

Tim Hughes

Applying Cognitive Perspectives on Decision-Making to the Policy Advice Process: a practitioner's view

Behavioural economics and the related fields of cognitive and social psychology are now very much in the mainstream, as the highly visible success of the Behavioural Insights Unit in the United Kingdom attests. A robust and diverse range of findings about the limits of human thinking challenges policy practitioners to reconsider how they both design and advise on policies. This challenge is particularly relevant given that the training and background of policy advisors typically does not include these fields, with political science, law and conventional economics much more common.

A range of recent books have popularised many concepts from these fields and are leading an increasing number of people outside academia to revisit the way we conceive of thinking and decision-making. For example, *The Wisdom of Crowds* (Surowiecki, 2004), *Blink* (Gladwell, 2005), *The Black Swan* (Taleb, 2007), *Predictably Irrational* (Ariely, 2008), *Nudge* (Thaler and Sunstein, 2008), *Thinking Fast and Slow* (Kahneman, 2011) and *The Signal and the Noise* (Silver, 2012) all underline the limitations of rational accounts of thinking and decision-making.

Perhaps reflecting the new public popularity of these fields, it has become fashionable in certain circles to consider ways to incorporate the findings of cognitive psychology and behavioural economics into the design of policies (e.g. Ministry of Economic Development, 2006; Dolan et al., 2010), often under the label libertarian paternalism or choice architecture (Thaler, Sunstein and Balz, 2010). The argument is often that small changes in the design of policies can nudge choices in a desired direction

Tim Hughes studied economics at the University of Auckland and has worked in policy roles at the Department of Corrections and the Ministry of Social Development. He is currently working in the sector investment team at the Ministry of Justice.

without the need for compulsion. Perhaps the best known example is the design of KiwiSaver, where the default option is automatic enrolment, with people required to opt out instead of opt in.

Cognitive science has also influenced the academic analysis of policy, particularly of the way cognitive biases can affect the agenda-setting process and the public communication of policies (Araral et al., 2013).

In this article I take a practitioner's perspective on how the insights of cognitive psychology and behavioural economics also have important implications for the nuts and bolts of policy advice. I focus particularly on the implications of the biases and heuristics literature for the way we structure choices and provide information to decision-makers. I argue that this literature provides a formal language and a toolkit to help policy advisors better understand the implications of the advisory choices they make.

The biases and heuristics tradition

Neoclassical economics typically relies on assumptions of rationality in agents in a strict sense. The rational actor paradigm has influenced a range of fields, including political science and organisational decision-making (Jones, Boushey and Workman, 2006). At the same time, there is a long tradition of treating the decision-making process as descriptively or normatively quasi-rational, or boundedly rational. For example, Herbert Simon (1956) observed that decision-makers often settle for a good enough option, or satisfice, rather than seeking an optimal solution among all possible alternatives.

The biases and heuristics tradition has built further evidence to counter the assumptions of the rational decision model. Led initially by Daniel Kahneman and Amos Tversky in the 1970s, there is now a wide range of findings showing experimentally where judgements differ descriptively from those normatively prescribed by the rational actor paradigm. Researchers have identified a range of heuristics, or thinking shortcuts, that we commonly use, and which can lead to judgemental biases in many circumstances.

The biases and heuristics tradition has now been incorporated into a general set of dual process theories in psychology. It is now understood that there are essentially two quite different thinking systems we use, including to make decisions (Chaikain and Trope, 1999).

System 1 is the fast system of *Thinking Fast and Slow*, and which operates automatically to produce an ongoing stream of unconscious judgements using the brain's basic associative machinery. System 2 is the slow, deliberate, conscious type of reasoning that more closely adheres to the type of thinking prescribed by a rational choice model. In general we use the rapid processes of system 1 wherever possible, as system 2 is effortful to use.

One classic example of the difference between the two systems is the following problem from the Cognitive Reflection Test (Toplak et al., 2011):

A bat and a ball cost \$1.10 in total.
The bat costs \$1 more than the ball.
How much does the ball cost?

System 1 will typically generate an obvious and intuitive heuristic answer unbidden: 10 cents. When asked this question, most subjects will in fact give that answer and express confidence that it is the right answer, including large proportions of highly select university students at MIT, Princeton and Harvard (Frederick, 2005). But only a small amount of thought with system 2 is enough to convince oneself that the correct answer is actually 5 cents. The difficulty is that to arrive at the right answer system 1 needs to be deliberately overridden, and often isn't.¹

Our minds use many heuristics to generate what Gerd Gigerenzer calls 'fast and frugal' responses to the world, and this strategy is largely adaptive. In some

cases, however, reliance on heuristics can lead to biased or inaccurate thinking. If our minds can be fooled by such a simple task as the bat and ball problem, it seems reasonable to ask whether similar errors of judgement can be made in more complex, policy-relevant domains, by both advisors and decision-makers.

Application to policy advisors and decision-makers

We are not in a position where we can definitively say that policy analysts and political decision-makers are riddled with thinking biases and judgemental errors. Some of the biases and heuristics literature has been criticised as narrow and artificial, with many results relying largely on American undergraduates and

... decision-makers can take different decisions based on exactly the same information depending on how that information is presented.

relatively few studies conducted in real-life conditions (Klein, 1998; Swoyer, 2002). Others have suggested that the language of bias is a misleading way to describe the phenomena revealed by the literature (e.g. Gigerenzer, Hoffrage and Kleinbolting, 1991).

However, several phenomena have been demonstrated across multiple populations, and there is robust evidence of several phenomena that should be concerning regardless of whether we use the language of 'bias'. For example, we should be concerned by findings that decision-makers can take different decisions based on exactly the same information depending on how that information is presented.

We might also wonder whether the high levels of cognitive ability that are typically present among policy advisors and decision-makers protect against biases. In recent years, researchers have focused on establishing boundary conditions and individual differences in many of the most important biases, to see whether certain individuals or situations

reduce biases. Indeed, not all biases appear to be universal. But Stanovich and West (2008) and others have found that several biases are uncorrelated with cognitive ability, and that cognitive ability does not always eliminate other biases even where it reduces them.

There is also evidence that people tend to think they are less subject to cognitive biases than they are, and that if anything this 'bias blind spot' may be larger for more intelligent people (West, Meserve and Stanovich, 2012). It seems reasonable to assume that the intelligence of policy advisors and decision makers is not sufficient to protect against all thinking biases or the bias blind spot.

Even if intelligence is not protective against all thinking biases, one could

Shanteau (1992) has offered a general framework for when expert intuition is likely to lead to effective judgements. In general, he argues, experts with demonstrably good performance, such as weather forecasters, chess masters and physicists, tend to work in high-validity environments with repetitive tasks, often related to physical processes, where feedback on judgement performance is readily available. In contrast, experts with observed poor performance, such as clinical psychologists and stockbrokers, often work in situations where tasks are unique and are related to human behaviour, and where feedback on results is rare, limited, or long-separated from the judgement itself.

According to these criteria, public

search for the best options in a problem space.

To the extent that advice is given or decisions are taken quickly, on partial information, on gut feel or the strength of the narrative case for change, they are likely to be subject to system 1 judgements that are known to be subject to many important biases. It then becomes important for advisors and decision-makers to consider how these biases might be affecting decision-making.

Framing bias and preference reversals

Over two dozen separate cognitive biases have been identified, too many to list here. But to demonstrate the relevance of the literature to policy practice, I will discuss briefly one of the most robust areas of the literature: framing bias and preference reversals.

One of the most important general findings for policy practitioners to consider from the biases and heuristics tradition is that decision-makers' preferences are typically not fixed, and vary depending on how they are elicited (Lichtenstein and Slovic, 2006). A wide range of different studies has shown that two or more options, faced with one's choice of option can change systematically depending on how the options are communicated, even when the content of the options is unchanged. This suggests that an important part of the policy advisor's role is to support decision-makers to construct their preferences, rather than just to generate options for appraisal against a fixed master list of values.

Policy advisors are familiar with the language of framing, but many are perhaps not aware that it is possible to make formal predictions about how certain types of frame will affect decision-making. The classic example of framing effects in cognitive psychology is the Asian disease problem, first introduced by Tversky and Kahneman (1981) and subsequently demonstrated in numerous experiments. In this experiment, subjects are provided identical information and asked to choose between two policy options. Half the participants are randomly assigned to a condition where the outcome information is presented

Policy advisors are familiar with the language of framing, but many are perhaps not aware that we are able to make formal predictions about how certain types of frame will affect decision-making.

argue that expertise or experience is sufficient to protect us. Klein (1998) argues the development of expertise leads to effective intuitive judgement in most cases. Klein has shown that in naturalistic settings, decision-making is typically dominated by intuitive, system 1 judgements, particularly under time pressure, and that in many situations this naturalistic decision-making provides good results, as seen among firefighters, nurses and military commanders.

Against this, judgemental biases have also been observed in certain domain-specific areas of knowledge, suggesting that specialist knowledge does not necessarily protect against thinking biases. For example, doctors have been shown to be subject to framing effects when considering the risks of different treatment options (McNeil et al., 1982), and researchers with advanced statistical training have been found to display errors of mental prediction that conflict with basic statistical rules (Tversky and Kahneman, 1971).

policy would seem to be an environment where expert intuition is unlikely to produce good judgements on its own, on the part of either advisors or decision-makers, except perhaps in more technical areas such as construction policy. Even if it were possible to develop effective intuitive judgement in the low-validity environments common to public policy, this is unlikely to work in practice because many decision-makers and advisors will be exposed to a specific policy environment only briefly, will work on relatively few policy decisions each year, and will often not be sure what effect a policy has had, particularly where there is no follow-up evaluation. This is somewhat concerning, since most policy decisions can be characterised as taken on at least a semi-intuitive, narrative basis. Since ministers typically take decisions on a greatly reduced or simplified set of information – a ten-page limit, or about 3-4,000 words, is common for policy advice – it is hard to argue that decision-makers make a comprehensive, rational

either in terms of deaths incurred or lives saved:

Introductory information (both conditions)

Imagine that the US is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programmes to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programmes are as follows

First condition (frame)

Programme A: If programme A is adopted, 200 people will be saved.

Programme B: If programme B is adopted, there is a 1/3 probability that 600 people will be saved, and 2/3 probability that no people will be saved.

Second condition (frame)

Programme A: If programme A is adopted 400 people will die.

Programme B: If programme B is adopted there is a 1/3 probability that nobody will die, and 2/3 probability that 600 people will die.

Under the first condition, the majority of subjects typically prefer programme A (are risk-averse), whereas under the second condition the majority typically prefer programme B (are risk-seeking). This particular type of framing bias is recognised as a result of prospect theory, according to which people are more sensitive to losses than to gains, and are typically risk-averse in the domain of gains and risk-seeking in the domain of losses (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992). Prospect theory helps to explain status quo bias, whereby people are more sensitive to the losses associated with change than they are to the potential gains – typically twice as sensitive (Samuelson and Zeckhauser, 1998).

Status quo bias occurs when the status quo is implicitly or explicitly used as the reference point against which changes are measured. In some cases, reference points other than the status quo may be

appropriate to consider. For example, as Sunstein (2002) observes:

In environmental regulation, it is possible to manipulate the reference point by insisting that policymakers are trying to restore water or air quality to its state at time X; the restoration time matters a great deal to people's choices. (Sunstein, 2002, p.221)

It is not at all clear in this example or any other whether the status quo reference point or the alternative reference point is normatively superior in any objective sense. The key thing for non-partisan civil servants to observe is that the choice of reference point can affect choices, and

Our inborn tendency is to generalise that something is good or bad, rather than a complex mixture of the two.

so the selection of a reference point has an ethical component.

Prospect theory also provides a useful perspective on trial initiatives, as a way of shifting the perceived status quo (or reference point) in increments, of blunting the psychological impact of potential losses by allowing for the possibility that they can be reversed.

The area of risk is one that appears particularly sensitive to framing effects. Because of the way the associative machinery of system 1 operates, it is difficult for us to maintain conflicting ideas about a policy or anything else. Our inborn tendency is to generalise that something is good or bad, rather than a complex mixture of the two. In practice, this means that it is cognitively difficult to fully accept the risks associated with a favoured policy. Stanovich and West (2008) described this as the non-separability of risk and benefit judgements. In a different context this phenomenon is also known as the halo effect.

Along these same lines, Shafir (1993) found that the way people are asked to evaluate options can affect how they evaluate them and can lead to preference

reversals. Where people are asked to choose from a set of options they typically think about the positive features of the options. Where people are asked to *reject* options, they focus on the negative features of the options. So, an option with strong positive and negative features can be both preferred and rejected over a more average option depending on how the choice is made.

The precise way risk is communicated also appears to be important. For example, Slovic et al. (2000) and Dieckmann, Slovic and Peters (2009) have documented how any expression of risk in terms of relative frequencies (such as framing a risk of cancer in terms of one person in a million) can raise the perceived risk

by inducing affect-laden images of one or more people suffering, particularly for less numerate people. Expressing the same risk as a probability (.000001) leads to a lower perceived risk. Further, less numerate people appear more likely to rely on the narrative evidence accompanying a numerical estimate of likelihood, and are more likely generally to interpret risk as high, even when the objective probability is very low.

Another form of framing effect comes from the inclusion of decoy options into a choice set. Ariely and Wallsten (1995) showed that where two options are quite dissimilar, introducing a decoy option that is similar to one of the options but clearly inferior to it biases decisions towards the option that is linked to the decoy. Similarly, Sunstein (2002) has observed a general tendency to extremeness aversion. With two options, say a small and a medium option, introducing a third, 'large' option can bias decisions towards the medium option.

The way in which options are considered also seems to be important. Hsee et al. (1999) review a range of findings where choices can differ

systematically depending on whether an option is considered in isolation, in series with another option, or side by side with the alternative. Hsee et al. offer an evaluability hypothesis, whereby preference reversals across evaluation mode are particularly likely where an important criterion is difficult to evaluate in isolation. They quote, for example, Desvovages et al. (1992), who find that when options are considered in isolation, subjects show no greater willingness to pay to protect 20,000 endangered birds over 2,000 or even 200 endangered birds.

It also appears important whether an option is considered from within its own domain of value or in a wider context. Kahneman and Ritov (1994) quizzed members of the public and

Implications of framing bias and preference reversals for policy advice

There are several important implications here, most of which are likely to be immediately apparent to practitioners. Perhaps the most important is to remember that advisors' own thinking and preferences are likely to be biased in many situations, and as professionals we have a duty to guard against them and practice cognitive humility.

Another important implication is that most ways of presenting information to decision-makers have the potential to subtly bias decision-making one way or another, and often in predictable ways. In one's own work or in offering second opinion advice or peer review, knowledge of these effects will help advisors

... it appears important that advisors are careful in the way they communicate risk, to ensure the risks are well understood without leading to their cognitive exaggeration.

found preference reversals among a series of comparisons between environmental outcomes and human outcomes, such as protecting spotted owls versus improving earthquake safety. Under single evaluation, subjects tended to state a higher willingness to pay for the environmental outcome, but when choosing between the two in joint evaluation they tended to prefer the human outcome.

A related type of preference reversal is known as the 'less is more' bias. Slovic et al. (2002) and Stanovich and West (2008) found that subjects, considering options in isolation, rated more highly an option that would save 98% of 150 lives than an option that saved 150 lives. This is partly explained by the affect heuristic, whereby 98% of a good thing sounds good and in itself creates positive affect that can bias decision-making. It is also another illustration of the evaluability hypothesis, and the difficulty of independently assessing the value of '150 lives' without any kind of comparator.

understand the influence of one's advisory choices, and to acknowledge one's ethical role in the co-production of government policy.

There is clearly the potential for these phenomena to be used to push a particular agenda. Indeed, it is possible to argue that all advice is intrinsically biased in one way or another, consciously or unconsciously. But we do not necessarily need to yield to full relativism.

Payne, Bettman and Schkade (2006) suggest that an effective decision analyst will support decision-makers to construct preferences that are robust to manipulation, by explicitly offering multiple perspectives and different frames wherever possible.

More generally, it appears that the benchmark or standard chosen for comparison of policy proposals is very important. For example, decisions are likely to be influenced by how and whether policy advisors:

- compare options to the status quo or a different reference point;

- provide qualitative or quantitative assessment of trade-offs;
- provide alternative options to the proposal, including more positive options;
- compare the policy proposal to others in different domains of public value.

A final direct implication is about the importance of clearly communicating risks or downsides of policy options to support sound decision-making. There appears to be an underlying cognitive bias, for both advisor and decision-maker, towards overlooking or downplaying the downsides of an option that is favoured. One option to deal with this problem may be to follow a strategy that leads to the rejection, rather than selection, of options, as per Shafir's (1993) findings that this can lead to a greater focus on downsides. At the same time, it appears important that advisors be careful in the way they communicate risk, to ensure the risks are well understood without leading to their cognitive exaggeration.

Further applications of cognitive psychology to the policy process

In this article I have illustrated only some of the most obvious applications of cognitive psychology to policy practice. Other applications could include:

- The importance of the availability bias (overemphasising salient events and issues in analysis) and affect bias (over-reliance on emotive affect) for agenda setting and strategy, such as through the briefing to the incoming minister process.
- Drawing upon a wider range of formal tools for systematically evaluating options, such as:
 - elimination by aspects (Tversky, 1972) – an approach to choosing between options whereby options are eliminated if they fall below a threshold on the most important attribute, then on the second-most important attribute, and so on;
 - the Delphi method – a structured system for collating individual forecasts or predictions into a group consensus forecast;
 - prediction markets (Surowiecki, 2004) – markets where people

trade 'shares' that pay out if a given outcome comes true, for example www.ipredict.co.nz.

- The importance of training in basic analytical concepts where misunderstanding can lead to analytical biases – such as regression to the mean and the rules of conjunctive probability, both of which have been shown to be poorly understood by many people. (Kahneman, 2011)

Within the traditional eightfold path of policy making (Bardach, 2000), the 'tell your story' phase is perhaps the most important part of the process to consider from the perspective of biases and heuristics, because it entails a simplification and perhaps a shift to faster, system 1 thinking, with greater risk of bias. Even the most rational, exhaustive analysis will not necessarily lead to rational decision-making if distilled to the simplicity of an elevator conversation, as is commonly required of policy advisors.

Overall, the major general lesson for public policy from the biases and heuristics tradition may be the finding that we are all hardwired to respond well to simple causal stories because of the way system 1 operates. The problem with this is that simple, convincing stories can be misleading: 'Paradoxically, it is easier to construct a coherent story when you know little, when there are fewer pieces to fit into the puzzle' (Kahneman, 2011, p.201). The cognitive seductiveness of a simple story should give practitioners pause when we consider the recent proliferation of advice offered orally or via one-page briefings. So too should the common practice of providing only one option for consideration by Cabinet, or one full option with one or two thinly-described straw man alternatives, particularly where there is no regulatory impact statement in support, or when the Cabinet paper is used as a communications tool, either explicitly or implicitly, in anticipation of release under the Official Information Act.

There is a tension between advice that is considered good because it tells a simple, compelling story and advice that is good by some other criteria because it adequately communicates all

complicating and difficult information to decision-makers. Simple advice is appealing for many reasons, but is perhaps more likely to activate system 1 judgements and is, I would argue, more susceptible to cognitive biases. Advice that is robust to framing and other biases is likely to be substantially more complex and place greater demands on decision-makers, with more emphasis on effortful use of system 2.

We also need to acknowledge that this will not always be possible. Research into management styles (Tetlock, 2000) reminds us that some decision-makers are not very concerned by potential biases: Tetlock found that many organisational managers defend simple, heuristic-based errors and prefer simple, decisive leadership styles that reduce the

relying on individuals to tend to their own biases is unlikely to be an effective strategy. The procedures used during the policy development and advice process are likely to be important in either exaggerating or mitigating thinking biases.

The greatest risk of biased analysis and advice is perhaps when the development and advice process occurs with a particular policy in mind. In this case, the confirmation bias can trigger a range of other thinking biases, leading analysts to unconsciously focus on evidence that supports the intended policy, to frame the problem in a way that supports the change, to communicate the policy in a favourable way, to oversimplify the policy, and to underplay risks or trade-offs. Social norms and group dynamics can be particularly powerful impediments to

But the language and concepts of cognitive science are not often encountered explicitly, and are not typically part of the formal training and development of advisors.

information load on top management and avoid unnecessary argumentation.

Where constraints of time or space mean advisors are denied the luxury of offering advice from a range of frames or perspectives, perhaps the lesson is to acknowledge the risk of bias and communicate that clearly, and to ensure that the analysis and lower-level discussion and debate that underpins the final advice has been sufficiently robust that the agency offering the advice has a good understanding of its limitations and potential biases. Related to this is the need to carefully examine the language used to summarise concepts, as certain word choices have strong connotations.

Procedural approaches to considering cognitive biases

Awareness of cognitive biases may be insufficient to address them. As West, Meserve and Stanovich (2012) note, people tend to be much better at identifying biases in others than in themselves, so

overcoming individual-level biases where confirmation bias is at work.

A policy development and advice process that is robust to thinking biases would ideally focus on trying to disconfirm proposed options, in the same way science proceeds by generating hypotheses then seeking ways to disprove them. By deliberately considering ways in which a proposed policy might fail or generate adverse effects, we will be more likely to identify flaws in design, recognise how framing effects might be leading to inflated expectations of success, and gain an accurate understanding of risks.

Procedural options to support a disconfirming strategy include:

- the pre-mortem, a deliberate group task which asks decision-makers or advisors to imagine that a policy decision has been implemented and has failed, and to write down all the reasons why failure might have occurred (Klein, 1998);

- wider collaboration and consultation with individuals with a different point of view, including those who are likely to oppose a proposed policy and who have an incentive to identify flaws in the argument for it (for example, the external panel used by Treasury in the development of its long-term fiscal statement);
- quality assurance processes containing a checklist of some of the most important biases: Gawande (2010) has recently made an eloquent case for the importance of simple checklists to improve the performance of even highly-skilled professionals such as surgeons, and checklists have been used to great effect in aviation to improve the performance of pilots.

Procedural checks are perhaps most important for reactive, time-pressured policy, where biases are more likely. Where rapid or reactive advice is needed, a simple procedural check to alert decision-

makers to potential biases could be a signed statement by the author, covering how extensively they have examined any relevant literature, how many options have been fully developed, and so forth. This could provide reason to pause for decision-makers, perhaps buying time for more considered policy advice with less risk of biased analysis.

Conclusion

Effective practitioners will already intuitively understand much of what I have discussed in this article. But the language and concepts of cognitive science are not often encountered explicitly, and are not typically part of the formal training and development of advisors. For some advisors and decision-makers, the first reaction to the biases and heuristics literature may be that it is important for the design of policies, but has less relevance for the thinking processes of advisors and decision-makers. But the evidence to date suggests that policy makers may be just as

subject to biases in their thinking as policy takers.

Overall, it appears that much of the cognitive psychology literatures are developed and widely accepted enough now that formal training and a new language could usefully be included in practitioners' toolkits, and that as a professional community policy advisors, public officers and decision-makers would do well to reflect on the way we practice our craft in light of these findings.

¹ This could be seen as an example of the attribute substitution heuristic, whereby the mind substitutes an easy problem ($\$1.10 - \1) for a somewhat harder problem ($x + y = \$1.10, y - x = \1 , solve for x).

Acknowledgements

Thanks to Nick Carroll, Tony Burton, Jason Raven, Aphra Green and Brook Barrington for comments on earlier drafts. Mistakes are mine alone.

References

- Aralal, E., S. Fritzen, M. Howlett, M. Ramesh and X. Wu (2013) *Routledge Handbook of Public Policy*, New York: Routledge
- Ariely, D. (2008) *Predictably Irrational*, New York: HarperCollins
- Ariely, D. and T. Wallsten (1995) 'Seeking subjective dominance in multidimensional space: an explanation of the asymmetric dominance effect', *Organisational Behavior and Human Decision Processes*, 63 (3), pp.223-32
- Bardach, E. (2000) *A Practical Guide for Policy Analysis: the eightfold path to more effective problem solving*, New York: Chatham House
- Chaikan, S. and Y. Trope (1999) *Dual-process Theories in Social Psychology*, New York: Guilford Press
- Desvougues, W., R. Johnson, R. Dunford, K. Boyle, S. Hudson and K. Wilson (1992) *Measuring Non-use Damages Using Contingent Valuation: an experimental evaluation accuracy*, Research Triangle Institute monograph 92-1, Research Triangle Institute Park, NC: RTI Press
- Dieckmann, N.F., P. Slovic and E.M. Peters (2009) 'The use of narrative evidence and explicit likelihood by decision makers varying in numeracy', *Risk Analysis*, 29, pp.1473-88
- Dolan, P., M. Hallsworth, D. Halpern, D. King and I. Vlaev (2010) *MindSpace: influencing behaviour through public policy*, London: Institute for Government
- Frederick, S. (2005) 'Cognitive reflection and decision making', *Journal of Economic Perspectives*, 19 (4), pp.25-42
- Gawande, A. (2010) *The Checklist Manifesto*, New York: Henry Holt
- Gigerenzer, G., U. Hoffrage and H. Kleinbolting (1991) 'Probabilistic mental models: a Brunswickian theory of confidence', *Psychological Review*, 98 (4), pp.506-28
- Gladwell, M. (2005) *Blink*, New York: Little, Brown
- Hsee, C., G. Loewenstein, S. Blount and M. Bazerman (1999) 'Preference reversals between joint and separate evaluations of options: a review and theoretical analysis', *Psychological Bulletin*, 125 (5), pp.576-90
- Jones, B., G. Boushey and S. Workman (2006) 'Behavioral rationality and the policy processes: toward a new model of organisational information processing', in B. Peters and J. Pierre (eds), *Handbook of Public Policy*, London: Sage Publications
- Kahneman, D. (2011) *Thinking Fast and Slow*, New York: Farrar, Straus and Giroux
- Kahneman, D. and I. Ritov (1994) 'Determinants of stated willingness to pay for public goods: a study in the headline method', *Journal of Risk and Uncertainty*, 9, pp.5-28
- Kahneman, D. and A. Tversky (1979) 'Prospect theory: an analysis of decision under risk', *Econometrica*, 47, pp.263-91
- Klein, G. (1998) *Sources of Power: how people make decisions*, Cambridge, MA: MIT Press
- LeBoeuf, R., E. Shafir and J. Bayak (2010) 'The conflicting choices of alternating selves', *Organisational Behavior and Human Decision Processes*, 111, pp.48-61
- Lichtenstein, S. and P. Slovic (2006) *The Construction of Preference*, Cambridge: Cambridge University Press

- Maclean, D. (2002) 'Some morals of a theory of nonrational choice', in R. Gowda and J. Fox (eds), *Judgements, Decisions and Public Policy*, Cambridge: Cambridge University Press
- McNeil, B., G. Pauker, H. Sox and A. Tversky (1982) 'On the elicitation of preferences for alternative therapies', *New England Journal of Medicine*, 306, pp.1259-62
- Ministry of Economic Development (2006) *Behavioural Analysis for Policy: new lessons from economics, philosophy, psychology, cognitive science, and sociology*, Wellington: Ministry of Economic Development
- Payne, J., J. Bettman and D. Schkade (1999) 'Measuring constructed preferences: towards a building code', *Journal of Risk and Uncertainty*, 19, pp.243-70
- Samuelson, W. and R. Zeckhauser (1998) 'Status quo bias', *Journal of Risk and Uncertainty*, 1, pp.7-59
- Shafir, E. (1993) 'Choosing versus rejecting: why some options are both better and worse than others', *Memory and Cognition*, 21 (4), pp.546-56
- Shafir, E. (2002) 'Cognition, intuition and policy guidelines', in R. Gowda and J. Fox (eds), *Judgements, Decisions and Public Policy*, Cambridge: Cambridge University Press
- Shanteau, J. (1992) 'Competence in experts: the role of task characteristics', *Organisational Behavior and Human Decision Processes*, 53, pp.252-66
- Silver, N. (2012) *The Signal and the Noise*, New York: Penguin
- Simon, H. (1956) 'Rational choice and the structure of the environment', *Psychological Review*, 63, pp.129-38
- Slovic, P., M. Finucane, E. Peters and D. MacGregor (2002) 'The affect heuristic', in T. Gilovich, D. Kahneman and D. Griffin (eds), *Heuristics and Biases: the psychology of intuitive judgement*, Cambridge: Cambridge University Press
- Slovic, P., J. Monahan and D. MacGregor (2000) 'Violence risk assessment and risk communication: the effects of using actual cases, providing instructions, and employing probability vs. frequency formats', *Law and Human Behavior*, 24, pp.271-96
- Stanovich, K. and R. West (2008) 'On the relative independence of thinking biases and cognitive ability', *Journal of Personality and Social Psychology*, 94 (4), pp.672-95
- Sunstein, C. (2002) 'Toward behavioural law and economics', in R. Gowda and J. Fox (eds), *Judgements, Decisions and Public Policy*, Cambridge: Cambridge University Press
- Surowiecki, J. (2004) *The Wisdom of Crowds*, New York: Little, Brown
- Swoyer, S. (2002) 'Judgement and decision making: extrapolations and applications', in R. Gowda and J. Fox (eds), *Judgements, Decisions and Public Policy*, Cambridge: Cambridge University Press
- Taleb, N. (2007) *The Black Swan*, New York: Random House
- Tetlock, P. (2000) 'Cognitive biases and organizational corrections: do both disease and cure depend on the politics of the beholder?', *Administrative Science Quarterly*, 45 (2), pp.293-326
- Tetlock, P. (2005) *Expert Political Judgement: How good is it? How can we know?*, Princeton: Princeton University Press
- Thaler, R. and C. Sunstein, *Nudge*, New Haven: Yale University Press
- Thaler, R., C. Sunstein and J. Balz (2010) *Choice Architecture*, available at SSRN: <http://ssrn.com/abstract=1583509> or <http://dx.doi.org/10.2139/ssrn.1583509>
- Toplak, M., R. West and K. Stanovich (2011) 'The Cognitive Reflection Test as a predictor of performance on heuristics-and-biases tasks', *Memory and Cognition*, 39, pp.1275-89
- Tversky, A. and D. Kahneman (1992) 'Advances in prospect theory: cumulative representation of uncertainty', *Journal of Risk and Uncertainty*, 5, pp.297-323
- Tversky, A. and D. Kahneman (1971) 'Belief in the law of small numbers', *Psychological Bulletin*, 76 (2), pp.105-10
- Tversky, A. and D. Kahneman (1974) 'Judgement under uncertainty: heuristics and biases', *Science*, 185 (4157), pp.1124-31
- Tversky, A. and D. Kahneman (1981) 'The framing of decisions and the psychology of choice', *Science*, 211 (4481), pp.453-8
- Tversky, A. (1972) 'Elimination by aspects: a theory of choice', *Psychological Review*, 79, pp.281-99
- West, R., R. Meserve and K. Stanovich (2012) 'Cognitive sophistication does not attenuate the bias blind spot', *Journal of Personality and Social Psychology*, 103 (3), pp.506-19