

Cognitive-Affective Interactions in Human Decision-Making: A Neurocomputational Approach

Hongbin Wang (Hongbin.Wang@uth.tmc.edu)

Chance Coble (Chance.M.Coble@uth.tmc.edu)

School of Health Information Sciences, University of Texas Health Science Center
7000 Fannin Suite 600, Houston, TX 77030, USA

Paul Bello (Paul.Bello@rl.af.mil)

Air Force Research Laboratory, Information Directorate
525 Brooks Rd, Rome, NY 13441, USA

Abstract

To develop realistic computational cognitive models of human decision making, it is essential to incorporate various traditionally non-cognitive but critically important soft factors into the model development, such as cultural biases, intuition, emotion, social norms, etc. In this paper we attempt to explore human cognitive-affective interactions in strategic thinking through a neurocomputational modeling approach. To empirically justify the model, we carried out a neuroimaging experiment using the ultimatum game.

Introduction

In the Art of War Sun Tzu tells us, if the enemy is too strong, flee from him. In general, why do we run away if we notice that we are in danger? Is it because we are afraid? Or is it because we decide that it is the logical thing to do based on our reasoning and strategic thinking? Or is it because of both?

While good decision-making is among the prime examples of human intelligent behavior, how people do that has been a long-standing open question. The dominant theoretical framework (e.g., Subjective Expected Utility Theory) proposes that people always weigh the costs and benefits, or the feasibility and desirability, of actions so as to choose the action that can maximize the overall utility. While theoretically and formally appealing, psychological evidence typically shows that people systematically deviate from the prescription of this normative theory. As noted by Viner long time ago (1925), many factors may contribute to human performance besides utility maximization:

Human behavior, in general, ... ,not under the constant and detailed guidance of careful and accurate hedonic calculations, but is the product of an unstable and irrational complex of reflex actions, impulses, instincts, habits, customs, fashions and hysteria. [p.373-374]

Despite these observations, the relation between reason and emotion has been a tangled one and how their interaction leads to decision making and strategic thinking remains elusive. From antiquity onward, philosophers have

often opposed reason and emotion (e.g., "Reason wishes the decision that it gives to be just; anger wishes to have the decision which it has given seem the just decision", Seneca, On Anger). According to Seneca, anger is typically too shortsighted to engage in strategic interaction with reason. In his famous book, The Theory of Moral Sentiments, Adam Smith (1759) argued that human behavior was determined by a struggle between the "passions" and the "impartial spectator." The passions included basic instincts such as hunger, pain, sex, and emotions such as fear and anger. According to Smith, while behavior was often under the direct control of the passions, people could override passions' control through the impartial spectator – an outsider and a “moral hector who, looking over the shoulder of the economic man, scrutinizes every move he makes” (see Ashraf, Camerer, & Loewenstein, forthcoming). The outsider in Smith's theory is, of course, only imaginary. But the struggle between reason and emotions is true and ubiquitous in human everyday decision-making.

Efforts have been taken in recent years to augment traditional computational models with so-called behavior moderator functions (e.g., Biddle, Henninger, Franceschini, & Jones, 2003). The dominant approach is to compute emotion through an artificial appraisal function, which is often heuristics-based or rule-based (e.g., Gratch & Marsella, 2004; Marinier III & Laird, 2004). Such an appraisal can then influence cognition through new rules or parameter tuning.

The complex nature of emotions, however, makes any attempts to compute emotions difficult. Elster (Elster, 1999, 2004) suggests that emotion be described from the following six perspectives:

- Physiological arousal: the departure from the physiological baseline;
- Physiological expressions: making emotions felt by one person known to others;
- Valence: the pain and pleasure that accompany emotions;
- Cognitive antecedent: emotions are triggered by beliefs;
- Intentional objects: emotions are about something, or directed toward something.
- Action tendency: emotions lead to action. For example, anger causes the object of anger to suffer (revenge),

hatred causes the object of hatred to cease to exist, fear leads to flight or fight, and love leads to approach and touch the other.

Among the six aspects, valence has the most obvious relevance to reasoning and decision making – the pain and pleasure of emotions are simply negative and positive utilities that may contribute to the overall expected utility of a choice. Valence reveals the hedonic property of emotions. Elster suggests that the level of pain and pleasure derived from emotions are inversely related to the probability of the events that generate these emotions. However, according to Elster, emotion is the sum of all the six aspects and leaving any one out can be misleading.

The Emotional Brain

Recent advances in cognitive neuroscience help clarify the neural underpinnings of emotions and their relationship with other higher-level cognitive functions (Damasio, 1995; LeDoux, 1996, 2000). The fact that the thinking brain (e.g., the neocortex) evolved from the emotional brain (e.g., the limbic system) suggests that there was an emotional brain long before there was a rational one and has important implications on cognitive-affective interaction. On the one hand, the neocortex apparently allows for the subtlety and complexity of emotional life. There is more neocortex-to-limbic system in primates than in other species, explaining why humans are able to display a far greater range of reactions to our emotions and suggesting higher-level thinking might somehow be able to govern emotional responses (LeDoux, 1996). On the other hand, as the root from which the newer brain grew, the emotional areas are enervated to all parts of the neocortex, giving emotional centers immense power to influence the functioning of the rest of the brain. Rite Carter (1998) put it this way in his book *Mapping The Mind*:

At the center of the brain lies a cluster of strange-shaped modules that together are known as the limbic system. This is the powerhouse of the brain — generator of the appetites, urges, emotions and moods that drive our behavior. Our conscious thoughts are mere moderators of the biologically necessary forces that emerge from this unconscious underworld; where thought conflicts with emotion, the latter is designed by the neural circuitry in our brains to win [p.54]

Interestingly, more recent studies using functional neuroimaging have largely conformed to these evolutionary hypotheses though more details have also been revealed. Various structures, such as the amygdala, substantia nigra (SN), anterior cingulate cortex (ACC), and the prefrontal cortex (PFC), have all been found active in processing emotionally competent stimuli and/or the execution of emotions (LeDoux, 1996; Rolls, 1999). It has been found that these systems overlap dramatically with the areas that are modulated by the mesencephalic dopamine system, which has long been believed to represent a “hedonia”

center in the brain and contribute to reward processing, reinforcement learning and decision making (Holroyd & Coles, 2002; Liang & Wang, 2003; O'Reilly & Munakata, 2000; Schultz, 2002; Shizgal & Arvanitogiannis, 2003; Sutton & Barto, 1998). The mesencephalic dopamine system is composed of a collection of nuclei including SN and the ventral tegmental area (VTA). These nuclei project widely to the basal ganglia and the ACC and PFC. How the valence function is computed through these brain areas using biologically plausible mechanisms may hold the key to understand how emotions affect decision making.

The Ultimatum Game

The ultimatum game (Camerer, 2003; Davis & Holt, 1993; Thaler, 1988) is an ideal task for studying cognitive-affective interaction in decision making in the sense that it simultaneously involves rational reasoning, emotional responses, and strategic thinking. As a step toward systematically understanding the cognitive-affective interactions in human decision making, we have examined the brain activities and computational algorithms underlying human decision making using the ultimatum game

The rules of the game are quite simple. Two players have to agree on how to split a sum of money. The proposer makes an offer. If the responder accepts, the deal goes ahead. If the responder rejects, neither player gets anything. In both cases, the game is over. Obviously, rational responders should accept even the smallest positive offer, since the alternative is getting nothing. Proposers, therefore, should be able to claim almost the entire sum. In a large number of human studies, however, conducted with different incentives in different countries, the majority of proposers offer 40 to 50% of the total sum, and about half of all responders reject offers below 20% (Camerer, 2003).

One dominant explanation for the responder's typical rejection of offers below 30% of the total sum resorts to the people's willingness to pursue fairness (Nowak, Page, & Sigmund, 2000). According to this view, unfair offers often invoke various psychological and emotional responses, which can influence or even overwrite rational decision making – yet another manifestation of the well-documented and long-standing opposition between reasons and passions. Recent neuroimaging evidence in the emerging field of neuroeconomics supports the explanation (Glimcher & Rustichini, 2004; Sanfey, Rilling, Aronson, Nystrom, & Cohen, 2003). It has been found that unfair offers activate specific brain regions including the anterior insula, which are often associated with disgust (Sanfey, Rilling, Aronson, Nystrom, & Cohen, 2003).

While intuitively appealing, why and how this is so remains unclear. Our neurocomputational study has the potential to expose explicitly the representations and computational mechanisms underlying the decision making process and offer new insights that go beyond the descriptive opposition of reasons and emotions.

The Experiment

The original ultimatum game is 2-player gain-based game. That is, 2 players are asked to split an amount of money. In the current experiment, in order to explicitly examine the emotional component in human decision-making, several key independent variables are manipulated.

The first one is the framing factor. In addition to the gain framing, we add a loss framing. That is, the players are asked to share a monetary cost or penalty. The proposer proposed an amount for the responder to share. Again, the responder can either accept or reject. If she/he accepts, deal moved forward. If she/he rejects, everybody in the group has to pay the whole amount (e.g., \$10).

The second variable is the number of players, which can be 2 or 3 or 5. The purpose of this manipulation is to change the rational expectation of the responder. For example, if the total amount is \$10 and the number of players is 5, the rational expectation would be \$2.

The third variable is the amount of proposal. In the gain framing condition, it would be the offer the responder would get if he/she accepts. In the loss framing condition, it would be the share of cost the responder would have to pay if he/she accepts. The amount can be \$0, \$1, \$2, \$3, ..., \$10.

A within-subject 2(gain or loss) x 3 (# of players) x 11 (proposal \$) design is adopted. The order of gain or loss condition is balanced among subjects. In each framing task, subjects perform each trial condition twice, resulting in a total of 66 trials. They are presented in a completely random order. Subjects respond using a mouse, clicking left button for acceptance and right button for rejection. The EEG data is collected using the 128-channel EGI system with Net Station. Six subjects were paid to participate in the experiment.

The trial layout and various time parameters are shown in Figure 1.

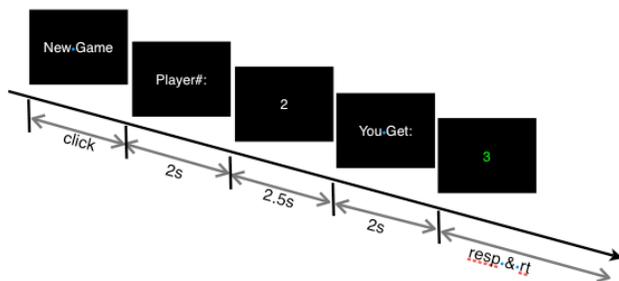


Figure 1: Trial layout

Behavioral Results

The reaction time data show that subjects make responses fairly quickly (~1s), though there is a difference between the gain and loss conditions. The response data are analyzed separately for the gain and loss frames.

Gain Table 1 (left) shows the average response rejection rate in each condition (a combination of group size and offer amount). To illustrate, the cells are color coded, where the

fair offers are coded blue, unfair offers are coded in other colors. In addition, we distinguish two types of unfair offers. Sometimes the offers are less than the fair offer, which, if accepted, were unfair and unfavorable to the responders. These conditions are coded in red. In other cases the offers were more than the fair offer, which, if accepted, are favorable to the responders but unfair to the other group members. These conditions are coded in green. It is clear that where subjects typically reject those unfavorable unfair offers, they accept those favorable unfair offers.

Table 1: Rejection rate in the gain (left) and loss (right) framing

	2	3	5		2	3	5
0	1.00	0.92	1.00	0	0.08	0.00	0.00
1	1.00	1.00	0.83	1	0.08	0.00	0.00
2	0.83	0.67	0.17	2	0.00	0.00	0.00
3	0.67	0.42	0.08	3	0.00	0.08	0.50
4	0.67	0.08	0.00	4	0.08	0.42	0.58
5	0.00	0.00	0.00	5	0.17	0.58	0.67
6	0.00	0.00	0.00	6	0.50	0.83	0.92
7	0.00	0.00	0.00	7	0.83	1.00	1.00
8	0.00	0.00	0.17	8	0.92	1.00	1.00
9	0.00	0.08	0.08	9	1.00	1.00	1.00
10	0.17	0.08	0.17	10	1.00	1.00	1.00

There are 3 important issues that need to be noted here. First, while the typical 30% rejection rule found in the literature holds for 2-player condition, it does not hold for those more than 2-player conditions. It seems that subjects simply tend to reject unfair and unfavorable offers in general. Second, if the above point can be explained as the subjects' general tendency to pursue fairness, as has been suggested in the literature, subjects' tendency in accepting those offers that are favorable to themselves but unfair to their peers rebuts this explanation. In general, subjects tended to accept those offers. Third, it is interesting to note that when the offer is \$9 or \$10, there are some trials in which subjects reject the offer. It seems that subjects might have realized that such offers were too unfair to other group members and feel morally guilty if accepted. This is confirmed by subjects' verbal report after the experiment. As a matter of fact, it turns out that only one subject did it. Except this single subject, everybody else accepted such quite "generous" offers.

Loss The loss framing shows almost exactly opposite pattern in terms of the rejection rate (see Table 1, right). Again, the general pattern is that subjects tend to reject unfavorable and unfair offers but accept favorable and

unfair offers. The subject who showed higher moral standard and made a difference in the gain framing condition did not make a difference here. Later debriefing confirmed that she did not feel obligated to share some cost when offered to pay \$0.

ERP Results

To probe the brain activities and dynamics underlying subjects' decision making, we collect the EEG data while subjects are doing the experiment so that the ERP (event-related potential) analysis can be conducted (Luck, 2005; Picton et al., 2000; Rugg & Coles, 1995). An ERP is a pattern of electrical activity on the scalp, generated by the brain, that occurs in response to a cognitive event during an experimental condition. Again, data are analyzed separately for gain and loss frames, and only the gain results are reported here.

Accept vs Reject (Gain) As we described above, one of the most important contrasts we are interested in is based on subjects' responses. We would like to know the neural differences underlying when subjects accept the offer vs when they reject the offer. From the behavioral data, we know that subjects typically reject the unfair and unfavorable offers while accept those fair offers and those unfair but favorable offers. Note that based on the standard game theory, subjects should accept any offers that are larger than \$0. So those rejecting trials might represent a case in which emotional factors have played a role. This has been confirmed by a recent fMRI study (Sanfey et al., 2003). However, due to the low temporal resolution of fMRI technique, how and when these factors work in the decision making process is unclear. We would like to explore this issue by comparing ERPs of the accept trials and reject trials.

To do so, we have made two manipulations. First, we choose not to use those trials with \$0 offers. The reason is that it makes a rational sense to reject those offers anyway. The result of this manipulation is that we have only 60 trials for each subject. Second, we exclude one subject's EEG data in our analysis due to the fact that too many eye blinks are detected in his EEG data.

Based on the 5 subjects' data, we successfully derive the ERP results, which shows that in the first 150ms after the onset of the offer, ERPs were generally negative. More positive ERPs appeared after that, starting in the frontal part of the brain, and generally propagated backward and reached the parietal region fairly quickly (by about 200ms). These positive ERPs stayed on afterward in the whole brain regions, only they were getting more and more positive.

Given this general pattern, however, it is easy to see the difference between the two response conditions. The most noticeable difference occurs in the period of 150 to 270 ms after the offer onset. In the accept condition, positive ERPs gradually extend backward to the parietal region and get more positive during this period. However, this is not the case in the reject condition, where the parietal region ERPs

stay negative during this period. Only later, at about 270 ms after the offer onset, the reject condition gradually catches up the accept condition in that more positive ERPs appear in the posterior part of the brain.

A large body of evidence has shown that the parietal lobe is closely associated with numerical concept, mathematical thinking, and utility calculation (Dehaene, Spelke, Pinel, Stanescu, & Tsivkin, 1999; Hubbard, Piazza, Pinel, & Dehaene, 2005; Sanfey, 2004). Our finding here in a sense is consistent with previous findings. However, the difference between the two conditions during a specific time window offers an interesting new finding. Given the context, it seems plausible to explain the difference in terms of possible cognitive-affective interactions. Rejectable offers, which are often unfair and unfavorable, induced some sub-cortical emotional activations during that time window, which can delay or short-circuit the rational utility calculation that is supposed to occur. This in turn eventually result in reject responses.

A Neurocomputational Model

Our computational model centers on the concept of expectation. Just like Adam Smith separates "passions" and "impartial spectator", we hypothesize that there is a distinction between rational self and moral self. One major difference between the two is they typically possess different expectations. In the \$10-sum ultimatum game, for the responder, while \$0 is the expectation for the rational self, \$5 is the expectation for the moral self. As a result, an offer of \$3 would lead to a positive difference (e.g., \$3-\$0) for the rational self and a negative difference for the moral self (e.g., \$3-\$5). Since the two differences are in opposite signs, a conflict resolution mechanism then becomes necessary.

Yet another difference between the rational self and moral self is that while the former is emotionless, the latter leads to feelings and emotions, suggesting the moral self is more closely associated to the emotional centers of the brain.

How expectations are derived is a problem of learning. Learning to expect as a basic learning mechanism has been explored extensively in computational modeling frameworks. The well-known Rescorla-Wagner learning rule (Rescorla & Wagner, 1972) can be thought as a mathematical description on how an organism learns to predict the outcome given the cue and has been widely applied to explain human and animal learning. Specifically, it uses the difference between the predicted outcome value and the true outcome value as a training error to systematically improve the prediction. Its equivalent form in machine learning domain, the delta rule, has been the cornerstone in training artificial neural networks to perform nontrivial learning tasks. Temporal difference (TD) learning (Sutton, 1988; Sutton & Barto, 1998) is an extension of the Rescorla-Wagner rule to the continuous time domain. Instead of using the difference between the prediction value and the true outcome value as the train error, in TD learning a TD error is used, which is the difference between the

prediction value at the time t and the prediction value at time $t+1$. As a result, in TD learning an organism does not have to wait until the final outcome is available to learn – it can learn at any time within the trial now simply by comparing the prediction values at any two successive time points, a simple bootstrapping strategy. More complex bootstrapping strategies can be adopted, such as comparing the prediction value at time t with the average prediction value at all later time points, which results in different type of TD error. TD learning as a general computational formulation of human, animal, and machine reinforcement learning and sequential decision making has been generally supported. Recent evidence from cognitive neuroscience has shown that there might exist suitable mechanisms in the brain, particularly in the frontal systems, which implement a similar TD type learning (Holroyd & Coles, 2002; O'Reilly & Munakata, 2000).

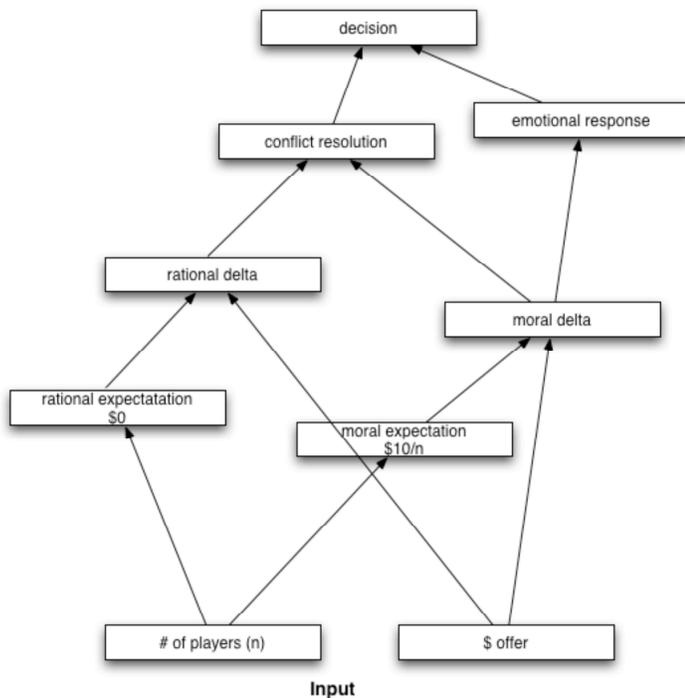


Figure 2: A sketch of the neurocomputational model of the ultimatum game

Recent evidence from cognitive neuroscience research supports our claim that different brain systems join to help make decisions (Damasio, 1995; Lee, 2006; Rolls, 1999; Wagar & Thagard, 2004). We can map the theory to functional brain regions based on these available evidence. Expectations are represented in the parietal lobe (Dehaene, 2002; Dehaene, Spelke, Pinel, Stanescu, & Tsivkin, 1999). Deltas are calculated in the midbrain areas such as VTA and SN (Shizgal & Arvanitogiannis, 2003). The delta signals are broadcasted, via neurotransmitter dopamine, to the frontal cortex (for rational self) and basal ganglia (for moral self). The basal ganglia broadcasts its signals to the limbic system, leading to emotions and feelings (Damasio, 1995,

2001). When the two selves tend to make different decisions (eg, deltas in opposite signs), this invokes ACC, which monitors and resolves the conflict (Brown & Braver, 2005; C. S. Carter et al., 1998).

This mapping naturally leads to the model sketch depicted in Figure 2, which can be implemented and evaluated in biologically realistic neural network modeling frameworks such as leabra (O'Reilly & Munakata, 2000).

Conclusion

With only six subjects, the current study so far has already provided promising results and generated innovative theoretical insights about cognitive-affective interactions in human decision making. It shows rejected offers and accepted offers induced different brain activities. In particular, it seems that some processes in the parietal lobe, which may be utility calculation related based on previous literature, are delayed or short-circuited in the time window of 150-270 ms after the offer onset in the reject response condition. However, it is important to note that although it is reasonable to speculate that the lower sub-cortical emotion-related brain circuits is responsible for such an intervention, without further data analysis and source localization, we cannot make this claim for sure. Clearly, further work is necessary.

Acknowledgments

This work is partially supported by an Air Force Summary Faculty Fellowship awarded to HW in 2005. We would like to thank Dr. Tim Busch and Dr. James Peterson for their collaboration and counseling.

References

- Ashraf, N., Camerer, C. F., & Loewenstein, G. (forthcoming). Adam Smith, Behavioral Economist. *Journal of Economic Perspectives*, *x*, x-x.
- Biddle, E. S., Henninger, A., Franceschini, R., & Jones, R. M. (2003). Emotion Modeling to Enhance Behavior Representation: A Survey of Approaches. In *Proceedings of the '03 Interservice/Industry Training, Simulation and Education Conference (IITSEC)*. Orlando, FL.
- Brown, J. W., & Braver, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science*, *307*, 1118-1121.
- Camerer, C. F. (2003). *Behavioral game theory: Experiments in strategic interaction*. New York: Russell Sage Foundation.
- Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, *280*, 747-749.
- Carter, R. (1998). *Mapping the mind*. Berkeley, CA: University of California Press.

- Damasio, A. R. (1995). *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Putnam.
- Damasio, A. R. (2001). Fundamental Feelings. *Nature*, *413*, 781.
- Davis, D. D., & Holt, C. A. (1993). *Experimental Economics*. Princeton: Princeton University Press.
- Dehaene, S. (2002). Single-Neuron Arithmetic. *Science*, *297*, 1652-1653.
- Dehaene, S., Spelke, E., Pinel, P., Stanescu, R., & Tsivkin, S. (1999). Sources of Mathematical Thinking: Behavioral and Brain-Imaging Evidence. *Science*, *284*, 970-974.
- Elster, J. (1999). *Alchemies of the Mind : Rationality and the Emotions*. New York: Cambridge University Press.
- Elster, J. (2004). Emotions and Rationality. In A. S. R. Manstead, N. Frijda & A. Fischer (Eds.), *Feelings and Emotions*. New York: Cambridge University Press.
- Glimcher, P. W., & Rustichini, A. (2004). Neuroeconomics: The consilience of brain and decision. *Science*, *306*, 447-452.
- Gratch, J., & Marsella, S. (2004). A Domain-independent Framework for Modeling Emotion. *Cognitive Systems Research*, *5*, 269-306.
- Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the Error-Related Negativity. *Psychological Review*, *109*(4), 679-709.
- Hubbard, E. M., Piazza, M., Pinel, P., & Dehaene, S. (2005). Interactions between number and space in parietal cortex. *Nature Reviews Neuroscience*, *6*, 435-448.
- LeDoux, J. E. (1996). *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. New York: Simon & Schuster.
- LeDoux, J. E. (2000). Emotion circuits in the brain. *Annual Review of Neuroscience*, *23*, 155-184.
- Lee, D. (2006). Neural basis of quasi-rational decision making. *Current Opinion in Neurobiology*, *16*(2), 191-198.
- Liang, H., & Wang, H. (2003). Top-down anticipatory control in prefrontal cortex. *Theory in Biosciences*, *122*, 70-86.
- Luck, S. J. (2005). *An introduction to the Event-Related Potential technique*. Cambridge, MA: MIT Press.
- Marinier III, R. P., & Laird, J. E. (2004). *Toward a Comprehensive Computational Model of Emotions and Feelings*. Paper presented at the ICCM-2004.
- Nowak, M. A., Page, K. M., & Sigmund, K. (2000). Fairness Versus Reason in the Ultimatum Game. *Science*, *289*, 1773-1775.
- O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience*. Cambridge, MA: MIT Press.
- Picton, T. W., Bentin, S., Berg, P., Donchin, E., Hillyard, S. A., Johnson, R. J., et al. (2000). Guidelines for using human event-related potentials to study cognition: recording standards and publication criteria. *Psychophysiology*, *37*(2), 127-152.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-98). New York, NY: Appleton-Century-Crofts.
- Rolls, E. T. (1999). *The Brain and Emotion*: Oxford University Press.
- Rugg, M. D., & Coles, M. G. H. (1995). *Electrophysiology of mind: Event-related brain potentials and cognition*. Oxford: Oxford University Press.
- Sanfey, A. G. (2004). Neural computations of decision utility. *Trends in Cognitive Sciences*, *8*(12), 519-521.
- Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., & Cohen, J. D. (2003). The neural basis of economic decision-making in the Ultimatum Game. *Science*, *300*, 1755-1758.
- Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, *36*(2), 241-263.
- Shizgal, P., & Arvanitogiannis, A. (2003). Gambling on Dopamine. *Science*, *299*, 1856-1858.
- Smith, A. (1759). *The theory of moral sentiments*. London: A. Millar.
- Sutton, R. S. (1988). Learning to predict by the method of temporal differences. *Machine Learning*, *3*, 9-44.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Thaler, R. H. (1988). Anomalies: The ultimate game. *Journal of Economic Perspectives*, *2*, 195-206.
- Viner, J. (1925). The utility concept in value theory and its critics. *The Journal of Political Economy*, *33*(4), 369-387.
- Wagar, B. M., & Thagard, P. (2004). Spiking Phineas Gage: A Neurocomputational Theory of Cognitive-Affective Integration in Decision Making. *Psychological Review*, *111*(1), 67-79.