Government Office for Science

Foresight

Mental Capital and Wellbeing:

Making the most of ourselves in the 21st century

State-of-Science Review: SR-E3

Neuroeconomics

Professor Ray Dolan

Wellcome Trust Centre for Neuroimaging

UCL

Summary

Neuroeconomics can be defined as the application of neuroscientific methods to the study of human choice behaviour. This emerging discipline is concerned with proximate causes of behaviour, where the principal interest is in the exercise of choice in decision-making, particularly choice under uncertainty (Camerer, 2003). The field owes its beginnings to a landmark publication by Platt and Glimcher (1999) who showed that activity within neurones in the posterior parietal cortex was sensitive to a manipulation of standard economic variables, specifically magnitude and probability of reward. Neuroeconomics ultimately represents an intellectual convergence between the distinct disciplines of economics, psychology and neuroscience, predicated on the idea that fundamental advances in understanding human economic choice behaviour must ultimately take account of brain activity.

1. Decision-making

Choices between gambles is the model for much of behavioural economics, the 'fruit fly of decision theory' (Kahneman and Tversky, 2000). The optimal strategy for making decisions, where there is uncertainty in outcomes, has long been of interest to mathematicians. The earliest formal solution proposed that decisions under uncertainty should be motivated by maximisation of expected value, a multiplicative function of probability times the value of the outcome. However, this initial formulation ignored the fact that objective and subjective value do not increase linearly.

Daniel Bernoulli, a Swiss mathematician, is credited with formalising a subjectivist approach to decision-making by proposing that people assess the *psychological satisfaction* of wealth, which he termed utility, rather than actual wealth itself (Bernoulli, 1954). In simple terms, this means that satisfaction associated with a change in wealth is closely tied to one's total wealth. This reduces to an observation that, while utility increases rapidly with an initial increase in wealth, it then slows with increasing wealth, giving rise to a concave utility function. Thus, a utility function is a mapping between objective and subjective value. The widely used concept of *marginal utility*, which is often thought of as the cornerstone of all economics, describes a change in utility associated with a unit change in a good, for example, money. It is of interest that a concavity in the utility function has been invoked as an explanation for the well-known phenomenon of risk aversion.

Utility is a highly abstract concept and this raises the problems of how it can be quantified or measured, for example, either on a cardinal or ordinal scale. Empirically, utility is usually derived from observations of an animal or person's behaviour, such that, if option X is consistently chosen over another option, then it is considered to have a higher utility. In the context of choice behaviour, the utility of a particular option for action reflects the desirability of its consequences. Thus, a decision-maker faced with a choice between a set of options should weigh the utility of each choice and choose an option that offers the maximal expected utility. Indeed, maximisation of expected utility, transitivity of preferences (if A is preferred to B and B to C, then A is preferred to C) and description invariance (a decision choice should not be influenced by how it is presented) are at the core of von Neumann and Morgenstern's classic modern formulation of decision-making (von Neumann and Morgenstern, 1947). Despite limitations (Loomes and Sugden, 1983; Loomes, 1988), this subjective utility theory remains the guiding perspective within behavioural economics (Camerer, 2003).

2. Psychological perspectives on decision-making

Early accounts of decision-making were primarily concerned with gambling. However, a wider psychological perspective has led to the concept of utility being extended beyond the monetary domain to reflect the ability of a good to satisfy a want. This extended, and more abstract, concept can be attributed to Jeremy Bentham who suggested that utility should encompass the balance between pleasure and pain. In effect this extended view gives rise to a utility function that reflects a mapping between attainment or consumption and a metric of subjective happiness. Prospect Theory provides one of the best known modern articulations of this more psychological perspective, though at its core it represents a descriptive theory rather than a radical replacement of standard subjective utility theory (Kahneman and Tversky, 1979).

Standard applications of utility theory assume that the outcome of risky prospects are evaluated as states of wealth. Prospect Theory assumes that outcomes are either positive (gain) or negative (loss) deviations from a reference point (Kahneman and Tversky, 1979). As the value function within Prospect Theory is steeper in the loss compared to the gain domain, this can explain the phenomenon of loss aversion. In other words, there is an inequality between the pleasure of a gain and the pain of an equivalent loss. Loss aversion is revealed by two common observations in risky choice: a tendency to reject gambles that offer equal probabilities to win and lose the same amount, coupled with an increase in this aversion with the size of the stakes. Secondly, there is a non-linear weighting of probability such that low probabilities are over-weighted and high probabilities under-weighted, described within Prospect Theory by an S-shaped probability-weighting function.

3. Individual difference in decision-making

Behavioural observations have highlighted a marked individual variation in susceptibility to discrete decision options. A common expression of this behavioural variation is reflected in a sensitivity to variance in outcomes around an option's expected value. This susceptibility can be used to characterise individuals as 'risk-seeking' or 'risk-averse', such that a risk-seeking individual will show preference for a gamble with wide variance, while a risk-averse individual will show preference for an option with minimal variance.

Decision-making also involves weighing the relative value of future compared to immediate rewards, and in this context it might seem intuitive that reward loses a constant proportion of its initial value with time. However, this intuition is contradicted by evidence that discounting the value of future reward is best captured by a hyperbolic function (Ainslie, 1975). This reduces to the fact that a reward loses a large proportion of its value at early stages in time and less value with increasing delay. A prediction from hyperbolic discounting, which is not a feature of exponential discounting, is that preferences between a small, more immediate, and a larger, more delayed reward should be observed to reverse, depending on the time that a choice is made. In effect, hyperbolic discounting explains preference reversals.

Discounting in time can, thus, lead to conflict such as deciding whether to opt for immediate gratification and forego reaping a greater, but delayed, gratification. This is captured by a discount function and, in the example cited above, an individual who foregoes a delayed later (and larger) reward can be considered to show a high discounting rate for the future. Thus, variation in susceptibility for choices across time can provide a formal definition of impulsiveness, reflected in a preference for small immediate, relative to larger, delayed gains.

4. Brain pathology and aberrant economic behaviour

The observation that selective brain injury can impact on economic decision-making, independent of any effect on general intelligence, has been a key observation exemplified in the famous case of Phineas Gage (Harlow, 1868). Following an unfortunate accident, in which a tamping iron was blasted though his skull resulting in a focal brain injury, Gage appeared to lose his ability to make advantageous decisions where the decision process involved weighing up costs and benefits. Many similar patients have now been described in whom the common region damaged is the orbital prefrontal cortex (OFC).

Antonio Damasio, following careful observation of such patients, formulated what is now known as the 'somatic marker hypothesis' which can be seen as a biological theory of choice (Damasio, 1994). The theory proposes that signalling the prospective consequences of options for action can assist in selection of an advantageous response option. According to the theory, patients with ventromedial prefrontal lesions do not make advantageous real-life decisions because they have lost the ability to incorporate predictions regarding the emotional consequence of an action into the decision process. This prospective signalling assumes considerable importance when subjects ponder decisions where there is increasing risk or ambiguity.

The importance of this theory was that it provided a clear link between an economic evaluation of choice and the functions of discrete brain regions. It also opened up the possibility that functions implemented in specific brain regions, such the orbital prefrontal cortex, might encode a distinct computational process related to economic choice; for example, encoding a discount function used to represent the relative value of immediate and delayed rewards.

5. Neuroscience, decision-making and economic behaviour

Rational decision-making involves choosing between available alternatives so as to maximise ensuing benefits. In situations of uncertainty, the problem for the brain is to estimate the expected value (reward probability x magnitude) of alternative actions. As indicated, across a range of disciplines including engineering, economics, and computational science, there has been a realisation that this is a common issue that converges with key perspectives in neuroscience. For example, within computational reinforcement learning theory there is a strong affinity with ideas that decision-making involves integration of reward, reward magnitude and reward timing to provide a representation of action desirability (Glimcher and Rustichini, 2004).

One solution for decision-making under uncertainty is to build predictions based on past experience. This form of prediction-making has been given a formal mathematical description within a class of computational theories referred to as 'reinforcement learning' (Sutton and Barto, 1987; 1981; 1998). In reinforcement learning, a choice of action can be reduced to three steps: firstly, estimating how much reward a particular action will yield; secondly, choice selection based on a comparison of action values among alternatives; lastly, updating action values by virtue of errors in the estimate of the action values. This updating, via an error signal, derives from a signalled discrepancy between predictions of reward and actual realised rewards (in simple terms, this is calculated as the difference between an actual outcome *minus* an expected outcome). Within reinforcement theory, optimisation of choice is achieved by learning the value of each action at each state, estimated in terms of the probability that reward will be received starting from that action. Advantageous actions are then chosen simply by comparing their values, referred to as state-action value function or simply a value function.

6. Linking decision theory to neurobiology

As outlined above, reinforcement learning addresses a fundamental problem of an agent who has to learn optimal choices through trial and error interactions with an environment. It provides one of the best examples of a fit between a computational learning model and empirically observed behaviour of neuronal populations. The critical demonstration has been observations from single neuron recordings in monkey which show that phasic activity within dopaminergic neurones, recorded while an animal performs a simple reward-learning task, encode an error in reward prediction (Schultz, 1997). The observation that dopamine neurones encode a prediction error for reward has provided a considerable impetus for exploring the possibility that other key economic variables may be encoded in neuronal firing patterns (Montague et al., 1996).

The landmark study of Platt and Glimcher (1999) showed that, in monkeys performing a decision task, an experimental variation in the amount, and probability, of reward was encoded in firing rates within a population of neurones in posterior parietal cortex, hitherto thought to be primarily involved in sensorimotor integration. This co-variation provided the first evidence that key economic variables, including the expected value of actual choices, are encoded in neuronal responses. Subsequent work in monkeys has provided examples of encoding of economic variables, such as a demonstration that activity in dopamine neurones determines choice behaviour and that neurones in orbital prefrontal cortex provide a value assignment during economic choice (Padoa-Schioppa and Assad, 2006).

7. Does the human brain encode economic variables?

This question can be addressed using classical lesion deficit models or modern functional neuroimaging techniques. For example, using a lesion deficit approach Sanfey et al. (2003) have shown that, in patients who suffered ventromedial prefrontal lesions (VMPFC), they could distinguish two groups. One group behaved on a gambling task akin to normals manifesting risk aversion, while the other patients displayed a risk-seeking pattern, choosing preferentially from high variance options in a gambling task. The inference from this type of investigation is that the VMPFC may encode risk, formalised as variance around a mean.

Studies of human economic behaviour have been greatly enhanced by neuroimaging tools such as functional magnetic resonance imaging (fMRI). Key early human fMRI experiments have established a robust relationship between distinct patterns of brain activity and reward anticipation and receipt (Knutson et al., 2001), expectation of reward and loss (Breiter et al., 2001), ambiguity (Dickhaut et al., 2003), current versus delayed rewards (McClure et al., 2004). Responses consistent with prediction errors, the key teaching signal within reinforcement learning, have been seen in the context of both Pavlovian and Instrumental learning (O'Doherty et al., 2003). Recently, it has been shown that a positive prediction error is expressed not only for gains but also for losses (Seymour et al., 2007).

A key role for prediction errors is not only to update the value of a state but also to guide actions towards options that maximise reward. An important question is whether there is evidence that dopaminergic prediction errors guide human action learning. A direct linkage to the activity of dopamine during reward learning has recently been established during an instrumental learning task where subjects had to choose between actions associated with different probabilities of obtaining reward (Pessiglione et al., 2006). At outcome, the magnitude of a prediction error measured in the striatum was influenced by a dopaminergic manipulation. Relative to placebo, administration of drugs that enhance (levodopa) or reduce (haloperidol) dopaminergic tone either attenuated or enhanced the prediction error signal. In accord with the proposal that prediction errors guide action choice, subjects under levodopa learnt to choose the most rewarding action relative to placebo, while those under haloperidol treatment were less good than those under placebo in choosing the optimal action. Furthermore, by taking the directly measured reinforcement signal magnitude and incorporating this into a standard learning model, the behavioural choices actually observed under the different drug conditions could be precisely captured. Thus, a dopamine-mediated prediction error is used as a teaching signal to learn expected action values and to favour optimal choice in humans.

The power of brain imaging technologies, when combined with computational theory, for understanding the neural underpinnings of human economic behaviour can also be illustrated in relation to some of the most complex issues in decision-making, particularly where the variables of interest are subjective. This can be illustrated by an example of economic choice where the reward value of distinct options is unknown or can only be approximated. Classically, these situations pose a conflict between an imperative to exploit what is currently estimated on the basis of experience to be the current best option, *versus* an imperative to sample an uncertain, and potentially more rewarding, alternative. This scenario is widely known as the 'explore-exploit dilemma' and has been addressed in computational theory and optimal control in engineering.

If a decision rule involves maximisation of expected value, then this involves choosing the option with the estimated greatest value (exploitation). But it can also entail picking randomly among the alternatives (exploration), as these may be more rewarding. An alternative choice model involves choosing actions with a probability that is weighted by their estimated values, the 'soft max decision rule', where exploration and exploitation are governed by their relative value. A third class of rule involves uncertainty bonuses which augment the probability of an uncertain choice being chosen. In certain implementations of this problem, a soft max solution can account for subjects' actual choices. Based upon application of this decision rule, it is possible to calculate subjective variables such as value predictions, prediction errors and choice probabilities (Daw et al., 2006).

Computational model-derived measures are now being used in conjunction with brain imaging to identify neural regions where activity is correlated with the internal signals of computational models. In making decisions under uncertainty, a computational approach affords a characterisation of neural activity as supporting exploratory or exploitative decisions based upon whether an actual choice was one predicted by the model to be determined by the highest expected value (exploitative), or a response with a lower expected value (exploratory). Without a computational model, it is extremely difficult to approximate these subjective variables. In the example cited, it has now been shown that orbital prefrontal cortex (OFC) activity covaries with exploitative actions driven by value, while anterior frontopolar cortex mediates exploratory actions (Daw et al., 2006). This latter region can be inferred to provide a control mechanism facilitating switching between exploratory and exploitative strategies to allow a less deterministic sampling of the environment and possibly reveal a richer seam of rewards.

8. Emotional heuristics and rational models of choice

A number of striking violations of 'rationality' in decision-making suggest the operation of distinct decision processes in the human brain. A classic example is the 'endowment effect' described by Richard Thaler (Thaler et al., 1988). This can be illustrated by an example of vintage wine which an owner will refuse to sell for £200 but, paradoxically, would not pay as much as £100 to replace. In economic theory, an individual's buying and selling price for a good are assumed to differ solely due to transaction costs and an income effect, both of which are insufficient to account for the degree of deviation seen with the endowment effect. One psychological explanation for the endowment effect comes from Prospect Theory where it is suggested that the carriers of utility are not states (owning or not owning the wine) but changes in states elicited by either acquiring the wine or giving it up. Within this frame of reference, giving up is weighted more than acquisition by virtue of loss aversion.

The potential of brain imaging technologies, such as fMRI, for neuroeconomics is perhaps best exemplified by their ability to ground models of economic choice within the framework of neurobiology. For example, a fundamental axiom in economics is logical consistency across decisions, regardless of how choices are presented – an axiom challenged by a wealth of empirical data most notably in the 'framing effect' (Kahneman and Tversky, 2000; 1979; Tversky and Kahneman, 1981). One theoretical consideration is that the framing effect results from a systematic bias in choice behaviour arising from an affect heuristic (Slovic et al., 2005).

A neurobiological account of the framing effect has now been provided by an fMRI study where subjects had to choose between a 'sure' or a 'gamble' option presented in one of two different 'frames' (De Martino et al., 2006). This manipulation engendered a marked framing effect evident in subjects being risk-averse in the Gain frame (tending to choose the sure over the gamble option), and risk-seeking in the Loss frame (tending to prefer the gamble option). Using fMRI, bilateral amygdala activity was seen when subjects' choices were influenced by the frame, consistent with the operation of an emotional decision system elicited by the frame. Inter-subject variability in susceptibility to the frame enables construction of a 'rationality index' that provides a metric of susceptibility to the frame. Here, a significant correlation was evident between decreased susceptibility to the framing and enhanced activity in orbital prefrontal cortex (OFC), a brain regions, which, when damaged, leads to impaired decision-making (De Martino et al., 2006).

9. Neuroeconomics and behavioural pathologies

Understanding basic mechanisms of psychiatric disease remains the major challenge within clinical neuroscience. The complexity of behavioural pathologies has meant that, although much empirical neurobiological research has been directed at conditions such as depression, addictions, pathological gambling and attention-deficit hyperactivity disorder (ADHD), our actual fundamental knowledge base in relation to the core pathophysiology of these conditions is extremely thin. One critical issue is the lack of a convincing theoretical base to connect behavioural aberrations seen in these disorders to underlying pathological processes. The appeal of neuroeconomics is that it provides a rich theoretical stratum that can motivate more systematic approaches to normative accounts of decision-making and link these to behavioural pathologies.

Arguably, a number of psychiatric disorders can be characterised as displaying aberrant decision-making as their core feature. For example, conditions such as psychopathy, addictions, pathological gambling, and ADHD are all amenable to theoretical accounts derived from neuroeconomics. One strong prediction is that some forms of gambling might reflect extreme discounting for the future or, alternatively, a lack of normal risk aversion. Indeed, an example of this type of approach to the problem of ADHD has already been published: specific deficits in reinforcement learning were seen in some subjects (Williams and Dayan, 2005). In depression, the lack of motivation and anhedonia that characterise this condition might reflect insensitivity to reward outcomes.

There are also good grounds for inferring that key economic control variables such as a disposition to explore or exploit, a propensity to discount future rewards, or indeed sensitivity to variance in outcomes, might all be under modulatory neurotransmitter control, for example, via noradrenaline or serotonin (Doya, 2002). These are the very same neurochemical systems that are targets of much of the pharmacological treatments commonly used in psychiatry. A more theoretical based understanding of these conditions, informed by neuroeconomics, opens up the possibility of pharmacological interventions targeted at core deficits in conditions, rather than the non-specific, standard approach adopted within clinical practice. It is even conceivable that there are distinct, or indeed multiple, deficits within a single diagnostic entity allowing for variability of treatments that are highly targeted at core deficits.

10. Conclusions

Neuroeconomics is a new discipline whose future influence in motivating neuroscientific investigation of human choice behaviour is likely to be considerable. The field represents a multidisciplinary convergence of interest around the topic of decision-making, where the central importance of brain data in informing theory is explicitly acknowledged. The appeal of neuroeconomics lies in its ability to provide a theoretically motivated approach to human decision-making, but where theory is constrained by empirical brain data. Neuroeconomics has considerable societal importance in that its core intellectual quest is likely to provide a theoretical basis for a more informed understanding of behavioural deficits seen in a range of psychiatric, disorders; these are the very disorders whose neurobiological underpinnings have remained elusive. Neuroeconomics also has wider, public policy implications in that critical decisions in relation to key issues such as take-up in organ donation schemes may be critically influenced by how public appeals are framed (in effect, the framing of policies) or what constitutes the default status in the minds of potential donors.

References

Ainslie, G. 1975. Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychol Bull,* 82: 463-496.

Bernoulli, D. 1954. Exposition of a new theory on the measurement of risk. *Econometrica,* 22:23-36.

Breiter, H.C., et al. 2001. Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron,* 30:619-39.

Camerer, C.F. 2003. *Behavioural Game Theory: Experiments in Strategic Interaction.* Princeton, New Jersey: Princeton University Press.

Damasio, A.R. 1994. Descartes' error and the future of human life. *Sci Am,* 271:144.

Daw, N.D. et al. 2006. Cortical substrates for exploratory decisions in humans. *Nature,* 441:876-9.

De Martino, B. et al. 2006. Frames, biases, and rational decision-making in the human brain. *Science,* 313:684-7.

Dickhaut, J. et al. 2003. The impact of the certainty context on the process of choice. *Proc Natl Acad Sci USA,* 100:3536-41.

Doya, K. 2002. Metalearning and neuromodulation. *Neural Netw,* 15:495-506.

Glimcher, P.W. and Rustichini, A. 2004. Neuroeconomics: the consilience of brain and decision. *Science,* 306:447-52.

Harlow, J.M. 1868. Recovery from the passage of an iron rod through the head. *Massachusetts Medical Society,* 2:327-346.

Kahneman, D. and Tversky, A. 1979. Prospect theory: an analysis of decision under risk. *Econometrica.* 47:263-291.

Kahneman, D. and Tversky, A. 2000. *Choices, Values, and Frames.* New York: Cambridge Univ. Press.

Knutson, B. et al. 2001. Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *J Neurosci,* 21:RC159.

Loomes, G. and Sugden, R. 1983. A rationale for preference reversal. *The American Economic Review.* 73:428-432.

Loomes, G. 1988. Further evidence of the impact of regret and disappointment in choice under uncertainty. *Econometrica,* 55:47-62.

McClure, S.M. et al. 2004. Separate neural systems value immediate and delayed monetary rewards. *Science,* 306:503-7.

Montague, P.R., Dayan, P. and Sejnowski, T. J. 1996. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J Neurosci,* 16:1936-47.

Morris, G. et al. 2006. Midbrain dopamine neurons encode decisions for future action. *Nat Neurosci,* 9:1057-63.

O'Doherty, J.P. et al. 2003. Temporal difference models and reward-related learning in the human brain. *Neuron,* 38:329-37.

Padoa-Schioppa, C. and Assad, J.A. 2006. Neurons in the orbitofrontal cortex encode economic value. *Nature,* 441:223-6.

Pessiglione, M. et al. 2006. Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature,* 442:1042-5.

Platt, M.L. and Glimcher, P.W. 1999. Neural correlates of decision variables in parietal cortex. *Nature,* 400:233-8.

Sanfey, A.G. et al. 2003. Phineas gauged: decision-making and the human prefrontal cortex. *Neuropsychologia,* 41:1218-29.

Schultz, W. 1997. Dopamine neurons and their role in reward mechanisms. *Curr Opin Neurobiol,* 7:191-7.

Seymour, B. et al. 2007. Differential encoding of losses and gains in the human striatum. *J Neurosci,* 27:4826-31.

Slovic, P. et al. 2005. Affect, risk, and decision-making. *Health Psychol,* 24:S35-40.

Sutton, R.S. and Barto, A.G. 1981. Toward a modern theory of adaptive networks: expectation and prediction. *Psychol Rev,* 88:135-70.

Sutton, R.F. and Barto, A.G. 1987. A temporal difference model of classical conditioning. *In Proceedings of the 9th Annual Conference of the Cognitive Science Society:* 355-378. Hillsdale, New Jersey: Erlbaun.

Sutton, R.S. and Barto, A.G. 1998. *Reinforcement Learning: an Introduction.* Cambridge, Massachusetts: MIT Press.

Thaler, D.E., Rolls, E.T. and Passingham, R.E. 1988. Neuronal activity of the supplementary motor area (SMA) during internally and externally triggered wrist movements. *Neuroscience Letters,* 93:264-9.

Tversky, A. and Kahneman, D. 1981. The framing of decisions and the psychology of choice. *Science,* 211:453-458.

von Neumann, J. and Morgenstern, O. 1947. *The Theory of Games and Linear Programming.* New York: Wiley.

Williams, J. and Dayan, P. 2005. Dopamine, learning, and impulsivity: a biological account of attention-deficit/hyperactivity disorder. *J Child Adolesc Psychopharmacol,* 15:160-79; discussion 157-9.

All the reports and papers produced by the
Foresight Mental Capital and Wellbeing Project may be downloaded from the Foresight website
(www.foresight.gov.uk).
Requests for hard copies may also be made through this website.

First published September 2008.

The Government Office for Science.