Herding, social influences and behavioural bias in scientific research

Simple awareness of the hidden pressures and beliefs that influence our thinking can help to preserve objectivity

Michelle Baddeley

The mission of scientific research is to understand and to discover the cause or mechanism behind an observed phenomenon. The main tool employed by scientists is the scientific method: formulate a hypothesis that could explain an observation, develop testable predictions, gather data or design experiments to test these predictions and, based on the result, accept, reject or refine the hypothesis. In practice, however, the path to understanding is often not straightforward: uncertainty, insufficient information, unreliable data or flawed analysis can make it challenging to untangle good theories, hypotheses and evidence from bad, though these problems can be overcome with careful experimental design, objective data analysis and/or robust statistics. Yet, no matter how good the experiment or how clean the data, we still need to account for the human factor: researchers are subject to unconscious bias and might genuinely believe that their analysis is wholly objective when, in fact, it is not. Bias can distort the evolution of knowledge if scientists are reluctant to accept an alternative explanation for their observations, or even fudge data or their analysis to support their preconceived beliefs. This article highlights some of the biases that have the potential to mislead academic research. Among them, heuristics and biases generally and social influences in particular, can have profoundly negative consequences for the wider world, especially if misleading research findings are used to guide public policy or affect decision-making in medicine and beyond.

The challenge is to become aware of biases and separate the bad influences from

the good. Sometimes social influences play a positive role—for example, by enabling social learning. Condorcet’s “jury principle” is another example of the power of collective wisdom: the collective opinion of a jury—in which each individual juror has just a slightly better than average chance of matching the correct verdict—is more likely to reach the correct verdict, but only if the individuals’ judgements are uncorrelated. In other situations, social influence and collective opinions are unhelpful—for example, if people follow a group consensus even though they have private information which conflicts with the consensus. If researchers are aware of these pitfalls and the biases to which they might be prone, this greater awareness will help them interpret results as objectively as possible and base all their judgements on robust evidence.

Many mistakes that people make are genuine and reflect systematic cognitive biases. Such biases are not necessarily irrational, in the sense of being stupid, because they emerge from the sensible application of heuristics or quick rules of thumb—a practical approach to solve problems that is not perfect or optimal, but is sufficient for the task at hand. Herbert Simon, an American Nobel laureate in economics, political scientist, sociologist and computer scientist, analysed rationality and his insights are helpful in understanding how heuristics are linked to socio-psychological influences affecting experts’ beliefs. Simon distinguished substantive rationality—when decisions have a substantive, objective basis—usually based around some mathematical rule—from procedural rationality, when

decision-making is more sensible, intuitive, based on prior judgements, and “appropriate deliberation” [1].

“Bias can distort the evolution of knowledge if scientists are reluctant to accept an alternative explanation for their observations, or even fudge data or their analysis to support their preconceived beliefs”

Using heuristics is consistent with Simon’s definition of procedural rationality: heuristics are reasoning devices that enable people to economise the costs involved in collecting information and deciding about their best options. It would be foolish to spend a week travelling around town visiting supermarkets before deciding where to buy a cheap loaf of bread. When planning a holiday, looking at customer reviews may save time and effort, even if these reviews give only a partial account. Similarly, heuristics are used in research: before we decide to read a paper, we might prejudice its quality and make a decision whether or not to read it depending on the authors’ publication records, institutional affiliations or which journal it is published in. A more reliable way to judge the quality of a paper is to read all other papers in the same field, but this would involve a large expenditure of time and effort probably not justifiable in terms

of the net benefit. Most of the time, heuristics can work well, but sometimes they generate systematic cognitive/behavioural biases, which can be grouped into three main sets, as originally identified by Tversky and Kahneman: the availability heuristic, the representativeness heuristic and anchoring/adjustment [2].

Availability heuristics are used when people form judgements based on readily accessible information—this is often the most recent or salient information—even though this information may be less relevant than other information which is harder to remember. A well-known example of the availability heuristic is people's subjective judgements of the risks of different types of accidents: experimental evidence shows that people are more likely to overestimate the probability of plane and train crashes either when they have had recent personal experience or—more likely—when they have read or seen vivid accounts in the media. Objectively, car and

pedestrian accidents are more likely, but they are also less likely to feature in the news and so are harder to recall.

Problems emerge in applying the availability heuristic when important and useful information is ignored. The availability heuristic also connects with familiarity bias and status quo bias: people favour explanations with which they are familiar and may therefore be resistant to novel findings and approaches. Research into the causes of stomach ulcers and gastric cancer is an illustrative example. The conventional view was that stress and poor diet causes stomach ulcers, and when Barry Marshall and colleagues showed that *Helicobacter pylori* was the culprit—for which he received the Nobel prize with Robin Warren—the findings were originally dismissed and even ridiculed, arguably because they did not fit well with the collective opinion.

The representativeness heuristic is based on analogical reasoning: judging events and processes by their similarity to other events and processes. One example relevant to

academic research is Tversky and Kahneman's "law of small numbers," by which small samples are attributed with as much power as large samples. Deena Skolnick Weisberg and colleagues identified a similar problem in the application of neuroscience explanations: their experiments showed that naïve adults are more likely to believe bad explanations when "supported" by irrelevant neuroscience and less likely to believe good explanations when not accompanied by irrelevant neuroscience [3].

Finally, Kahneman and Tversky identified a category of biases associated with anchoring and adjustment heuristics. People often anchor their judgements on a reference point—this may be current opinion or the strong opinions of a research leader or other opinion former. Adjustment heuristics connect with confirmation bias: people tend to interpret evidence that connects with their preconceived notions of how the world works. In this case, beliefs will be path dependent: they emerge according to what has happened before. As explored in more detail

below, anchoring and adjustment heuristics can help to explain the impact of herding and other social influences—when individuals anchor and adjust their own judgements around socially determined reference points, including other experts' opinions.

Social influences come in two broad forms: informational influence—others' opinions that provide useful information—and normative influence—agreeing with others based on socio-psychological and/or emotional factors. Informational influences are the focus of economic models of “rational herding” and social learning, based on Bayesian reasoning processes, when decision-makers use information about the decisions and actions of others to judge the likelihood of an event. Such judgements are regularly updated according to Bayes's rule and therefore are driven by relatively objective and systematic information.

“Adopting consensus views may mean that potentially useful private information, especially novel and unexpected findings, is ignored and discarded, and so is lost to subsequent researchers”

Social learning models are often illustrated with the example of choosing a restaurant. When we see a crowded restaurant, we infer that its food and wine are good because it is attracting so many customers. But a person observing a crowded restaurant may also have some contradictory private information about the restaurant next door: for example, a friend might have told them that the second restaurant has much better food, wine and service; yet that second restaurant is empty. If that person decides in the end to go with the implicit group judgement that the first restaurant is better, then their hidden private information (their friend's opinion) gets lost. Anyone observing them would see nothing to suggest that the empty restaurant has any merit—even if they have contradictory private information of their own. They too might decide, on balance, to go with the herd and queue for the first restaurant. As more and more people queue for the first restaurant, all useful private

information about the superior quality of the second restaurant is lost.

This problem can be profound in scientific research. Adopting consensus views may mean that potentially useful private information, especially novel and unexpected findings, is ignored and discarded and so is lost to subsequent researchers. Evidence that conflicts with established opinions can be sidelined or deemed unpublishable by reviewers who have a competing hypothesis or contrary world view. In the worst cases, the researchers who uncover evidence that fits well with unfashionable or unconventional views may be ostracised or even punished, with well-known historical examples, not least Galileo Galilei, who was convicted of heresy for his support of the Copernican heliocentric model of the solar system.

Herding, fads and customs can also be explained in terms of reputation building. When people care about their status, conformity helps them to maintain status, while departing from social norms carries the risk of impaired status. Reputation also survives a loss better if others are losing at the same time. If financial traders lose large sums when others are losing at the same time, they will have a good chance of keeping their job; but if they lose a large sum implementing an unconventional trading strategy, there is a good chance they will lose their job, even if, overall, their strategy is sound and more likely to succeed than fail. People often make obvious mistakes when they observe others around them making similar mistakes. In social psychology experiments, when asked to judge the similarity in length of a set of lines, subjects were manipulated into making apparently obvious mistakes when they observed experimental confederates deliberately giving the wrong answers in the same task—they may agree with others because it is easier and less confusing to conform [4].

The propensity to herd is strong and reflects social responses that were hard-wired during evolution and reinforced via childhood conditioning. Insights from neurobiology and evolutionary biology help to explain our herding tendencies—survival chances are increased for many animals when the group provides safety and/or gives signals about the availability of food or mates. Some neuroscientific evidence

indicates that, partly, herding activates neural areas that are older and more primitive in evolutionary terms [5]. Herding also reflects childhood conditioning. Children copy adult behaviours, and children who have seen adults around them behaving violently may be driven by instinctive imitation to behave violently too [6].

Social influences, including social pressure, groupthink and herding effects, are powerful in scientific research communities, where the path of scientific investigation may be shaped by past events and others' opinions. In these situations, expert elicitation—collecting information from other experts—may be prone to socially driven heuristics and biases, including group bias, tribalism, herding and bandwagon effects. Baddeley, Curtis and Wood explored herding in “expert elicitation” in geophysics [7]. In geologically complex rock formations, uncertainty and poor/scarcely data limit experts' ability to accurately identify the probability of oil resources. Bringing together experts' opinions has the potential to increase accuracy, assuming Condorcet's jury principle about the wisdom of crowds holds, and this rests, as noted above, on the notion that individuals' prior opinions are uncorrelated. Instead, expert elicitation is often distorted by herding and conventional opinions and conformist views will therefore be overweighted.

“The propensity to herd is strong and reflects social responses that were hard-wired during evolution and reinforced via childhood conditioning”

Perverse incentives exacerbate the problem. When careers depend on research assessment and the number of publications in established journals, the incentives tip towards following the crowd rather than publicising unconventional theories or apparently anomalous findings [8]. When herding influences dominate, the accumulation of knowledge is distorted. Using computational models, Michael Weisberg showed that, with a greater proportion of contrarians in a population, a wider range of knowledge will be uncovered [9]. We need contrarians

as they encourage us to pursue new directions and take different approaches. The willingness to take risks in research generates positive externalities, for example: new knowledge that the herd would not be able to discover if they stick to conformist views.

.....

“Scientific evidence can and should be interpreted keeping these biases in mind”

.....

In the worst case, social influences may allow fraudulent or deliberately distorted results to twist research if personal ambition, preoccupation with academic status and/or vested interests dominate. A recent illustration was the Diederik Stapel case: Stapel is a social psychologist and he manipulated his data from studies of the impact of disordered environments on antisocial behaviour. Marc Hauser, a former professor of psychology at Harvard University, published influential papers in top journals on animal behaviour and cognition, until an investigation found him guilty of scientific misconduct in 2011. Both were influential, leading figures in their field, and their results were unchallenged for many years, partly because members of their research groups and other researchers felt unable to challenge them. Their reputations as leading figures in their fields meant it took longer for whistle-blowers and critiques questioning the integrity of their data and findings to have an impact.

Deliberate fraud is rare. More usually, mistakes result from the excessive influence of scientific conventions, ideological prejudices and/or unconscious bias; well-educated, intelligent scientists are as susceptible to these as anyone else. Subtle, unconscious conformism is likely to be far more dangerous to scientific progress than fraud: it is harder to detect, and if researchers are not even aware of the power of conventional opinions to shape their hypotheses and conclusions, then conformism can have a detrimental impact in terms of human wellbeing and scientific progress. These problems are likely to be profound, especially in new fields of research. A research paper that looks and sounds right and matches a discipline's

conventions and preconceptions is more likely to be taken seriously irrespective of its scientific merit. This was illustrated in the case of the Sokal hoax in which a well-written but deliberately nonsensical research paper passed through refereeing processes in social science journals, arguably because it rested well with reviewers' preconceptions. Another salient example is tobacco research: initial evidence about a strong correlation between cigarette smoking and lung cancer was dismissed on the grounds that correlation does not imply causation, with some researchers—including some later hired as consultants by tobacco companies—making what now seems an absurd claim that the causation went in reverse with lung cancer causing cigarette smoking [10].

Other group influences reflect hierarchies and experience, if, for instance, junior members of a research laboratory instinctively imitate their mentors, defer to their supervisors' views and opinions and/or refrain from disagreeing. When researchers—particularly young researchers with careers to forge—feel social pressure to conform to a particular scientific view, it can be difficult to contradict that view, leading to path dependency and inertia.

Scientific evidence can and should be interpreted keeping these biases in mind. If researchers support an existing theory or hypothesis because it has been properly verified, it does not mean that the consensus is wrong. More generally, social influences can play a positive role in research: replicating others' findings is an undervalued but important part of science. When a number of researchers have repeated and verified experimental results, findings will be more robust. Problems emerge when the consensus opinion reflects something other than a Bayesian-style judgement about relative likelihood. When researchers are reluctant to abandon a favoured hypothesis, for reasons that reflect socio-psychological influences rather than hard evidence, then the hypothesis persists because it is assigned excessive and undue weight. As more and more researchers support it, the likelihood that it will persist increases, and the path of knowledge will be obstructed. Journal editors and reviewers, and the research community

more generally, need to recognise that herding and social influences can influence judgement and lead them to favour research findings that fit with their own preconceptions and/or group opinions as much as objective evidence.

Conflict of interest

The author declares that she has no conflict of interest.

References

1. Simon H (1979) From Substantive to Procedural Rationality. In *Philosophy and Economic Theory*, Hahn FH, Hollis M (eds). Oxford: Oxford University Press
2. Tversky A, Kahneman D (1974) Judgment under uncertainty: heuristics and biases. *Science* 185: 1124–1131
3. Weisberg DS, Keil FC, Goodstein J, Rawson E, Gray JR (2008) The seductive allure of neuroscience explanations. *J Cogn Neurosci* 20: 470–477
4. Asch SE (1956) Studies of independence and conformity: a minority of one against a unanimous majority. *Psychol Monogr* 70: 1–70
5. Baddeley M (2010) Herding, social influence and economic decision-making: socio-psychological and neuroscientific analyses. *Phil Trans R Soc B* 365: 281–290
6. Bandura A, Ross D, Ross SA (1961) Transmission of aggression through imitation of aggressive models. *J Abnorm Soc Psychol* 63: 575–582
7. Baddeley M, Curtis A, Wood R (2005) An introduction to prior information derived from probabilistic judgments; elicitation of knowledge, cognitive bias and herding. In *Geological Prior Information: Informing Science and Engineering*, Curtis A, Wood R (eds), Special Publications No. 239, pp 15–27. London: Geological Society
8. Baddeley M (2013) Herding, social influence and expert opinion. *J Economic Methodol* 20: 37–45
9. Weisberg M (2013) Modeling herding behaviour and its risks. *J Economic Methodol* 20: 6–18
10. Harford T (2014). Big data: are we making a big mistake? *The Financial Times Magazine*, 28 March 2014. Downloaded 28 May 2015 from <http://www.ft.com/cms/s/2/21a6e7d8-b479-11e3-a09a-00144feabdc0.html>