

Focus: Decision making and cognitive biases

BlueSci explores biases in human cognition, from their evolutionary background to their socio-political implications

FRIDAY, 16 MAY 2014

Humans, and organisms in general, are faced with different possibilities and trade-offs in every aspect of their lives. When to reproduce or divide? When to fight or flee? Which political party to elect? Rationality is held by many people as intrinsic to human behaviour. We assume that rational decision-making underlies our cognition, and that we make decisions by finely weighing options and thoughtfully calculating outcomes. However, decision-making processes in human and non-human animals are pervaded with biases. In this article we discuss some of the most striking cognitive biases known in human decision-making, focusing on the economic, social and political implications that they can have. We then contrast these processes with less-biased decision-making programs in computers, and close the article by discussing the potential evolutionary explanations behind cognitive biases: why are they there?

According to economists, if humans are 'rational' and free to make choices then they should behave according to the 'Rational Choice Theory' (RCT), which states that people should make decisions by determining the value of a potential outcome, how likely this outcome is, and multiplying these factors. Daniel Kahneman and Amos Tversky were two of the first social psychologists to empirically prove that this is frequently not the case, and that decision-making processes have many biases. After their first work in the 1970s, the list of known cognitive biases has increased greatly, and their work set out a paradigm in psychology and economics of humans as irrational decision-making agents.

Let's start by describing a bias in an everyday situation. We may think that the opinion we have of somebody depends on a deep scrutiny of what we know about them, but in reality this is not the case. We tend to have limited information about people, and we extrapolate this information to other aspects of their personality or abilities. This phenomenon, known as the 'halo effect', was first described by Edward Thorndike in 1920 after he asked commanding officers to evaluate their soldiers according to the physical qualities and intellectual capabilities they possessed. He found that whenever a soldier scored well in one category, he would rank well in all categories. Basically we assume that because someone is good in task A, then he or she will also be good in task B. The halo effect has been studied in different scenarios with remarkable results being found. For instance, students have been shown to rank essays as better written if they think the author is attractive than if they perceive the author as unattractive. Another study carried out in 1975 found a correlation between the attractiveness of criminals and the harshness of the sentence they would receive: the more attractive the criminal, the less likely they were to be convicted, or the lighter the sentence that they would receive.

Political campaigns too can be influenced by the halo effect. A study carried out by Melissa Surawski and Elizabeth Ossoff in 2006 confirmed this by asking participants to rate the political skills of candidates after showing them photographs and playing audio clips of them. They found that physical appearance predicted the candidate's competency better than voice. A recent study on US congressional elections also showed that winning was strongly affected by inferences of abilities based on facial appearance.



Another cognitive bias that can have profound socio-political implications is the 'bandwagon effect', where individuals support the opinion of the majority of their peers. This effect is particularly relevant in politics, where voters may alter their decision to match the majority view and hence be on the winner's side. For example, in the 1992 US presidential election, Vicki Morwitz and Carol Pluzinski conducted a study in which they exposed a group of voters to national poll results indicating that Bill Clinton was in the lead, while keeping another group unaware. They found that a number of voters in the former group who intended to vote for George W. Bush changed their preference after seeing these results, while the latter group didn't change their decision. Studies have also looked at a counter-acting bias known as the 'underdog preference', a rarer phenomenon than the bandwagon effect, wherein some people may support the less-favoured candidate or sports team in a match to feel fair.

How about when we have to decide between two options? If we were to make a rational decision when comparing two alternative options then the relative proportion of choices made between them should be the same regardless of whether they are presented to us on their own, or whether a less-preferred, third option is presented alongside. Nevertheless, this less-preferred third option can make us change our opinion about the other two, better options. This is the 'decoy effect', and it happens when a preference between two options changes when shown a third option that is asymmetrically-dominated, meaning it is inferior in some way and superior in another to one option but inferior in all aspects to the other option. To illustrate this, imagine trying to decide between two train tickets to London, which vary in price and duration of the journey:

	Ticket B
	£30
	40 min

It may be difficult to decide between these two, and some people may decide for ticket A and some for ticket B depending on their priorities. But what would happen if we included a third option that was more expensive than both A and B, but took longer than A?

Ticket B	Ticket C
£30	£45
40 min	35 min

In this case, the asymmetrical decoy (Ticket C), would make option A more appealing than option B. Interestingly, the decoy effect has also been described in non-human animals. For instance, Melissa Bateman

and her colleagues studied the foraging behaviour of hummingbirds by placing birds in two different scenarios. In one scenario the birds had to choose between two alternative artificial flowers, and in the other scenario they had to choose between the same two flowers and an asymmetrical decoy:

A	Scenario B	
	Quantity	Conc
30%	15 μ l	40%
30%	45 μ l	30%
10%	10 μ l	35%

The target flower offered less food but a higher sucrose concentration than the competitor. The asymmetrical decoy in scenario B was worse than the target in both aspects and better than the competitor in terms of sucrose concentration, but not in the total amount of food that it provided. Birds that were exposed to the decoy option showed a higher preference for the target flower than birds that did not have access to the decoy.

The decoy effect has been studied in many different circumstances with remarkable effects being found. For instance, Joel Huber, a marketing professor at Duke University, showed how it can affect our preference for restaurants. He asked people if they would prefer to have dinner in a five-star restaurant that was farther away or in a closer three- star restaurant. With these alternatives some people would prefer the five-star restaurant and some the three-star restaurant. But once he included a third option, a four-star restaurant that was farther away than the first two restaurants, then more people preferred the five-star restaurant over the three-star one.

A well-known scenario where biases blur our choice is when we are confronted with risky decisions. Let's imagine that flooding in a city prompted the council to announce two emergency schemes to limit the damage from the flooding: we could either pile up sandbags everywhere, which would save about a quarter of the city from flooding, or we could invest in an experimental and risky flood barrier, which has a twenty five per cent chance of saving the entire city and a seventy five per cent chance of failing and saving nothing. Which solution would you choose? Now imagine that an updated risk assessment comes in about these two solutions. It turns out that piling up sandbags is actually

going to cause three quarters of the city to be flooded, while the flood barrier has a seventy five per cent chance of causing the whole city to be flooded, and a twenty five per cent chance of preventing the flooding completely. Which solution would you choose this time?

Most people would have chosen the sandbags the first time, and the flood barrier the second. But as some of you may have noticed, the options were identical in both cases. All that changed was the framing: describing the positive effects of the measures the first time, and the negative effects the second. This 'framing effect' is related to our instinctive loss aversion. We hate losing out on something we already have, and will take riskier decisions (the uncertain flood barrier) to try and maintain it. On the other hand, when it comes to gains, we prefer a decision with a certain and reliable outcome (the sand bags). This causes our behaviour to change, in identical circumstances, depending on whether an outcome is presented as a potential loss or a potential gain. These are just some examples of dozens of known, named and heavily researched systematic biases that exist in human decision-making. But is there a way of making un-biased decisions?



Thanks to the work of Kahneman and Tversky, computer-based decision-making has come to replace human decision-making in many walks of life, reducing the biases in some choices. Think of Moneyball, the story of the team that revolutionized baseball by ditching the tips of talent spotters in favour of statistical analysis in order to buy undervalued players and turn their prospects around. These techniques, once employed only by statistically-gifted mavericks, have become commonplace in courtrooms and hospitals, where a failure to recognise biases can have drastic consequences. A popular way of getting computers to do this relies on a branch of maths called 'Bayesian inference'.



Generally, theories can predict how experiments should work, but real data is often messy and experiments frequently go wrong. Scientists need to work backwards from the experiments to decide which theories are best, and to modify these theories to fit better with the observed results.

The mathematical principle behind how scientists do this is termed Bayesian inference: what should we believe, and how confident of it should we be, given the data? It's basically a method of updating probabilities when new information is acquired. Bayesian inference requires 'priors', or information of how likely a theory is before starting an experiment. Priors used in Bayesian inference

can be based on previous experience or general knowledge of the world. Not expecting a particular result more than any other can be the best choice in some cases.

A common human bias is that of ignoring 'statistical priors' when evaluating new information. For instance, in one of Kahneman and Tversky's experiments they gave subjects descriptions of several people's personalities, and asked them to guess if they were engineers or lawyers. The descriptions gave no occupation-specific information, but one group of subjects were told that the people came from a population with seventy per cent lawyers and thirty per cent engineers. The other group were told that the percentages were reversed. Bayesian inference states that if there are more engineers present, we should guess that more of the people described are engineers. However, in the experiment the two groups entirely ignored the information on the population as a whole and gave the same responses, no matter how vague the description was. Only if no personality description was given did the two groups give predictions that resembled the population divide.

Bayesian theory is often used as the basis of machine-learning, where computers predict or discover things about the world from large data-sets. For instance, naïve Bayesian classifiers are used to infer whether people who have been arrested are likely to commit crimes again if released on bail. To do this, first a program is trained with existing data on people previously released on bail, and information on whether or not they committed crimes before the trial was over. The program then establishes relationships between certain variables, or risk-factors, such as criminal records, social and economic situation, and reoffending rates. Afterwards the program can be used to predict the outcome of releasing criminals, with the first set of learnt relationships used as priors. When comparing the output of the program with decisions made by judges, it transpired that judges' subjective decisions had been completely misguided. Between their personal biases, and the fact that some people could not afford to post bail, the judges effectively released people at higher risk of re-offending while on bail.

The outstanding prevalence of biases in decision-making processes begs the question of why we, and animals in general, have these biases to begin with? Could there be an advantage to them or are they just an evolutionary solution to compensate for our limited cognitive ability?

Every individual, from the smallest virus to the largest whale, faces trade-offs during its lifetime. Resources are limited, and individuals are selected to maximise their lifetime reproductive success: getting the highest benefits with the lowest costs, according to the resources available. This is the basic principle of natural selection, the main driving force of evolution. Cognition is no exception. Our brain has been moulded by natural selection throughout our evolutionary history to adapt to the environment under which it evolved. Under this perspective the mind is 'adaptively rational', comprising a set of tools designed by natural selection to deal with situations our ancestors encountered. This concept has led to the idea of 'ecological rationality', which suggests that cognitive biases are a result of adaptive solutions to the decision-making problems of our evolutionary past, that maximised the ratio between the benefits and costs of decision-making.



Two of the strongest arguments linked to ecological rationality are 'heuristics' and 'error management'. Heuristics are efficient solutions to problems when information, time, or processing capabilities are constrained, while error management suggests that natural selection favours biases towards the least-costly error. An experiment that illustrates the idea behind heuristics was done by Andreas Wilke in 2006; he studied the foraging behaviour of people—when they would decide to leave their current resource patch and move onto a new patch—under different resource-distribution situations: random, evenly-aggregated or evenly- dispersed. He found that humans used the same set of rules to change resource patches independently of the distribution of food, always using rules that were particularly useful for aggregated resources. Although this result seems irrational, if we consider how resources are distributed in nature it becomes less puzzling; aggregated patches are more common in nature because species are not independent of one another—they tend to attract or repel each other. Using a straightforward rule that is useful in most patches can be a more efficient solution than carefully calculating the amount of food in each patch and deciding when to leave afterwards.

The idea of error management comes from the theory that eliminating all possible errors when making decisions may be impossible, so there may be selection to reduce the most-costly errors. For instance, being able to identify poisonous animals like snakes requires observation and identification of the object. There are two possible outcomes to this: a person could properly examine a potential threat and correctly classifying it as a snake, or could walk away without the certainty that it was in fact a snake. Either way, an error is likely to occur, so it may be more advantageous to have a bias that makes people run away from things that resemble snakes, even if sometimes they run away from lifeless objects, than to get close to the object to decide if it is actually a snake or not. In this second case, the likelihood of being bitten will be larger, as is the cost of the error. Error management may increase the overall rate of making errors, but minimises the overall costs of the errors made.


Ecological rationality leaves us with two ideas. One is that our brain has an optimal design that maximised the benefit-to-cost ratio of cognition during our evolutionary history. An optimal design need not necessarily be a 'perfect design', since the costs associated with producing and maintaining it could be too high, but another hypothesis that emerges is that, under some circumstances, biases could have a selective advantage. Whatever the reason behind our cognitive biases, one thing is certain: the study of how our brain works and how it evolved is a topic that has brought together the minds of economists, psychologists, philosophers, physicists, engineers, neuroscientists, medics, mathematicians and biologists. It is perhaps one of the most interdisciplinary lines of research, and one of the most interesting mysteries of our time.

Alex O'Bryan Tear is a 2nd year PhD student at the Department of Psychiatry

Robin Lamboll is a 1st year PhD student at the Department of Physics

Shirin Ashraf is a 2nd year PhD student at the Department of Immunology

Ornela De Gasperin is a 3rd year PhD student at the Department of Zoology



Browse international scientific events,
from conferences to seminars

nature
careers