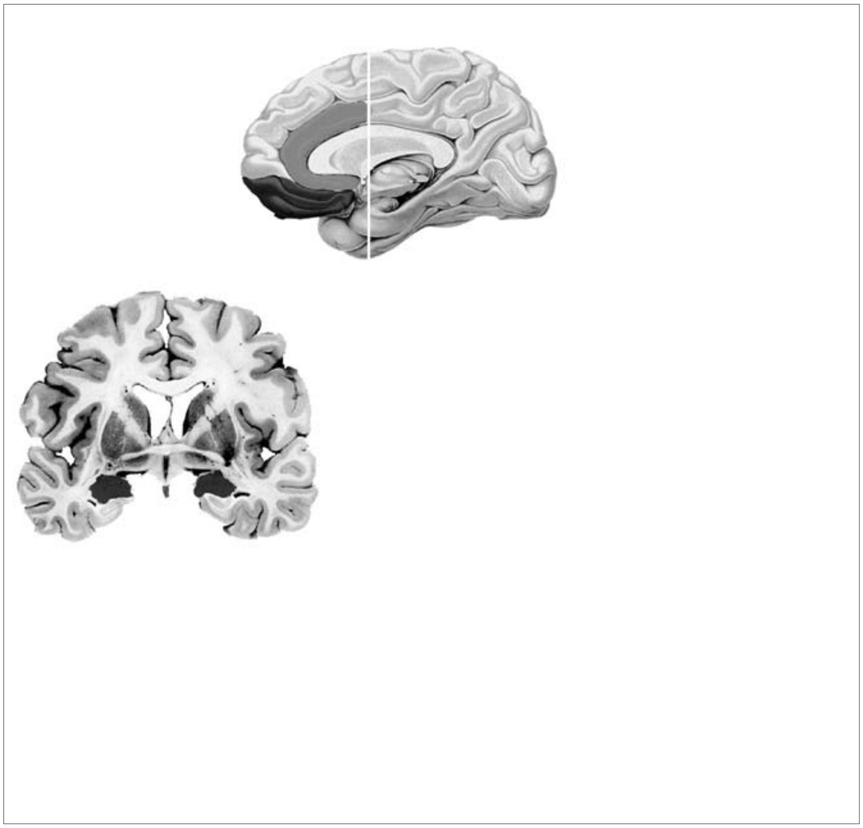

Image © Ralph Adolphs

Neuroeconomics

Colin Camerer

These remarks were given at a campus reception for Fellows and guests at the California Institute of Technology on February 15, 2005.

Colin Camerer is Rea A. and Lela G. Axline Professor of Business Economics at the California Institute of Technology. He has been a Fellow of the American Academy since 2003.

Neuroeconomics is a rapidly emerging area of study that brings together two very different fields to further understanding of decisions, game theory, and trading in markets. Until recently, economists have been content to treat the human brain as a "black box" and to express what the brain is doing in a "reduced form" mathematical equation. Lacking detail about brain mechanisms, most empirical studies of economic behavior have relied on measuring inputs, like prices, and predicting outputs, like how much people will buy. Now advances in genetics, brain imaging, and other techniques have made it possible to observe detailed processes in the brain better than ever before. The brain scanning that we carry out at the Broad Imaging Center at Caltech shows which parts of the brain are active when people make economic decisions. The results of this research will eventually enable us to replace the simple mathematical ideas that have been used in economics with more neurally detailed descriptions. The approach should also inform neuroscience by expanding the range of cognitive activities that are studied. To illustrate the kinds of experiments we undertake in neuroeconomics, I will describe several research projects conducted with my colleague Ralph Adolphs, Dan Tranel at Iowa, and two intrepid Caltech graduate students, Meghana Bhatt and Ming Hsu.

Let me begin with a few definitions. By economics, I mean precise, very stylized mathematical models of choice under scarcity. With only so much money and only so much time, how do you decide what to do? The trade-offs between goods and money are central, but more interesting are the trade-offs involving time and risk (the focus of one of our current studies, which I will talk about later). In economic theory, we assume that people act as if they can attach a number, called a "utility," to everything they might want, and they choose the goods with the highest utility number.

What gives economic models of aggregate behavior their precision is the concept of equilibrium, a word borrowed from physics. In what we call a "competitive equilibrium," prices adjust until they equalize supply and demand. In game theory, we use the term "equilibrium" in a somewhat different way, to mean accurate (or "rational") expectations. Players are in equilibrium when they have correctly guessed what others are planning to do and are making the best choices given their accurate guesses.

Neuroeconomics uses details of neural mechanisms to inform these ideas in economics of how we make choices under scarcity, and how equilibrium comes about. Neuroscientists are very opportunistic about using different tools: single neuron recording, the animal model, computational models, psychophysical measurement like skin conductance and EEGs, fMRI, and behavior of human patients with brain lesions. These tools enable you to be very precise about how brains might be computing something like a numerical utility. For example, some studies recording single neurons in monkey parietal and frontal cortex areas suggest that utilities are expressed by neural firing rates.

Game theory is a mathematical language for describing strategic interactions among players who choose strategies in order to get the outcomes they like most. Game theory has been applied to everything from biological competition among genes, to international politics where the players are nations. Despite the rapid growth of game theory as an analytical tool at many social levels, we know almost nothing about how the human brain operates when people are thinking strategically in a game. To study this we present people with a game, in the form of a matrix that shows how much two players earn if the row player picks one button representing his choice and the column player picks another button representing her choice (see Figure 1). We have lots of theories about which strategies they might pick, including the idea of equilibrium strategies that dominate most analytical game theory. Many studies show that players can learn to guess correctly what others will do, and

Neuroeconomics is a rapidly emerging area of study that brings together two very different fields to further understanding of decisions, game theory, and trading in markets.

choose equilibrium strategies themselves, but it takes learning for them to do so. So if equilibration takes trial-and-error learning, then when players think about a game for the first time they will probably *not* make equilibrium choices – that is, some players won't guess correctly what others will do. One behavioral theory is that people give up trying to guess what other players will do, and just choose a strategy that gives a high average payoff across all the other player's possible moves. This is called "one-step thinking." "Higher-order" thinkers might guess that other players will choose one-step strategies, and choose strategies that are the best responses to those. In brain circuitry, there is neural activity that may or may not correspond with these hypothesized processes.

In the normal form matrix, each player who is in the fMRI scanner picks a row, and another player, outside the scanner, picks a column. (This design also gives us a chance to see whether the experience of being in the scanner changes their behavior; it does not in our study.) In the example shown in Table 1 (one of the games our subjects actually played), the row player's strategy A is "dominated" by strategy B – that is, regardless of what the column player does, choosing B always gives a higher payoff than choosing A. If you are trying to earn money for yourself, there is no good reason to pick row A, and very few did. What will the column player do? A one-step column player will see that BB gives payoffs of either 86 and 47, while AA gives 41 and 74. The average payoff from BB is higher; so a one-step thinker will choose BB. This is not a bad choice (e.g., the one-step rule will never pick a dominated strategy like A for the row player). The one-step choice of BB is also a common one – 40 percent of the subjects chose it. But a player who chooses BB hasn't figured out that a rational row player will *rarely* choose A. So the idea that BB will pay "either 86 or 47" is wrong – the column player rarely earns the payoff of 86. In fact, if the column player "deletes" strategy A – that is, guesses the row player will never choose it – then the likely payoffs from AA and BB will be 74 and 47, respectively. AA now effectively dominates BB. Thus, in this game, the equilibrium strategies are to choose B, and to choose AA. But choosing AA requires the column player to think the row player will choose rationally.

	Row player payoff		Column player payoff	
	AA	BB	AA	BB
A	21	62	41	86
B	45	74	74	47

Table 1: A two-player "dominance-solvable" matrix game (Bhatt and Camerer, 2005).

The game in Table 1 is the simplest game we studied. Others require two or three steps of deleting dominated strategies one at a time, which requires many steps of iterated thinking. Before this study, we knew nothing about how the brain worked when making guesses about other players' guesses. To find out, the subjects actually perform three tasks for each game: The row player, for example, chooses a row, guesses what the column player will choose, and guesses what the column player will say that *she* – the row player – will choose.

Now think about how the brain might make these computations. Choosing a strategy requires looking at your own numerical payoffs, making a guess at what the other player will do (probably by looking at the other player's payoffs), making calculations of very low or high payoffs, or averages, and so on. If you are thinking strategically, and guessing the other player's choice correctly (i.e., in equilibrium), these same processes will be used to guess what the *other* player will do, by simulating their choice process. That is, if there is general choice circuitry in the brain, then when players are in equilibrium (because they are guessing correctly what others do) there should be a substantial overlap between activity during the task of choosing your own strategy and the task of guessing another player's strategy.

This is precisely what we see. Figure 1 shows areas of the brain that are *differentially* active when the row players made strategy choices, compared to when they made guesses about the column player's choice. The top "slice" of the brain shows that there is very little differential activity between choosing and guessing when they are mathematically in equilibrium (i.e., their guesses are correct); the only extra activity when they are choosing for themselves is in the ventral striatum, an all-purpose anticipated reward area (probably encoding the additional payoff in the choice task). The bottom slice of the brain shows differential activity when they are *not* in equilibrium. Here there is a lot more activity, in dorsolateral [top and side] prefrontal cortex (DLPFC) and paracingulate cortex. The fact that there is more activity suggests that when players are not guessing accurately, they are putting more thought into figuring out what to do than into figuring out what the *other* player will do. The figure shows precisely where this extra thought is occurring.

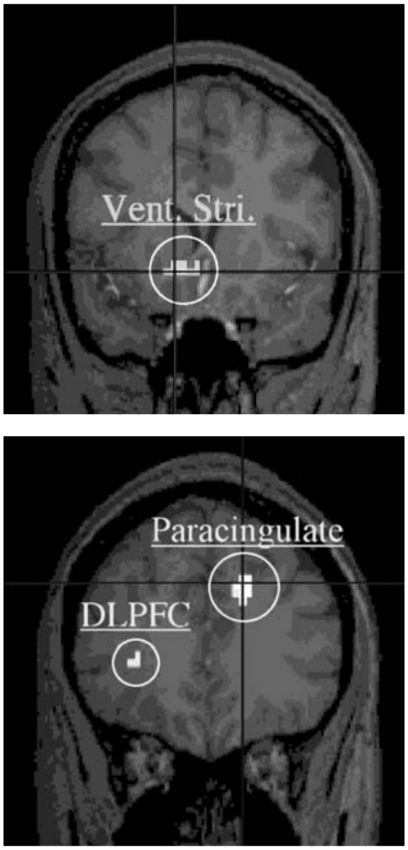

Figure 1: Equilibrium as a state of mind: differences in strategizing and guessing when game players are guessing accurately (in equilibrium, top) or guessing inaccurately (out of equilibrium, bottom) (Bhatt and Camerer, in press).

The point of this study is that when the brain is in equilibrium – which is a purely mathematical restriction on accuracy of beliefs about other players' choices – we can detect it in a pattern of neural activity. So equilibrium

Advances in genetics, brain imaging, and other techniques have made it possible to observe detailed processes in the brain better than ever before.

is not just a behavioral condition in which choices are optimal and beliefs rational; it is also a "state of mind" in a neural sense.

Let me turn to another study involving the question of trust. Large complex economies rely on trust every day. Studies indicate that trust of this sort seems to be highly correlated with economic growth. In Scandinavia, if you ask people whether "in general, people can be trusted," a vast majority say yes. Other countries, such as the Philippines and much of West Africa, are at the low end of the scale – only a small fraction say people can be trusted. Furthermore, the answer to this simple question is strongly correlated with economic growth across countries. So how trust works and how it's cultivated and understood is an important concept in the economy.

My colleagues and I have been analyzing a trust game in collaboration with Read Montague and other researchers at Baylor Medical School. A first player, who we call the investor, starts out with twenty currency units that are converted to actual dollars at the end of the game. (We always pay people actual money because it focuses their attention, and we often use very large sums of money to be sure they are seriously motivated.) Let's say the player invests fourteen units and keeps six. Whatever he invests triples. In this case, fourteen become forty-two, representing the return on a productive investment. The tripled amount rests in the hands of the second player, the trustee, who decides how much to keep and how much to give back. He can keep it all if he wants, so it is like investing in a foreign country with no legal protection against contractual breach.

The amount the first person invests is a measure of how much he expects the trustee to repay. The amount the trustee repays is a measure of trustworthiness. If the trustee repays less than fourteen out of the forty-two that was created by the investment, then the investor's trust did not quite pay. If the two players trusted one another, the original twenty would have become sixty. But if the first player thinks the second player is selfish, there's no reason to trust him.

We've been scanning the brains of two individuals: one at the Broad Imaging Center at Caltech, and the other at the Baylor Medical Center in Houston. They are actually playing with one another through an Internet connection and having their brains scanned as they play. This is the first time in fMRI scanning that anyone has ever taken two brain activity patterns and tried to correlate across them. The fact that the two brain activities could be correlated is not that surprising. For example, as I'm talking, you're listening. Language areas of our brains are both active, so naturally there would be some interesting correlations. However, what we found in our study is that something is going on in the two brains that is distinct from what is going on in one brain at a time. A kind of social brain pattern has occurred. The two brains are generating activity simultaneously in two different regions: a conflict resolution area called the cingulate and an expected reward area called the caudate (also seen in Figure 1). This correlation indicates that the two players are simultaneously thinking about what to do, and what it will pay.

Another interesting question to ask concerns how trust spreads. If one CEO does something terrible, do people think all CEOs do something terrible? Do they associate lack of trust with CEOs in a given state, or with a given skin color, or with an MBA from a particular school? How does trust generalize across social categories? Understanding exactly how this process works is very important for measuring and restoring trust. We know almost nothing about it, but can learn from behavioral experiments and imaging the brain.

Turning to another study, we've been interested in "known unknowns," or what decision theorists call "ambiguity." These are simply choices in which people can take a sure amount of money, or can make a bet on an actual event to win more than the sure amount. A typical event was whether the high temperature in New York on November 7, 2003, was above 50 degrees Fahrenheit. We picked 50 degrees because it is close to the average New York temperature on that day, and thus the players would likely have a fifty-fifty chance of winning the bet. Since the Caltech students in the study know something about temperatures in New York, this was a case of betting with the benefit of a lot of available information, just as an insurance company does when it uses a large sample. In another gamble – the known unknown – the player bets on the temperature in Dushanbe, Tajikistan, which they usually know very little about. Suppose people are often reluctant to bet that Tajikistan was warm on a particular day. In standard decision theory, events have crisp probabilities associated with them, so if you won't bet that Tajikistan was warm, you should believe it was cold and be willing to bet that it was cold. But interestingly, many people are unwilling to bet on *either side* of a low-knowledge ambiguous event. Our interest was the special brain activity occurring when players are evaluating these low-knowledge gambles. Our studies showed that when probabilities are ambiguous, there is additional activity in the amygdala. The amygdala is an area that is important in emotional learning and in expressing "vigilance" in the face of fear or discomfort. It is like a watchdog in the brain that responds rapidly, but rather stupidly, by barking whenever there is a threat. Seeing fearful faces rapidly activates the amygdala; when the amygdala is damaged, a person loses the ability to detect fear in the faces of others.

Despite the rapid growth of game theory as an analytical tool at many social levels, we know almost nothing about how the human brain operates when people are thinking strategically in a game.

In our study, when the player is betting on New York, there's not much fear; the player is just trying to weigh the odds. But when the player is betting on Tajikistan, the amygdala signal warns, "Be careful betting; you don't know anything about it." We see this as a neural way of resolving a longstanding debate in decision theory about the importance and source of "fear of the economic unknown."

Many decision-theorists have argued that you should talk yourself into not worrying about ambiguity. You should say, It's either

When the brain is in equilibrium – which is a purely mathematical restriction on accuracy of beliefs about other players' choices – we can detect it in a pattern of neural activity.

warm in Tajikistan or it's not warm. These are two separate events. If you don't know anything about temperatures there, you should treat it like a coin flip, and you should be just as willing to bet that the temperature is high as you would bet on a coin flip. The amygdala says be careful, we might be making a mistake. Talking yourself into treating these situations like coin flips requires a triumph of logic, probably in the frontal cortex, over the highly evolved "be careful" position that enabled organisms to survive danger for millions of years. Ironically, we identified a group of individuals who were immune toward the fear generated by ambiguity, as decision theory prescribes. They are not brilliant decision-theorists; they are but people with damage to the orbitofrontal cortex (just above the eye sockets). The amygdala projects neurally to the orbitofrontal cortex. These brain-damaged people do not receive the normal biological signals that transmit fear of the economic unknown, so they treat bets on Tajikistan like bets on New York. The fact that they behave "rationally" calls into question whether rationality should be defined as adherence to logical axioms (the traditional approach in economics) or as biological adaptation. More broadly, asking and answering questions like these has the potential to link biological and social sciences, which is the great promise of synthetic areas like neuroeconomics. ■

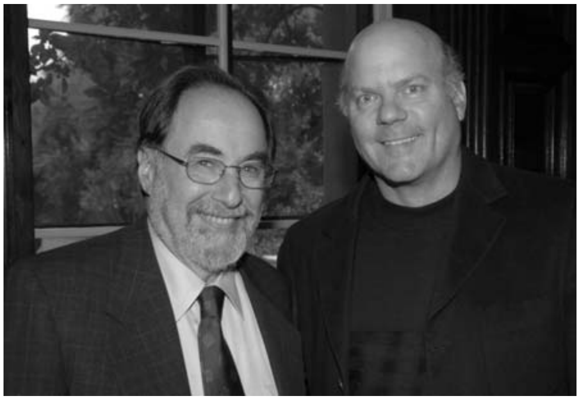

David Baltimore and Colin Camerer (California Institute of Technology)

Louis W. Cabot (Cabot-Wellington, LLC), Margo Marshak (California Institute of Technology), and Colin Camerer

John O. Ledyard, John D. Roberts, Herbert Bishop Keller, Rudolph Arthur Marcus, Paul C. Jennings, and John D. Baldeschwieler (California Institute of Technology)