the opposite is also true. So what is Holmes to do? If Holmes says, "Here's what's going to happen. Moriarty is going to think I'm going to Dover, so I'll get off at Canterbury." But of course, Moriarty is smart enough to know that, so he'll work through that same logic and he will also go to Canterbury, by which logic Holmes should of course then triple-cross him and go to Dover, by which—you get the idea. And so the solution—I'm simplifying it a tiny bit from the way Van Neumann presented it—Van Neumann argues what Holmes is actually supposed to do is flip a coin. The safest thing he can do is flip a coin—there's a fifty percent chance he'll get out alive. A lot of games have this quality. Although, they don't all require an even coin; it's always been true that a weighted coin gives you a better yield.

So as an evolutionary biologist, here's my answer: in the aggregate, thirty percent of people should do this. It can be predicted by game theory. Now you ask a great question, which is what is the psychological and, mapping one step down, neurobiological instantiation of that randomization process, or the social instantiation of that randomization process. I am deeply interested in that question. I think we want to build an answer to that at all levels.

Audience: I have a degree in both economics and finance, and I'm familiar with Dr. Nersessian and I'm familiar with the good doctor that preceded me. We're having an interdisciplinary conversation, here—a dialogue between different disciplines. I think it's very laudable that you want to understand economic endeavors scientifically and I think that's where the neuroscientists are sort of stepping in here, as if they have a greater rationality. The field of economics has been around for an awfully long time and there are lots of histories and approaches and attempts at understanding human behavior. Human behavior and economic behavior are ever changing. Group behavior situations change immediately. Uniform knowledge isn't fair. There are trading algorithms and all sorts of new approaches to the field of economics. The number of people that try to understand the markets and the rationality and the irrationality—if medicine had as much money as what is spent on research in the field of finance... I thought this gentleman's comments were very good in terms of it not being just fear and greed. There's understanding of cars and how cars work and should cars work and what do consumers want. And economics, I think, is no longer sheer economics. It's now political economics, because there is no such thing as the science of economics—it's so intertwined.

But the thing that bothers me here: when we talk about rationality and the neuroscientists are sort of talking about layering and they're bringing forth some new scientific understanding that's going to subsume the understanding of the economics in some way or to create some better understanding, I think we all look forward to it. Those of us who have been coming to the Philoctetes Center, we've heard a lot about neuroscience. We've heard about some of their promises, and it seems to me that it's rather in an adolescent period. We all hope that neuroscience gives us some greater understanding, but it seems here a little reductionistic and a little overly ambitious.

To try to understand rationality—we haven't talked about greed or fear; we haven't talked about optimism or pessimism; we haven't talked about conservatism in terms of behavior or flexibility; special, general masochism, all sort of psychological terms that do enter into this decision. And by the way, there is a course up at Harvard Business School and a man, Abraham Zaleznik, who's a psychoanalyst, is trying to teach something about behavior theory relative to

understanding psychoanalytic behavior of people and finance. That hasn't even been mentioned here. So I'm just saying I admire neuroscience, I admire your claim. I think it's a little premature to make these claims and to think that you can really undercut the huge amount of efforts and behavior that are trying to understand what we call rational. I just leave that point—what is rational? Is rational some neuroscientific term that we can say, “That's rational because the layer A, B, and C is in function.” Or is masochism and psychoanalytic terms, in terms of what is rational, a much better explanation for what rationality is.

Glimcher: Can I just say one thing because I think I must have misspoke terribly.

Phelps: I feel the same way, hearing that question.

Nersessian: You should both speak then.

Glimcher: I must have said something that led you to that conclusion, for which I apologize. The reason I have spent the last seven years trying to become an economist, and the seven years before it trying to become a psychologist, is not because I believe neuroscience will subsume economics or psychology. In fact, if you read my papers or my book or any of the things I've written, what I hope you'll see is that all three of those levels coexist in good theories. And my interest is in no way to replace economics. If I were going to do that, I would certainly not have wasted the time to be considered an economist by my colleagues, to try and publish in economics journals. My goal is to understand this process of our behavior in as many different ways as possible because it is my conviction that understanding it in more ways is more powerful than understanding it in less ways. When all is said and done, what you want is an explanation at every level, and as many ties between those explanations as you can get. They'll be incomplete; the mappings won't be perfect. That will be the strongest explanation of human behavior.

Walling off these disciplines, which is the alternative, strikes me as ostrich-head-in-the-sand “my discipline's fine.” Now, neurobiologists mostly believe that—don't get me wrong. I came from a tradition where they said—when I wrote my first economics paper for a neuroscientist, *Nature* wrote back and said, “Are you out of your mind? You want to write a utility function in a neuroscience paper? That is a waste of time. Neuroscientists understand choice completely and we know exactly how to study it and ties with economics are a waste of time.” So I feel so passionately about this because I spent so much time on it. I am not in any sense trying to replace economics or psychology; I am trying to unite them with neuroscience.

Phelps: I completely agree with everything Paul said. I want to add that I don't think neuroeconomics is in its adolescence, I think it's in its infancy. I agree—I felt like when you asked your question that you misunderstood everything that I was trying to say today, which is to say I'm interested in the psychological level of description and I think neuroscience can inform that. When we actually talk among ourselves about what motivates our behavior, we're never going to talk about our neurons. People talk now in law about the “brain excuse” or something like that, but ultimately, we care about behavior and the actions people take. So I don't think in any way are we trying to subsume the economic level of behavior; this is adding something to that discussion that could be informative.

Bisin: Could I just add one little thing from the other side? I made at some point the example of the glasses and I said that's rationality. That's supposed to generate the reaction that we have a very stripped down notion of what rationality is, which has been incredibly powerful on one hand, but of course it's incredibly reductionist. So it seems to me that what they are doing—what neuroeconomics is doing—is actually going the other way, to enrich the notion of rationality that we put into our model as opposed to what the question seemed to imply, to not be rich enough. The notion of rationality that we use is extremely simple and small, and what they're doing is just enriching it. So it seems to me the other way around.

Audience: One of the speakers earlier said that people had evolved over hundreds of millions of years to make rational choices. I think it might be more precise to say that collections of genes have evolved over hundreds of millions of years to make optimum choices, and that those optimum choices for collections of genes might have involved irrational choices or variance from rational choices by individuals bearing various portions of those collections of genes. I wonder how the speakers in any way feel they can adapt the rational-choice model to speak more broadly to the collections of genes and, in particular, are there any particular cultural, ethological, or other ways of grouping human behavior so that you can see variance in a broad way across those cultures or groups in the way decisions are made?

Nersessian: Does anybody want to tackle that question?

Phelps: No. I guess I never thought about that.

Glimcher: I assume you're thinking of Richard Dawkins' powerful ideas about the selection occurring not just at the level of species, as Darwin had originally proposed in the 1860s, but occurring at the level of individual genes, and of course we now know that selection occurs not just at the level of individual genes and whole species, but gene clusters. The evolution of data now about the hap maps and stuff like that seem to indicate that multi-gene loci evolve together. I spoke too quickly when I said that. Selection is definitely occurring at many levels, for sure.

As a back-of-the-envelope model, we could take Darwin's and think about selection at the level of the species. It is certainly not the only level at which selection occurs. One thing that's striking about this is that some behaviors which appear irrational when we think about what people tell us verbally they want—and this is natural for analysts, I assume—make perfect sense when we look more broadly either at the individual's welfare or look at a species' groups of behavior. Some of them won't explain easily this way. So the idea that I want to get across is rationality in the economic sense is about maximizing utility function. What goes in the utility function is a little underspecified.

One of the things that we might begin to think about when we think about utility functions is our evolutionary heritage. That's sort of the point I wanted to push, and I want to back away from selection at any one particular level, whether at the meme level, the deme level, the individual gene level, the hap map level, all of those are going on, of course.

Audience: Are there any experiments that show variation across large populations these kinds of rational decisions?

Glimcher: There's a great experiment that's unpublished by Read Montague's group at Baylor College of Medicine. And in this experiment, Read has his subjects play what looks loosely like a game. For those of you who grew up in psychology departments twenty years ago, what they're actually doing is an experiment of the Harvard psychologist, Richard Herrnstein. They're picking between two levers. The levers pay off money with probability and distributions that are a little complicated. It turns out that there are two good solutions to this. One solution is to half the times push the left lever and half the times push the right lever, basically in random order. There is a better solution, but it's hard to find. What you have to do is push the left lever nine times for every one time you push the right lever. The reason this is hard to find is that all the intermediate ratios stink. And it turns out that only half of Read's subjects find this better solution. Half do and half don't. So they drew blood from these people and they genotyped their dopamine receptors. And it turns out that all of the people that find the best solution have a particular isoform of the dopamine D-5, if I remember right, the subtype of receptor gene. And this particular isoform seems to be associated with propensity toward risky behavior.

Audience: The intellectual power that is represented at the table is overwhelming. I'm overwhelmed by it and yet it goes off in so many different directions, it's hard to focus, as you will appreciate. I also appreciate that it's very difficult to focus on any one thing, given what's at the table. It's too much to expect from it. At the same time, I want to make some random comments rather than anything terribly focused. My name is Michael Gropp and I've spent my time in ethics and moral theory for the last forty years. I'm coming out of retirement now to think again, which is why I'm here.

I'm remembering a philosophy professor when you tell us the story of Sherlock Holmes, one of my great heroes, and Moriarty, his archenemy. You say that he should have flipped the coin in order to make a decision as to what the best thing was to do. There's a professor by the name of K. D. Irani, as I remember him, at City College, the poor man's Harvard at the time. Irani stood there like this and we were discussing free will and determinism. He had his finger out like this and he said, "Which one of you can tell whether I'm going to do this or this with that finger?" And there was all kinds of debate and discussion for the next three or four weeks on the topic and he ended up by saying at one point—which impressed me a lot and is still with me and I'm bringing it here to the table—if we knew everything about K. D. Irani, we could predict what he would do. At the time, there was no discussion of neuroscience—I'd never heard the word before. I knew there were neurons, but I didn't know what they meant. Now you were describing them, but this was some time ago.

So I begin to wonder what you mean when you say that you, Paul, want to understand everything there is. When I sat behind you, I couldn't stop the challenge at the moment—to me it was a challenge—as to why you want to know all these things. What is it that you're interested in to know, and that goes for all of you at the table, and me too, living in the real world. What is it you really want to know and why do you want to know that? Then I heard the word "control" mentioned once or twice, and from a sociological point of view I left it alone because I was concerned that we'd be talking about a totalitarian question, which wasn't at the table. I didn't really want to introduce it, and yet I do. Where do you want to go with all of this? Maybe this is

not the time to answer all of that, but here you are in the various university disciplines, promoting ideas and organizations. Promoting ideas with what in mind? Where is it going to go?

Some of you remember that book by Vance Packard, *The Hidden Persuaders*, and I think it's from that book that I quote. Do you remember the word "subliminal?" Do you still use it, I don't really know. But the question was from the economist's point of view: how could you sell more Coca-Cola? That was the issue—economics. Economics in the movie theatre. So what they did was to put the subliminal phrasing on the bottom of the screen, below your consciousness level, but within the framework of reaching you on some level, and during the intermission of this very dull movie, sales of Coca-Cola went up fifty or seventy-five—I don't know how much it was—percent. Because you were approached in a way through understanding your neuroeconomic content. However you want to describe this—biologically, psychologically, psychoanalytically, whichever way you want to go—you are going to change the behavior of people. And it's what I wanted to know and what you have to ask, it seems to me, as moral agents—which you are, because you're teaching people what they should be thinking about, and making every effort to convince them that your position is a good one.

The interplay between you two gentleman is wonderful, but it masks something, I would suspect. I think it masks the difference between you in making authority in the world, up to a point. And that's all right; we're all there. So I didn't get up to scold—I feel like I'm scolding, but I don't mean to scold—and yet it concerns me a great deal about what you mean by control. That's the deeper issue. You're now about to give another tool to the community that will give somebody an opportunity to have control over what I purchase. If the box of Tide comes in a blue box, I'll say it's good and mild, and if it comes in a red box, I'll say it's harsh—I don't like the chemistry. If it comes in a yellow box, I won't buy it anymore because it burns my hand. I could go on indefinitely, so I'll stop at this point. Mine is just an observation. Thank you.

Nersessian: Does anybody want to respond to—

Glimcher: Liz, definitely.

Phelps: I think you're bringing up some of the ethical issues that will emerge from this type of research that we do, and this is actually a hotly debated area in the field of neuroscience in general. Actually there's a Society for Neuroethics, which I'm on the board of, which was formed to try to address some of these questions. Things like what you call neuroeconomics, which has been exploited, I guess, by people doing marketing. There are companies that now sell neuro-marketing, and in the field of law there are companies now selling brain science to look at lie-detection and things like that. I can't speak for everybody, but for me, and I'm pretty sure for Paul, our goals as basic scientists are to understand how things work, knowing that this will then be applied, hopefully, to the treatment to things like psychological disorders, where they're actually helping people. But of course, when you have more knowledge, it can be used now in ways that perhaps society wouldn't be so happy with. This is genetics and physics and all sorts of sciences.

So you're asking what our goals are. Again, my goal is basic knowledge. I'm a basic scientist; I don't do my science with a specific application in mind. But I don't think neuroscientists are

walking away from these issues and these debates. There has been a large growth in neuroethics and an interest in neuroethics over the last five years. The Society for Neuroscience has an annual neuroethics lecture. Again, there's the Society for Neuroethics that's formed that's starting to have their own conferences. The *American Journal of Bioethics* now has a section every month on neuroethics. So these issues are emerging, and how people use this science is something, again, that we can't necessarily control, which is true of all science. You know, we think it's going to be for good, but we can't guarantee it will be for good. The best we can do is to be part of those debates, to talk to the lawyers who are now dealing with court cases where everyone has a brain excuse for their behavior, and what that really means and doesn't mean.

Audience: I've got a question that's primarily directed to Bill Brown, which has to do really with empirical data and its development, which has to do with what's really going on in our financial services firms, the investment banks, the hedge funds, the private equity funds. In terms of having both their own economists, their own mathematicians developing their algorithms, studying the actual performance of the markets, and whether or not to some extent, both the thought that's going on there, the data that they put to use, as a matter of applied economics, and presumably the continuing data that they gather over time, is sort of an unkept source of information regarding macroeconomic behavior, which might in the end have some microeconomic connection, as well.

Brown: Well, certainly there's a huge amount of effort that goes on in these financial firms, as was mentioned, there are a lot of resources and there are returns from doing this work. Economists probably have a smaller and smaller portion of this work, as opposed to other people with quantitative backgrounds that aren't necessarily in economics or finance. But the definitions or dividing lines aren't very clear. Anyway, there is certainly a lot going on. For a long time—going back twenty or twenty-five years—financial firms recognize the value of data sets, of having information on pricing, having information on a whole range of things, defaults, useful things that you can make money on, and if you have it, you have an advantage over the other people. So effort has been put into collecting data on different aspects of things.

The question of whether it's useful at a truly macro level—I would be hard-pressed to find good examples of that. I certainly know where it's been tried. When scanning things came into stores, the question is, okay, you're at Wal-Mart, you get all this scanning data, this is hugely good stuff, economists can do things, we can explain a lot at the macro—I wouldn't say that's an area that's produced a whole lot. It's probably produced less than people thought. Most of the work is just pretty straightforward. You can make money by predicting how many home loans will be repaid in each zip code in the United States. That's not a very interesting question for some, but from a money point of view, that's our biggest financial market—mortgages and pre-payment rates make a huge difference to their value and there are a lot of guys out there trying to predict pre-payment rates for each mortgage.

Glimcher: That is a notoriously irrational market—

Brown: And it's very difficult.

Audience: But in developing that kind of data, are they looking back in time to look at prior experience and therefore private behavior within the various zip codes? So at the core, isn't there really some information regarding the behavior of individuals?

Brown: Certainly, although I'm sure access to that data is of interest to people with other interests than the firm, and I'm sure that happens at some level.

Audience: Having been a foot soldier in the military and a trader—I was a currency trader—I was self-employed for a while, and I'm very keenly aware of the significance of information and access to it. I did take a medical ethics course a long time ago and I certainly appreciate how significant the issues you raise are and, as you yourself suggested, Paul, it's just the pursuit of information. You have it in its entirety and hopefully you're able to show some sunlight on all aspects of it. I think that addresses some of the ethics questions, hopefully. I was curious because decades ago, when E.O. Wilson's book first came out, before social biology was a dirty word, and they had to change it to evolutionary psychology—

Glimcher: We still use it.

Audience: Okay, well, you're in the ivory tower, so I guess you can get away with it.

Glimcher: As you know, Wilson got a pitcher of ice cold water dumped on his head in the ivory tower.

Audience: I know! I think it was right after that time they started changing the word. I saw him about a month ago, coinciding with the release of his collected works. Since you have essentially made such a pretty package here in actually talking without naming the word consilience, you know. Again, I'm not sure how many people are really prepared for that—

Glimcher: Can I tell you—the last paragraph of my book ends with the word consilience.

Audience: It does.

Phelps: It's available for purchase in the lobby.

Audience: Okay. I thought it was fascinating because his book when it first came out had devoted some 500-plus pages to what you observe about behavior in nature and clearly, nobody goes to school for that or they'd be pushed in the direction for choice-making. But the last thirty pages, of course, are incredibly controversial. Since this is an area that you're living, perhaps you could do your bit publicly here for some of the notions that he's addressing, and what is the final word in your book.

Glimcher: He's my hero, flat out. As a biologist, as a sociobiologist, I think Wilson is one of the greatest thinkers of the last hundred years—one of the five or six most important thinkers of the last hundred years. In *Consilience*, what he argues is—exactly what I believe—that there is a tremendous amount of unity between fields of knowledge and that what has happened in the last couple of hundred years is we have revealed, starting at the really easy stuff like quantum

physics, the unities that connect knowledge at different levels. And it's his conviction that we'll be able to draw parallels and mappings as we move up, even to the level of whole social systems of behavior. What he argues in the book—very eloquently, I think, but I guess not everybody drew that conclusion—is that the really big challenge is the gap between the social sciences and the natural sciences. And bridging that gap is going to be really, really hard. At my university, there is a Dean for Social Sciences and a Dean for Natural Sciences and those are two separate divisions of school. And we're working on that problem, basically—

Audience: No wonder everything is subsumed by biology.

Glimcher: No, I don't think that happens. I know you're daring me, but the reason you know that won't happen is because there is still a Department of Chemistry. A hundred and ten years ago, a bunch of chemists and physicists had this battle about whether physics was going to become chemistry. Quantum theory emerged. It unites physics and chemistry, for sure. Both of those systems rely on it, but guess what, you still have to pick whether you want to be a professor of chemistry or physics. So do I hope a hundred years from now that there will be mappings between psychological levels, economic levels, and other scientific levels? I'm convinced there will be. But I think you'll still have to pick which you want to be a professor of.

Audience: I want to react to a couple of things Dr. Nersessian touched on. One was the talk fifteen years ago in New Orleans about neuroscience and its applications in unsolved problems in psychoanalysis and what the use is in knowing about vasopressin. I remember in 1950, sitting and listening to endless lectures by the discoverer of vasopressin and oxytocin, Vincent du Vigneaud, and thinking to myself: What use is this going to be? This is impossible. I don't understand how he arrived at this, how he found it, and what use is it going to be. And from there I go to the anatomy lesson and then I didn't really quite understand what the uses were of knowing all the two hundred bones. The only use of anatomy at that point in my career was to go home and try it at home, but beyond that, it really wasn't much use, and nobody would particularly countenance using anatomy in your first year of medical school in any useful way.

But if you move forward, in tomorrow's *New York Times* Magazine section, there's an outline of a woman with very high, uncontrolled hypertension—you could read one paragraph of that and if you knew enough anatomy, you could know why this woman had uncontrolled hypertension. So it took about fifty-seven years for some of that to pay off, and what I'm trying to get at is I think these metaphors and similes and fantasies are immensely powerful, but it takes an awfully long time to develop the information, to test the information, to validate the information, to show that it has predictive value.

Now, this week there was a discovery announced that may someday, a hundred years from now, rank with some of Newton's observations. Ventner, who in many ways has a very mixed background in the genetics community, took his own yacht and went around and, every two hundred miles, took water out of the ocean and used molecular biology techniques to—not culture bacteria from that, because that doesn't work very well—but to use much more advanced methods to look at the genetics of the millions of bacteria that are in all water and to outline a whole new structure that's going to have to be validated, that's going to have to be repeated. I think the urge to have some answers right now—because we need the answers right now, the

pain is right now, the hypertension is right now—to develop this and to test it and to show that it has predictive value, that takes a very long time.

Glimcher: Well, Liz and I are still young, so.

Nersessian: If there are no other questions we will stop. Thank you very much.