The Neural Basis for Self-Control

You're on a diet, but you really want a piece of that chocolate cake. What's going on in your brain as you struggle to resist temptation?

***Antonio Rangel* is an associate professor of economics. After graduating from Caltech with his BS in economics in 1993, he received his PhD from Harvard in 1997. He then became an assistant professor of economics at Stanford before returning to Caltech in 2006. He is the president of the Society for Neuroeconomics. The following, based on his talk at the 72nd Seminar Day in May 2009, was funded by the Moore Foundation and the Economic Research Service of the U.S. Department of Agriculture on Behavioral Health Economics Research on Dietary Choice and Obesity. This article was edited by Marcus Y. Woo.**

Many of the world's problems are the result of faulty decision making. If we could understand how the brain makes decisions, then maybe we could make better choices. I'm a professor of economics and part of a team of psychologists and neuroscientists

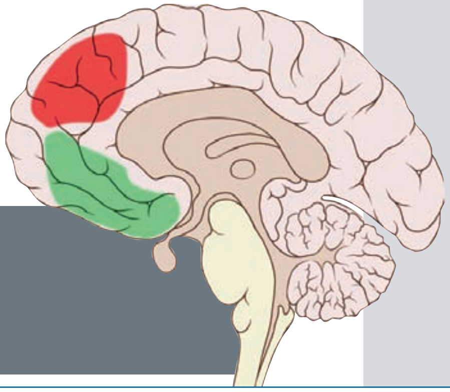

Researchers are finding that the ventromedial prefrontal cortex, or vmPFC, (green) is where the brain encodes how much you're inclined to make a particular choice. The dorsolateral prefrontal cortex, or DLPFC, (red) has to come online for the brain to exercise self-control.

By Antonio Rangel

at Caltech tackling the problem of how the brain makes decisions. In particular, I'm interested in self-control, which is at the core of many of the most pertinent public-policy and health issues in the United States. Think about addiction. Think about obesity, which is a very personal problem for me—I like sweets a little bit too much. Or think about the low savings rates of the United States.

All of these problems are the result of poor decision making, and have inspired us to ask three questions. First, why does the brain have a problem with self-control at all? Why has evolution developed a machine that becomes conflicted while trying to make good decisions—one of its most important tasks? Your visual system, for example, doesn't give you conflicting outputs. If I look at you, I see a clear, sharp image. Why doesn't your decision-making circuitry act like that? The second question is, how does the brain actually exercise self-control? The final question is the most important one: what is different between the brains of people who can and cannot exercise self-control? Can we make the second group more like the first?

We have been doing a lot of experiments to see what's going on in the brain while it's trying to exercise restraint. We primarily study the self-control involved in dieting. We're interested in obesity, and, more importantly, dieting, which is a simple paradigm that allows us to control a lot of variables.

Over the last two years, we've discovered that two important mechanisms in the brain kick in when we make decisions—such as choosing to eat an apple instead of a piece of cake. The first involves an area of the brain just a little above and behind the eyebrows called the ventromedial prefrontal cortex (vmPFC). When you are presented with a choice of objects—say a cookie or a piece of fruit—this region encodes a value for each object in the firing rates of neurons. The more frequently the neurons fire, the higher the value, and the more likely you'll end up choosing that object. It seems that this value signal determines your choice—regardless of whether you are inherently good at self-control.

The second mechanism we discovered determines the difference between a good and a bad self-controller. In people who have trouble with willpower, the signal seems to reflect only the immediate and effective value of things. When they see a cookie and a piece of fruit, they think, "Cookie: great taste. Health: who cares?" Health doesn't get incorporated into the value signal. A good self-controller, in contrast, also activates another area of the brain called the dorsolateral prefrontal cortex (DLPFC). This region modulates the values generated by the vmPFC so that they include long-term considerations, like keeping your weight down, guiding you to make more sensible choices. The difference between a good and a bad self-controller is not that the bad one likes cookies more. It's that this other area of the brain—the DLPFC—does not come online to effectively modulate activity in the vmPFC.

DECISIONS, DECISIONS . . .

We started with a very simple—yet useful—conceptual framework. Suppose you have to choose between several items. What does the brain need to compute? The brain has to first identify the items and then assign a value to each one. The brain then compares those values to make a choice. But, of course, you're not sitting there saying, "Hmm, the value of the cookie is five, the apple is a three. Therefore, I should go for the cookie." This process happens in as little as 300 milliseconds and can be totally unconscious.

In an experiment, we asked 19 test subjects to decide how much certain kinds of food are worth to them. While they made these decisions, we examined their brains with a functional magnetic resonance imaging (fMRI) machine, which helps map brain activity. The device is similar to the MRI scanner that you may have been in, unfortunately, to examine a torn ACL. Active areas of the brain consume more oxygen, and it turns out that the amount of oxygenated blood that comes into an area of the brain is proportional to neural activity there. Oxygenated blood and deoxygenated blood have different magnetic properties. The fMRI machine creates a three-tesla static magnetic field—a very, very strong and very stable magnetic field that allows us to detect tiny, localized changes in the brain's magnetic field and see where the value decisions are encoded. While the subject lay in the scanner, a screen displayed pictures of different snacks—junk food, like candy bars and

Listen to a podcast of Rangel discussing his research on the brain and decision making.

The difference between the self-controllers and non-self-controllers were striking. What was going on in their brains that was setting them apart?

chips—for four seconds. We had given the subjects money to buy these items, and they had to type in a bid of $0, $1, $2, or $3.

These weren't just hypothetical situations. The decisions were real. Before coming to the lab, the subjects were told to fast for four hours, and that they would have to stay in the lab for half an hour after the experiment. They were hungry, and the only thing they could eat was whatever they had bought from us. So that they didn't have to worry about budgeting their money, only one of their transactions—one we randomly selected—was implemented.

If I were a vision neuroscientist instead of a behavioral neuroscientist, the experiment would be very simple. I would show you a stimulus and vary an objective, easily quantifiable characteristic, such as color or contrast. Then I would see if there's an area of the brain where the activity is changing at a rate proportional to the change in that characteristic. But the difficulty in behavioral neuroscience is that it deals in subjective, hard-to-measure values. I need to use a trick—and this is where experimental economics meets neuroscience. In our experiment, we didn't sell the food at the price the subject offered, but used a method called a Becker-DeGroot-Marschak auction. In this procedure, you tell me that something is worth X dollars to you. I pick a random number, and if that number is less than or equal to X, you get the item and you pay the number that was chosen randomly. But if the random number was bigger than your bid, you keep your money and don't get anything. If you think about it, you can see that the optimal strategy is to bid the true value of the item—not a penny more, not a penny less—which allows us to get an objective measure.

We can measure changes every two seconds in brain activity within regions as small as a cube one to three millimeters per side. Even in such a small box, called a voxel, there are hundreds of thousands of neurons. But, as you'll see, these measurements are still very useful. We want to know if any of those boxes show a statistically significant response—above or below the normal background activity level—that's proportional to the value of the item. If so, then we can conclude that those voxels are probably involved in computing the value signal. We found that activity in the vmPFC correlates with that value.

We also ran an experiment in which the subjects had to choose between things they didn't like, such as canned vegetables and baby food. The experiment was the same as the one I described above, except that now the subjects had to pay *not* to eat the food. The bidding rules were the same, but now we're measuring how repellant something is, rather than how appetizing. At the end, we randomly chose one of the trials to implement, and unless the subjects won the auction, they had to eat the food, which wasn't the tastiest of treats. At one point we had pig's feet, but it was just too much for American students to swallow.

We asked a similar question with this experiment: were there any areas of the brain that encoded how disliked something was? In other words, did neuron firing rates increase proportionally with aversion? The answer, we found, was no—not even if you squint at the data. However, we found that there are areas of the brain in which activity *decreases* proportionally to how aversive something is. Furthermore, the relevant regions were the same as those of the previous experiment. If you put the brain scans from both experiments on top of each other—remember, each experiment had different foods, and was conducted on different days—the active areas overlap, suggesting that this region can encode a positive or negative value signal. For the appetizing case, the brain ramps up activity; in

Left: How much would you pay *not* to eat baby food (carrots, in this picture)? Top right: Do you have the self-control to resist a chocolate peanut-butter cup?

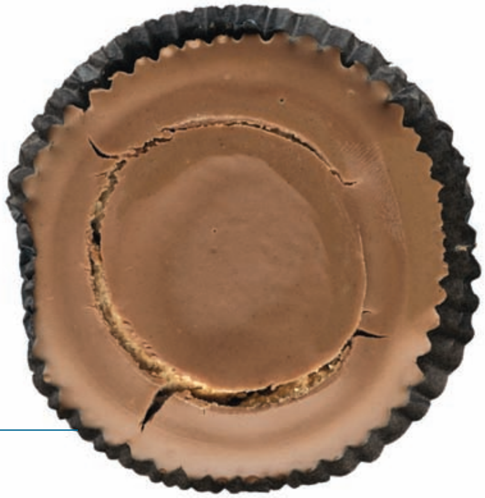

the aversive case, it ramps down activity.

But, you might say, maybe food is a special case—what about other things? So we did another version of the experiment in which we asked people to bid for an 80 percent chance to get food, Caltech paraphernalia, and cash. For example, subjects were asked how much they'd be willing to bid for an 80 percent chance at receiving $3 or a Caltech hat. This experiment showed that the items' values were all encoded in the same brain region, regardless of what type of item they were. There are also numerous other, more technical experiments that identify the vmPFC as the place where the brain encodes value. For example, Camillo Padoa-Schioppa and John Assad at Harvard measured the electrical activity in single neurons in the vmPFC of monkeys while the animals made choices between different types of juice. The researchers found that activity in these neurons encoded value signals, which is consistent with our findings.

The stimuli in this experiment were healthy items, like fruits, and unhealthy ones, like candy. Every subject had to give each item a health rating, independent of taste, on an integer scale of -2 to 2. Then they were asked to rate everything on taste, independent of any health considerations, on the same scale. We could then select a health- and taste-neutral item for each person, and once we had those reference items, we asked the subject to choose between a new food item and the neutral item, which remained constant with each trial. This way, we knew that it was the new item that drove whatever signals we detected. At the end of the experiment, we again randomly implemented one of their decisions, and they had to eat whatever they had chosen.

Based on the data, we then divided the subjects into self-controllers and non-self-controllers. Certain trials required self-control—if I were to show you an item that is tasty but unhealthy, you should say "no." But if I were to show you something that is tasty *and* healthy, you don't need to exercise any self-control. We categorized someone as a self-controller if he or she exercised discipline in at least half of the trials that required it. A typical self-controller is more likely to choose the healthier item; they're not really responsive to taste. But a non-self-controller—that's me—cares only about taste.

When no self-control was required, both groups behaved the same way. People always chose healthy things that they liked. In contrast, we could not get people to go for food that was disliked and healthy; no one was clamoring for broccoli. Still, the differences between the self-controllers and non-self-controllers were striking. What was going on in their brains that was setting them apart?

AN OFFER YOU CAN'T REFUSE?

This brings us to the DLPFC, which needs to come online for the brain to consider the long-term benefits of a particular choice. To study this, postdoc Todd Hare, Kirby Professor of Behavioral Economics Colin Camerer, and I recruited dieters and nondieters by offering them $100. As before, the subjects were asked to fast so that they would be hungry when they came to the lab. We then asked them to make several food choices while in the scanner.

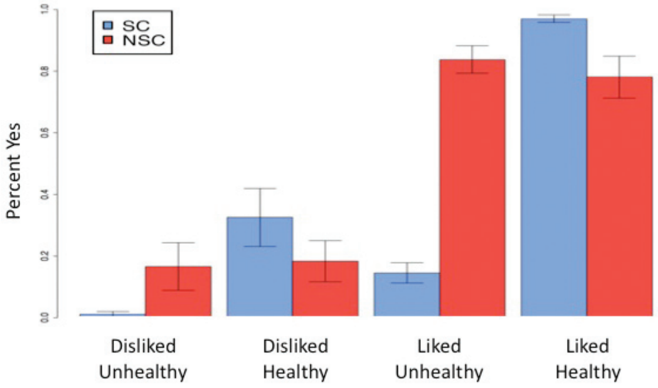

Percentage of self-controllers (blue) and non-self-controllers (red) who chose different combinations of healthy and liked food items. Those who were good at self-control were more likely to choose the healthier option, regardless of how much they liked it. This and the next two plots are from *Science*, Vol. 324, pp. 646-648, May 1, 2009. Reprinted with permission from AAAS.

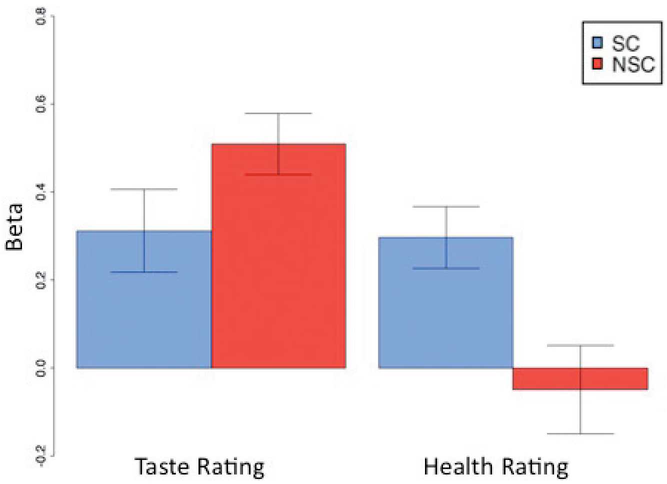

In self-controllers (blue), both the taste and health ratings of food influenced the activity of the vmPFC, as measured by the parameter beta. For non-self-controllers, only taste was important.

THE SELF-CONTROLLING BRAIN

We looked at whether taste or health ratings drove the value signals in the vmPFC for the two groups. In self-controllers, both health and taste ratings influenced vmPFC activity. In non-self-controllers, the vmPFC was only affected by taste. There was a large correlation between how much the subjects' health ratings affected their decision and how much the ratings affected vmPFC activity.

When the brain exercised self-control, the DLPFC was more active among self-controllers than non-self-controllers. Furthermore, the DLPFC is active whenever the brain successfully practices self-discipline, regardless of whether you're good at self-control. These results by themselves don't prove the powerful idea that the DLPFC modulates the vmPFC to drive choices. However, we then found more evidence supporting this idea.

Sitting a bit farther down the prefrontal cortex, behind the face, is a region called the IFG/BA46. This area appears to be an intermediary between the DLPFC and the vmPFC. For non-self-controllers, the DLPFC doesn't come online to modulate the vmPFC. But when a self-controller is confronted with, say, a piece of chocolate cake, the DLPFC comes online and inhibits the IFG/BA46. Lower activity in this intermediate region then lowers the activity of the vmPFC, and the subject declines the cake. So all of the pieces fit together just right.

The data suggests that a lack of self-control is the result of deficient DLPFC function. This is interesting for a couple of reasons. The DLPFC is involved in a host of other behaviors, such as emotional regulation—calming yourself down when something upsets you. This brain region also plays a role in cognitive control, which enables you to override certain ingrained responses. For instance, you might be accustomed to reaching for the light switch on the right when you enter your living room. But you don't want to reach to the right every time you enter a room. Furthermore, it's been shown that high IQ and behavioral self-control are highly correlated, and the DLPFC seems to be very important when people solve puzzles involved in IQ tests. Thus, we need to explore whether this region is involved in a series of regulatory mechanisms that are necessary for good emotional regulation, cognitive control, self-control, and cognition in general—that is, to be able to retrieve and use information.

Why is it that I can stay up until 3:00 in the morning working like an animal—I mean, I love my work, but it does require some self-control—but I just cannot say no to Pie 'n Burger [a popular local eatery that specializes in, well, pies and burgers]? We don't know, and the reason we don't know is that there is something deeper behind this. Self-control occurs when the DLPFC comes online and modulates the vmPFC, but what makes the DLPFC come online in the first place? And does it only come online in specific circumstances? Hopefully, we'll soon be able to answer those questions, and, within one or two decades, we may be able to stimulate the DLPFC directly to make sure that it gets deployed during important decision-making situations.

PICTURE CREDITS

36 — Bob Paz; 37 — Patrick J. Lynch ; 38-39 — Doug Cummings

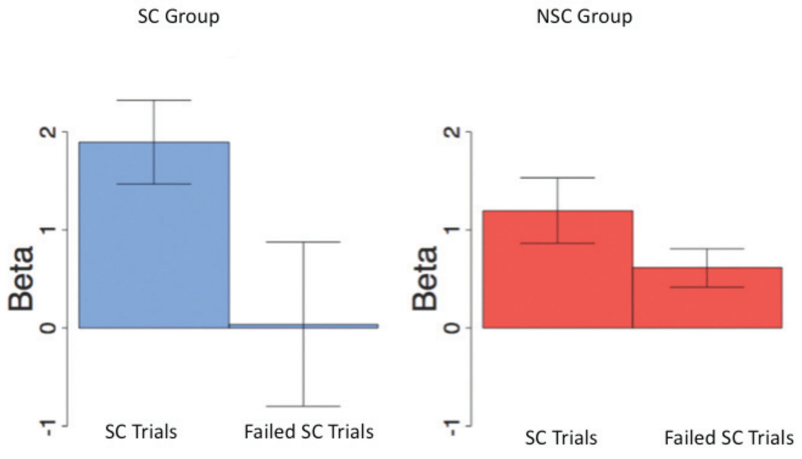

Both self-controllers (blue) and non-self-controllers (red) showed more activity in the DLPFC when they were successful at practicing self-discipline.