



Miserliness in human cognition: the interaction of detection, override and mindware

Keith E. Stanovich

Department of Applied Psychology and Human Development, University of Toronto, Toronto, Canada

ABSTRACT

Humans are cognitive misers because their basic tendency is to default to processing mechanisms of low computational expense. Such a tendency leads to suboptimal outcomes in certain types of hostile environments. The theoretical inferences made from correct and incorrect responding on heuristics and biases tasks have been overly simplified, however. The framework developed here traces the complexities inherent in these tasks by identifying five processing states that are possible in most heuristics and biases tasks. The framework also identifies three possible processing defects: inadequately learned mindware; failure to detect the necessity of overriding the miserly response; and failure to sustain the override process once initiated. An important insight gained from using the framework is that degree of mindware instantiation is strongly related to the probability of successful detection and override. Thus, errors on such tasks cannot be unambiguously attributed to miserly processing – and correct responses are not necessarily the result of computationally expensive cognition.

ARTICLE HISTORY Received 27 December 2017; Accepted 19 March 2018

KEYWORDS Heuristics and biases; miserly information processing; dual process theory

Introduction

That humans are cognitive misers has been a major theme throughout the past 50 years of research in psychology and cognitive science (see Dawes, 1976; Kahneman, 2011; Shah & Oppenheimer, 2008; Simon, 1955, 1956; Taylor, 1981; Tversky & Kahneman, 1974). When approaching any problem, our brains have available various computational mechanisms for dealing with the situation. These mechanisms embody a trade-off, however (Rand, Tomlin, Bear, Ludvig, & Cohen, 2018; Tomlin, Rand, Ludvig, & Cohen, 2015). The trade-off is between power and expense. Some mechanisms have great computational power – they can solve a large number of novel problems. However, these mechanisms take up a great deal of attention, tend to be slow, tend to

interfere with other thoughts and actions we are carrying out, and they require great concentration that is often experienced as aversive (Kahneman, 1973; Kurzban, Duckworth, Kable, & Myers, 2013; Navon, 1989; Westbrook & Braver, 2015).

Humans are cognitive misers because their basic tendency is to default to processing mechanisms of low computational expense. Humorously, Hull (2001) has said that “the rule that human beings seem to follow is to engage the brain only when all else fails—and usually not even then” (p. 37). More seriously, Richerson and Boyd (2005) have put the same point in terms of its origins in evolution: “In effect, all animals are under stringent selection pressure to be as stupid as they can get away with” (p. 135). Miserly cognitive tendencies have evolved for reasons of computational efficiency. But that same computational efficiency simultaneously guarantees that humans will be less than perfectly rational.

Miserly processing and human evolution

Of course, evolution guarantees human rationality in the dictionary sense of “the quality or state of being able to reason” because evolution built the human brain. What is meant here is that evolution does not guarantee perfect rationality in the sense the term is used throughout cognitive science – as maximising subjective expected utility. In contrast to maximisation, natural selection works on a “better than” principle. As Dawkins (1982) puts it:

Natural selection chooses the better of present available alternatives... The animal that results is not the most perfect design conceivable... It is the product of a historical sequence of changes, each one of which represented, at best, the better of the alternatives that happened to be around at the time. (p. 46)

The variation and selective retention logic of evolution “designs” for the reproductive advantage of one organism over the next, not for the optimality of any one characteristic (including rationality). Natural selection is geared to immediate advantage rather than long-term strategy. Human rationality, in contrast, must incorporate the long-term interests of the individual and thus it can diverge from the short-term strategies of evolutionary adaptation (Ainslie, 2001; de Sousa, 2007; Loewenstein, 1996; Nozick, 1993; Stanovich, 2004).

Organisms have evolved to increase the reproductive fitness of genes, not to increase the rationality of humans, and increases in fitness do not always entail increases in rationality. For example, beliefs need not always track the world with maximum accuracy in order for fitness to increase. Evolution might fail to select out epistemic mechanisms of high accuracy when they are costly in terms of organismic resources (for example, in terms of memory, energy or attention). A second reason that belief-forming mechanisms might not be maximally truth preserving is that:

a very cautious, risk-averse inferential strategy—one that leaps to the conclusion that danger is present on very slight evidence—will typically lead to false beliefs more often, and true ones less often, than a less hair-trigger one that waits for more evidence before rendering a judgment. Nonetheless, the unreliable, error-prone, risk-averse strategy may well be favored by natural selection. For natural selection does not care about truth; it cares only about reproductive success. (Stich, 1990, p. 62)

It is likewise in the domain of goals and desires. The purpose of evolution was not to maximise the happiness of human beings. As has become clear from research on affective forecasting (Gilbert, 2006; Kahneman, 2011), people are remarkably bad at making choices that make themselves happy. This should be no surprise. The reason we have pleasure circuits in our brains is to encourage us to do things (survive and reproduce, help kin) that propagate our genes. The pleasure centres were not designed to maximise the amount of time we are happy.

The instrumental rationality of humans is not guaranteed by evolution for two further reasons. First, many genetic goals that have been lodged in our brain no longer serve our ends because the environment has changed. For example, our mechanisms for storing and utilising energy evolved in times when fat preservation was efficacious. These mechanisms no longer serve the goals of people in our modern technological society where there is a McDonald's on practically every corner. The goals underlying these mechanisms have become detached from their evolutionary context (Li, van Vugt, & Colarelli, 2018). Finally, the cultural evolution of rational standards is apt to occur at a pace markedly faster than that of human evolution (Richerson & Boyd, 2005; Stanovich, 2004) – thus providing ample opportunity for mechanisms of utility maximisation to dissociate from local genetic fitness maximisation.

Thus, our evolutionary history does not guarantee that all of the miserly processing defaults in our brain result in optimal processing. Many of these evolutionary defaults of the cognitive miser were “good enough” in their day (our environment of evolutionary adaptation of thousands of years ago), but might not be serving us well now when our environments have radically changed. They were not designed for the type of situation in modern society that call for fine-grained analysis such as financial decisions, fairness judgements, employment decisions, legal judgements, etc.

Miserly processing in benign and hostile environments

For the reasons just outlined, it is clear why miserly processing ensures that at least some non-optimal behaviour will be produced. None of this is meant to deny that, on a token basis, miserly processing is astonishingly useful. Nevertheless, the benefits and costs of miserly processing depend on the nature of

the environment. For maximum effectiveness, miserly processing requires benign environments that contain cues that elicit practiced, adaptive behaviours. In hostile environments, however, defaulting to miserly processing can be costly. A benign environment is one that contains useful (that is, diagnostic) cues that can be exploited by various heuristics (for example, affect-triggering cues, vivid and salient stimulus components, convenient and accurate anchors). Additionally, for an environment to be classified as benign, it must also contain no other individuals who will adjust their behaviour to exploit those relying on miserly processing.

In contrast, a hostile environment is one in which there are few cues that are usable by fast-acting autonomous processes or one that presents misleading cues (Kahneman & Klein, 2009). Another way that an environment can turn hostile occurs when other agents discern the simple cues that are being used to trigger miserly defaults and arrange them for their own advantage (for example, advertisements, or the strategic design of supermarket floor space in order to maximise revenue).

Thus, we only want to engage in miserly processing when in benign environments. Our responses might be suboptimal when we default to miserly processing in a hostile environment. As Kahneman and Frederick (2002) have argued, humans often act as cognitive misers by engaging in attribute substitution – the substitution of an easy-to-evaluate characteristic for a harder one, even if the easier one is less accurate. For example, the cognitive miser will substitute the less effortful attributes of vividness or affect for the more effortful retrieval of relevant facts (Kahneman, 2003; Li & Chapman, 2009; Slovic & Slovic, 2015; Wang, 2009). But when we are evaluating important financial decisions – buying a home or obtaining a mortgage or insurance – we do not want to substitute vividness for careful thought about the situation. Modern society keeps proliferating such situations where shallow processing is not sufficient for maximising personal happiness. These situations are increasing in frequency because many structures of market-based societies have been designed explicitly to exploit miserly tendencies. Being cognitive misers in these situations will seriously impede people from achieving their goals (see Stanovich, 2004).

Heuristics and biases tasks were designed to assess costly miserly defaults in hostile worlds

The primary tasks used to assess miserly processing have been drawn from the so-called heuristics and biases tradition inaugurated by Kahneman and Tversky in the early 1970s (Kahneman & Tversky, 1972, 1973; Tversky & Kahneman, 1974). The term biases refers to the defaults that lead people to make systematic errors in choosing actions and in estimating probabilities. The

term heuristic refers to why people often make these errors: they use mental shortcuts (heuristics) to solve many problems. In short, the tasks developed in this literature were designed to tap the tendency towards miserly information processing. It is especially important to realise, though, that this assessment of miserly processing takes place in the context of a deliberately contrived hostile environment.

That heuristics and biases tasks are deliberately designed to be hostile environments for the miserly processor is one of the things that make them psychologically interesting and highly diagnostic of the dynamics of human reasoning. It is a key feature that makes heuristics and biases task items different from typical IQ test items which present much more benign task environments. For example, intelligence tests contain salient warnings that more than miserly processing is necessary. It is clear to someone taking an intelligence test that fast, automatic, intuitive processing will not lead to superior performance. Most heuristics and biases tasks do not strongly cue the subject in this manner. Instead, many such tasks suggest a compelling intuitive response that happens to be wrong. In heuristics and biases tasks, unlike the case for intelligence tests, the subject must detect the inadequacy of the intuitive response that is triggered with little effort. They must then suppress this response while selecting a better alternative.

It is not only the presence of an enticing lure response that distinguishes many heuristics and biases tasks from IQ test items. Additionally, the demand characteristics of IQ tests pretty much guarantee that the subject will be working at maximum efficiency on a problem that is unambiguously framed by its instructions. This is unlike heuristics and biases tasks where the subject must often choose a particular construal. The fact that many heuristics and biases tasks can be construed by the subject in different ways (a statistical interpretation versus a narrative interpretation, for instance) is often seen as a weakness of such tasks when in fact it is the design feature that makes the task diagnostic. In a probabilistic reasoning task from this literature, the entire point is to see how dominant or non-dominant the statistical interpretation is over the narrative interpretation. Likewise, the fact that many such problems have an intuitively compelling wrong answer is often seen as a misleading attempt to “trick” the participant. In fact, the presence of the compelling intuitive response is precisely what makes the problem diagnostic of the propensity to avoid miserly processing.

It is in this sense that the so-called artificiality of heuristics and biases tasks is a strength and not a weakness. It is a design feature, not a bug, because the modern world is, in many ways, becoming hostile for individuals relying solely on miserly processing. Einhorn and Hogarth (1981) long ago made the telling point that: “in a rapidly changing world it is unclear what the relevant natural ecology will be. Thus, although the

laboratory may be an unfamiliar environment, lack of ability to perform well in unfamiliar situations takes on added importance” (p. 82). Their point is that it is wrong to imply that because heuristics and tasks are abstract and not like “real life” then we need not worry that people do poorly on them. The issue is that, ironically, the argument that these laboratory tasks are not like “real life” is becoming less and less true. “Life”, in fact, is becoming more like the tasks!

Of course, it should not be inferred that miserly processing always leads us astray. As previously discussed, heuristics often give us a useful first approximation to the optimal response in a given situation, and they do so without stressing cognitive capacity. But we must not lose sight of the fact that the usefulness of the heuristics that we rely upon to lighten the cognitive load are dependent on a benign environment. In contrast, the environments of modern life in a technological society are often not benign. In modern life, we often must: decide which health maintenance organisation to join based on abstract statistics rather than experienced frequencies; decide on what type of mortgage to purchase; figure out what type of deductible to get on our auto insurance; decide whether to trade in a car or sell it ourselves; decide whether to lease or to buy; and think about how to apportion our retirement funds. These are just a few of the plethora of modern-day decisions and choices that are best made by *not* being miserly in our information processing.

The importance of knowledge structures (mindware) in understanding heuristics and biases tasks

Miserly processing, as well as heuristics and biases tasks, are much discussed in the literature on dual-process theories (Kahneman, 2011), but it is important to understand that the concept of miserly information processing is not necessarily tied to such models. As a processing characteristic, it will emerge as an issue whatever cognitive architecture is proposed (Arkes, 2016; De Neys, 2018; Newman, Gibb, & Thompson, 2017; Oaksford & Chater, 2014; Reyna & Brainerd, 2011; Sun, 2015). Likewise, whatever the architecture adopted, knowledge structures will be implicated in task performance. That is, to successfully avoid miserly processing in hostile situations, both procedural skills and declarative knowledge are required. For example, Kahneman and Frederick’s (2002) concept of attribute substitution was mentioned previously – miserly processing carried out via the substitution of an easy-to-evaluate characteristic for a harder one (even if the easier one is less accurate). If attribute substitution is to be avoided, it is necessary that the inadequacy of the substituted attribute be detected – and the possibility of responding based on a more diagnostic attribute must be recognised. This so-called detection process¹ must be followed by inhibitory processes that suppress

the intuitive, miserly response. These inhibitory processes are often termed the override mechanism.

Detection and override are the processing components of successful avoidance of attribute substitution. Often ignored in the literature is that successful performance on heuristics and biases tasks is dependent on stored knowledge of various types. These knowledge bases, rules, procedures and strategies have been referred to as *mindware*, a term coined by David Perkins in a 1995 book (Clark (2001) uses the term in a slightly different way from Perkins' original coinage). The mindware necessary to perform well on heuristics and biases tasks is disparate, encompassing knowledge in the domains of probabilistic reasoning, causal reasoning, scientific reasoning and numeracy.

In the study of heuristics and biases tasks, researchers have tended to emphasise their processing requirements and have not paid enough attention to the knowledge component in such tasks. Even more importantly, researchers have failed to realise how the knowledge component of such tasks interacts with the processing components. The critique of the heuristics and biases literature that follows is not meant to deny the enormous progress that has been made in understanding these tasks in the last 10 years (De Neys, 2014, 2018; De Neys & Glumicic, 2008; Evans, 2007, 2008, 2010, 2014; Ferreira, Mata, Dokin, Sherman, & Ihmels, 2016; Pennycook, Fugelsang, & Koehler, 2015; Stanovich, 2011; Stanovich & West, 2008; Thompson, 2014; Thompson & Johnson, 2014; Thompson & Morsanyi, 2012). For example, the important distinction between detection and override has been explored both theoretically and empirically in a way that makes clear how the earlier literature was mistakenly fusing these two critical concepts (De Neys & Glumicic, 2008; Stanovich & West, 2008). Nevertheless, this theoretical attention to distinguishing detection from override – itself of substantial scientific importance – may have contributed to a particular skewed view of heuristics and biases tasks: that they are *pure* indicators of miserly processing.

This assumption that the tasks involved are pure indicators of miserly processing has been a theoretically confusing aspect of the heuristics and biases literature. It has led to the widespread tendency in the literature to treat such tasks as if they do not involve learned mindware. This has been particularly true of discussions of Frederick's (2005) Cognitive Reflection Test (CRT) which

¹For the purposes of this essay, the term detection is used in the most ecumenical sense, in that it is not my purpose here to adjudicate among specific models of conflict assessment and resolution. The issues that I wish to address in this paper apply to both of the major dual-processing architectures that have been discussed in the literature – the default-interventionist architecture and parallel architectures (De Neys, 2012; Evans & Stanovich, 2013; Sloman, 1996). However, the present paper's emphasis on moving towards a more *continuous* concept of mindware presence is most compatible with the more recent default-interventionist architectures that emphasise the presence of multiple, conflicting Type 1 outputs (Bago & De Neys, 2017; De Neys, 2012; Pennycook, Fugelsang, & Koehler, 2015). Such architectures rely on automatic mechanisms that monitor conflict and that become bottom-up triggers of Type 2 processing (see Pennycook et al., 2015; Thompson, 2009, 2014; Thompson & Morsanyi, 2012).

is often treated as if it is a pure measure of miserly processing, even though work for some years now has shown that it is psychometrically complex (involving dispositions and numeracy, as well as cognitive capacity; see Liberali, Reyna, Furlan, Stein, & Pardo, 2012; Sinayev & Peters, 2015; Toplak, West, & Stanovich, 2011, 2014).

Virtually all of the tasks in the heuristics and biases literature involve mindware – some more than others (for a taxonomy, see Stanovich, West, & Toplak, 2016). This means that how well the relevant mindware is instantiated in memory will affect the performance observed on the task. The task will not just be an indicator of miserly processing but will also be an indicator of the depth of learning of the relevant mindware as well. This caution regarding these tasks is not just for the theorist employing dual-process theory. It will also hold for research using other cognitive architectures as frameworks (De Neys, 2018; Newman et al., 2017; Oaksford & Chater, 2014).

As a result of ignoring the role of mindware, many researchers treat heuristics and biases tasks as pure measures of miserly processing when in fact they are complex indicators. This leads to many conceptual errors being introduced into the literature (Evans, 2018). One that has become quite common is the assumption that dual process theories imply that all errors must be fast (the result of miserly processing) and that all correct responses must be slow (Achtziger & Alos-Ferrer, 2014; Jimenez, Rodriguez-Lara, Tyran, & Wengstrom, 2018; Moore, 2017). Figure 1, which employs a dual-process architecture introduced by Stanovich (2011), shows clearly that this inference does not follow. In the upper right of the figure is displayed the case that is over-represented in our literature. In the case represented there, a non-normative response from the autonomous mind (the so-called System 1 processes; see Kahneman, 2011; Stanovich, 1999) has been interrupted and the computationally taxing process of simulating an alternative response is underway. That simulation involves the computationally expensive process of accessing mindware for the simulation.

In contrast to this type of mindware access, indicated in the lower left of Figure 1 is a qualitatively different way that mindware can lead to a normative response. The figure indicates that within System 1 can reside normative rules and rational strategies that have been practiced to automaticity and that can automatically compete with (and often immediately defeat) any alternative non-normative response.

This idea is not new. The category of autonomous processing in cognitive science has long included the automatic triggering of overlearned rules (LaBerge & Samuels, 1974; Moors & De Houwer, 2006; Posner & Snyder, 1975; Shiffrin & Schneider, 1977). That is, automatized subprocesses include not only Fodorian (1983) modules, but in addition include many rules, stimulus discriminations and decision-making principles that have been practised to automaticity (Kahneman & Klein, 2009; Shiffrin & Schneider, 1977). However,

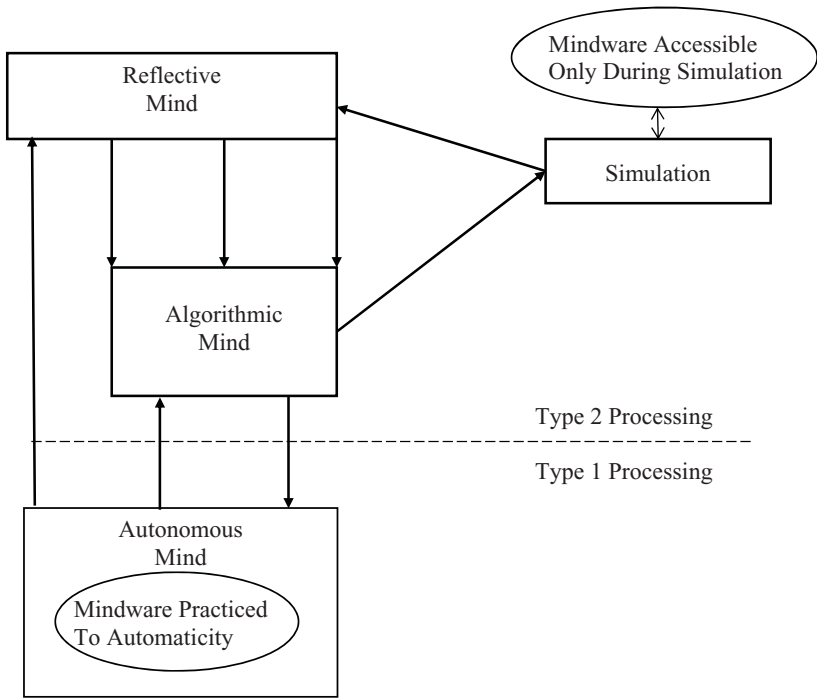


Figure 1. A simplified model showing both automatised mindware and mindware accessible during simulation.

System 1 contains the products of implicit learning as well. Much automatic associative learning becomes highly compiled and autonomously triggered from System 1, in addition to rules that are consciously practised. In short, not all knowledge in System 1 is automated through conscious practice, because some knowledge learned implicitly can become automated as well. Chapter 3 of Evans (2010) contains a discussion of how implicit associative knowledge can become stored in System 1.

The main purpose of Figure 1 is to concretely illustrate the idea that the normative mindware of correct responding is not exclusively retrieved during simulation activities, but can become implicated in performance directly and automatically from System 1 if it has been practiced enough. Some statistics instructors, for example, become unable to empathise with their students for whom the basic probability axioms are not transparent. The instructor can no longer remember when these axioms were not primary intuitions.

The instantiation of the normative mindware in System 1 will affect, for example, the feeling-of-rightness judgement that Thompson (2009, 2014) and Thompson, Prowse Turner, and Pennycook (2011) have shown is a bottom-up monitoring mechanism that has implications for performance on heuristics and biases tasks – especially if understood within the context of recent

default-interventionist architectures that emphasise the presence of multiple, conflicting Type-1 outputs (Bago & De Neys, 2017; Pennycook et al., 2015). The feeling-of-rightness output from the intuitive response might be lowered by the presence of the conflicting normative mindware in System 1, perhaps to a level that would initiate override and decoupling operations (Stanovich & Toplak, 2012). In the extreme, the feeling-of-rightness output of the normative mindware itself might reach a level so high that it is emitted as an automatic response from System 1 (see discussion below).

So it should be clear from Figure 1 that it does not follow from the output of a normative response that Type-2 processes were necessarily the genesis of the correct responding. Thus, it does not follow that a rapid response should necessarily be an incorrect one. Several recent papers have provided empirical evidence that rapid responses are not necessarily incorrect (Bago & De Neys, 2017; Newman et al., 2017; Trippas, Handley, Verde, & Morsanyi, 2016).

Figure 1 illustrates why the assumption that correct responses must necessarily be slow does not follow. Normative mindware of correct responding is not exclusively retrieved during simulation activities, but can become implicated in performance directly and automatically from System 1 if it has been practised enough. As we will see in the next section, the presence of automated mindware in System 1 complicates the interpretation of performance on heuristics and biases tasks.

In the remainder of this article, I will sketch some of the complex interactions between mindware, detection and override in heuristics and biases tasks. Importantly, the conclusion will be that these concepts are more intertwined than typically realised. These tasks have led to many important behavioural insights (Kahneman, 2011; Thaler, 2015) but, from a processing point of view, they are not pure indicators of any precise processing concept (such as miserly processing, for example). They are instead complex entities and they differ greatly from each other. In the sketch of the interdependence of mindware, detection and override that follows, a dual-process perspective will be adopted, but the points illustrated are applicable across a number of different architectures (De Neys, 2018; Newman et al., 2017; Oaksford & Chater, 2014).

The interdependence of mindware, detection and override

It is important to understand that the presence of mindware, detection of the need for override and sustained override capability are concepts that are intertwined in important ways. Here we will consider several such dependencies, each of increasing complexity.

First, the relevant (that is, normatively appropriate) mindware must be available to potentially come into conflict with the miserly response. Conflict detection abilities will not be assessed if the mindware necessary to generate

a conflicting response is not present. Other conflicts may of course still occur (for example, between the intuitive response and inappropriate mindware that is retrieved).

If the relevant mindware is in fact available, then detection of conflict is at least possible. However, even if the relevant mindware is present, if the subject does not detect any reason to override the intuitive response, then sustained override capability will not come into play in this instance. In short, whether or not a task (for a given person) assesses certain “downstream” capabilities (sustained override) depends on whether certain “upstream” capabilities (conflict detection, presence of mindware) are present.

Note in particular the trade-off-type relationship between the failure of override and the absence of mindware. In any taxonomy of errors in heuristics and biases tasks (Stanovich et al., 2016), these two error types² will be related to each other. The two errors are contingent on how well learned the mindware is. Errors made by someone with well-learned mindware are more likely to be due to override failure. Conversely, override errors are less likely to be attributed to people with little, or poorly learned, mindware installed. Of course, the two categories trade off in a continuous manner with an indistinct boundary between them. A well-learned rule not appropriately applied is a case of override failure. As the rule is less and less well instantiated, at some point, it is so poorly compiled that it is not a candidate for retrieval in the override process and thus the override error becomes an error due to a mindware gap. In short, a process error has turned into a knowledge error.

The next couple of figures serve to illustrate the interdependence and complex relationships between mindware presence, detection and sustained override capability. [Figure 2](#) is organised around a continuum reflecting how well the mindware in the relevant problem has been instantiated.³ At the far left of the continuum in [Figure 2](#), the mindware is totally absent. As the relevant mindware becomes practised and stored in long-term memory, it becomes available for retrieval by Type-2 processes. In the middle of the continuum (mindware learned but not fully automatised), the mindware must be retrieved by expensive Type-2 processing (see the upper right of [Figure 1](#)) in order to aid in creating what might be called a computed response to compete with what might be called the intuitive response that is naturally emitted by System 1 processes.

²Our taxonomy (see Stanovich, West, & Toplak, 2016) emphasises on the role of conflict and mindware because of our focus on modelling heuristics and biases tasks, many of which were designed to cause processing conflict. There are other ways to organise a taxonomy of rational thinking tasks. For example, Shah and Oppenheimer (2008) outline a taxonomy in terms of how heuristics bring about effort reduction in various tasks.

³It is critically important to note that all of these figures are task specific for a given subject. The degree of instantiation of mindware will vary from task to task within a subject and from subject to subject within a task.

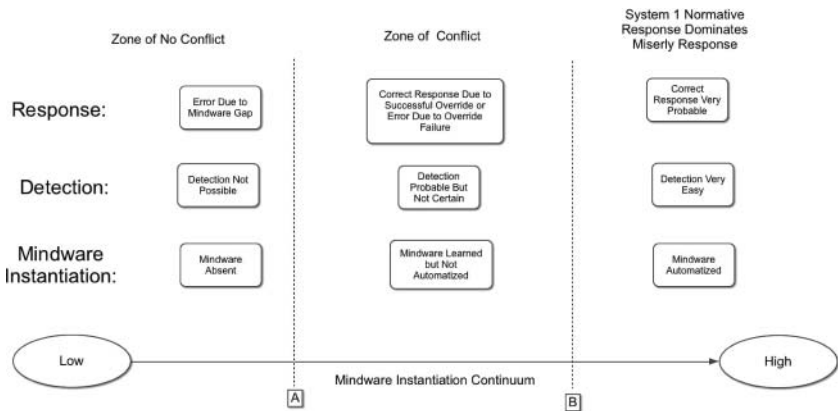


Figure 2. Processing states on the mindware continuum.

On the far right of the continuum (mindware automatised), the relevant mindware has been so overly practised that it has entered System 1 and is triggered automatically and autonomously. That is, it is the *normative* response that trips an automatic monitoring system (perhaps a feeling-of-rightness mechanism of the type discussed by Thompson, 2009, 2014) to such an extent that it can trigger a response. This mindware is so overly practised that it can often automatically trump the intuitive response from System 1 without a taxing override procedure needing to be invoked. In short, the far right of the continuum is the area where no sustained override is needed. Subjects will almost always get the problem correct while in this part of the continuum – when the normative mindware is so well instantiated. This situation contrasts sharply with that on the far left of the continuum. Here, the mindware is so little practised that no conflict detection will occur and the subject will always make an error due to a mindware gap. Override of the miserly response is not possible, because the mindware is not instantiated well enough for the necessity of override to be detected.

The middle section of the figure represents the zone of conflict between System 1 processes priming a miserly response and those priming a normative response. Here, whether the subject responds correctly or not will depend on the success of sustained override. This zone of conflict is defined by the area demarcated by thresholds A and B on the mindware instantiation continuum. To the left of dotted line A, the mindware is not well enough instantiated to prime a superior (that is, normative) response. To the right of dotted line A and to the left of dotted line B, the mindware is instantiated enough to activate some response priming but not enough to automatically trigger a normative response. To the right of dotted line B is the area discussed before, where an automatised normative response trumps the intuitive response.

It is important to note that Figure 2 illustrates that the probability of detection and the degree of mindware instantiation will be highly correlated. This theoretical conjecture is supported by a finding of Frey, Johnson, and De Neys (2018). They found that poor performance on no-conflict versions of heuristics and biases tasks (a measure of mindware instantiation) predicted detection failure on conflict versions of the same task. That the probability of detection and the degree of mindware instantiation will be highly correlated illustrates one insufficiently appreciated aspect of these tasks: *knowledge considerations and processing considerations are very difficult to separate in many of the heuristics and biases tasks that involve conflict.*

Further differentiation of heuristics and biases tasks: five different processing states

Figure 3 presents the logic of heuristics and biases tasks in even more detail. Here, the four letters W through Z mark the criterion values on the mindware continuum that help define five different processing states. Criterion W marks the place on the mindware instantiation continuum where conflicts become possible. To the left of this criterion, the mindware is so poorly learned that it should be considered absent and hence no conflict detection is possible. Conflict detection – that is, detection of a conflict between an intuitive (miserly) response and learned normative rules – is possible to the right of criterion W. However, conflict is only actually detected to the right of criterion X (by this particular subject on this particular task, see footnote 3). This criterion creates an area (to the right of W and to the left of X), where a computed normative response is possible because the mindware is there for retrieval, but no conflict is detected and hence no sustained override is even attempted. This processing state represents what might be called a detection error.

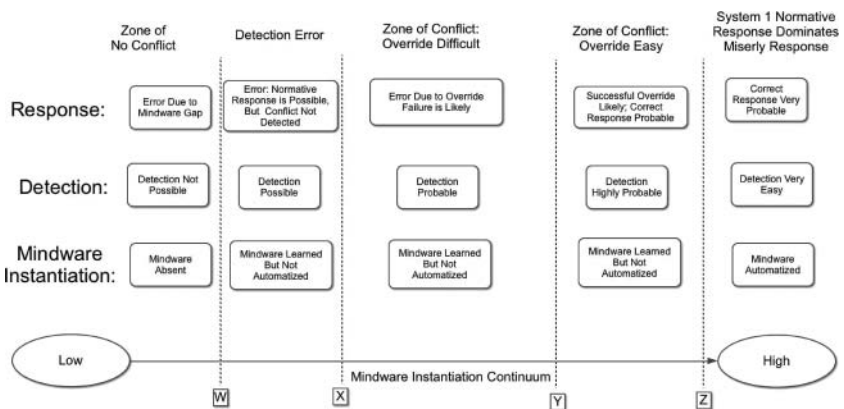


Figure 3. Processing states on the mindware continuum.

Conflict detection is more probable to the right of criterion X in [Figure 3](#). This makes override a possibility in that part of the graphic. Criterion Y demarcates successful from unsuccessful sustained override. To the left of criterion Y (and to the right of X) is the processing state that De Neys has explored in numerous studies (De Neys, [2006a](#), [2006b](#), [2014](#); De Neys & Franssens, [2009](#); De Neys & Glumicic, [2008](#); see also, Thompson & Johnson, [2014](#)). His research group demonstrated, with several heuristics and biases tasks, that various implicit measures of performance (decision latencies, unannounced recall, brain activation, autonomic arousal) indicated that conflict is sometimes detected even in cases where a subject did not successfully override the intuitive response.

It is important to note that in the heuristics and biases literature, some previous theorists have collapsed two areas in [Figure 3](#) (the two areas demarcated by criterion X) – calling both “failure to override”. They are differentiated here because it is important to distinguish sustained override failure – where Type-2 processing loses out to Type-1 processes in a conflict of discrepant outputs – from situations where Type-2 processing is not engaged at all.

Continuing rightward in [Figure 3](#), to the right of criterion Y and to the left of Z is the area of the mindware instantiation continuum where override is likely to be successful. Well-instantiated mindware makes detection of conflict easy and highly probable. The well-instantiated mindware also makes sustained retrieval during the simulation process easy, leading to a high probability of successful override. Finally, to the right of criterion Z is the processing state that we described before and illustrated in [Figure 2](#) – where no processing override of the intuitive response is necessary because the normative mindware is so well automated that it is the dominant response.

To summarise [Figure 3](#): the area to the left of W denotes an error due to missing mindware; the area between W and X an error due to detection failure; between X and Y an error due to override failure; between Y and Z a correct response achieved by sustained override; and to the right of Z a correct response due to automatic activation of the normative response from System 1. Moving from left to right in [Figure 3](#), the likelihood of deriving a correct response increases as the mindware becomes more automatised, conflict detection becomes easier, and override becomes easier as well. A novel contribution of [Figure 3](#) is that of demonstrating how the cognitive failures (and successes) differ as we move along the continuum from left to right.

The figure also illustrates how the degree of mindware instantiation is strongly related to the probability of successful detection and override. This is the reason why heuristics and biases tasks cannot be taken as pure processing indicators and why errors on such tasks cannot be unambiguously attributed to miserly processing. Additionally, it illustrates why correct responses are not necessarily the result of computationally expensive cognition.

Although no heuristics and biases task can be taken as a specific and unique indicator of a particular *type* of error, the three error types illustrated in [Figure 3](#) will be differentially associated with different types of tasks in the literature. For example, the failure of sustained override area in the figure (the area between X and Y) would be strongly indicated by subpar performance on tasks tapping aspects of self control and tasks that assess the temporal discounting of reward (Loewenstein, Read, & Baumeister, 2003). An error in the classic bat-and-ball problem from Frederick's (2005) CRT will often occur because of problems with conflict detection (the area between W and X). Finally, some tasks in the heuristics and biases literature are most often failed because of mindware gaps (the area to the left of W). Base-rate neglect and falling prey to the gambler's fallacy are probably most often a result of the absence of the mindware of probabilistic thinking (Stanovich, 2010, 2011).

The dependence between knowledge and process in heuristics and biases tasks can also be viewed from the hierarchical logic of [Figure 4](#), adapted from a discussion in Stanovich and West (2008). This figure is constructed in terms of which paths lead to the incorrect, miserly response and which paths lead to the normative response. [Figure 4](#) shows that detection and override are

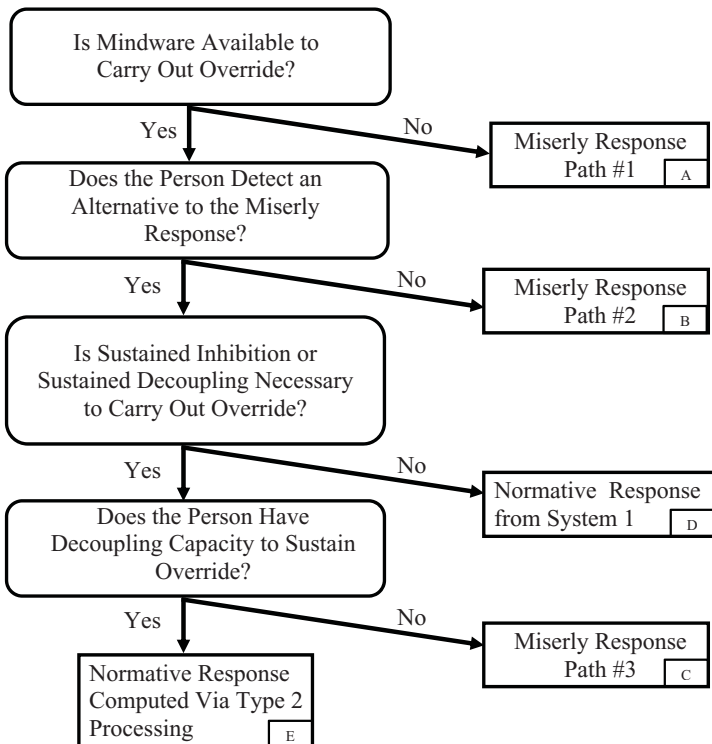


Figure 4. Dependence between knowledge and process in heuristics and biases tasks.

dependent on knowledge and that override is dependent on detection. We can see this by proceeding downward, sequentially. The question addressed in the first stage of the framework, at the top of [Figure 4](#), is whether, for a given task and subject, the mindware is available to carry out override (whether relevant declarative knowledge is available to be substituted for the intuitive response). If the relevant mindware is not available, then the person will produce a miserly response.

If the relevant mindware is in fact available, then the next question that becomes operative is whether or not the person detects the need to override the intuitive response. Even if the relevant mindware is present, if the subject does not detect any reason to override the intuitive response, then it will be emitted (this is path #2 to the intuitive response as labelled in the figure).

If the relevant mindware is present and if an alternative to the intuitive response has been noted, the question then becomes whether or not the task requires sustained inhibition (cognitive decoupling) in order to carry out the override of the intuitive response. If the normative mindware is so automated that it trumps the intuitive response, then no override is needed and a normative response is emitted.

In contrast, if the normative mindware is not instantiated in System 1 and if the task requires sustained decoupling in order to carry out override, then we must ask whether the subject has the cognitive capacity that will be necessary. If so, then the normative response will be given via Type-2 decoupling and simulation.⁴ If not, then the intuitive response will be given (path #3 to the intuitive response in the figure) – despite the availability of the relevant mindware and the recognition of the need to use it. In terms of [Figure 3](#), path #1 represents the area to the left of criterion W, path #2 represents the area immediately to the left of criterion X and path #3 represents the area immediately to the left of criterion Y.

[Figure 4](#) captures the dependence between mindware, detection and override. It makes clear that missing mindware renders questions about detection and override irrelevant. Mindware present at some minimal level at least provides the possibility of a detection error (or a correct response). The degree of mindware presence that enables (possible) detection then, indirectly, enables the possibility of sustained override (or override failure).

[Figure 4](#) also illustrates why correct responses are not necessarily slow and why incorrect responses are not necessarily fast. The five heuristics and biases response types are labelled A through E in [Figure 4](#) (A, B, C, for error responses, and D and E for correct responses). [Table 1](#) presents a 2×2 matrix

⁴Because it is tangential to the present discussion, I have not modelled another possibility suggested by Risen (2016) and by Stanovich (2004) – that successful Type-2 override outcomes may be rejected by a reflective third-order judgement (see Risen, 2016, on acquiescence, and Stanovich, 2004, pp. 228–243, on rational integration).

Table 1. Classification of errors and correct response types in Figure 4 as either fast or slow responses.

	Fast response	Slow response
Incorrect response	A: Mindware absent B: Detection failure	C: Override failure
Correct response	D: Automatised normative response stored in System 1	E: Normative response after successful override

crossing response accuracy with speed. The important thing to note is that there is a response type in each of the four cells represented. Fast incorrect responses can occur in two ways – when the relevant normative mindware is absent (A) or when the subject fails to detect the possibility of a response other than the miserly one (B). But slow error responses can occur as well – when override is attempted but fails (C). Optimal responses will indeed occur when the normative response is given after successful sustained override (E). But *fast* correct responses can occur as well – when the normative response is stored in System 1 and triggers automatically (D).

The framework outlined in Table 1 explicates the finding of Stupple, Pitchford, Ball, Hunt, and Steel (2017) that in the CRT the correlation between response time and accuracy is seriously reduced by “variance arising from responses that are neither correct analytic answers nor incorrect intuitive answers” (p. 15). This variance arising from responses that are neither correct analytic answers nor incorrect intuitive answers is represented by cells C and D in Table 1 – incorrect responses due to override failure (C) and correct responses due to normatively correct responses that are automatically triggered from System 1 (D).

Summary and research implications

Heuristics and biases tasks have been widely employed to assess the extent to which subjects engage in miserly information processing. The theoretical inferences made from correct and incorrect responding on such tasks have been overly simplified, however. It is not the case that making an error on such tasks can be unambiguously classified as indicating miserly processing. Likewise, correct responding on such tasks is not necessarily the result of computationally expensive processing.

That neither of these interpretations of responses on heuristics and biases tasks follow is clear from the framework developed here. That framework identified five processing states that are possible in most heuristics and biases tasks. Three of the states lead to response errors and two lead to correct responding. Because an error on such tasks may result from missing mindware in addition to miserly processing, an error does not unambiguously indicate that it was miserly processing rather than a mindware problem that was

the prime cause of the error. Even if it could be determined that mindware gaps were not responsible, the error could still result from two different processing states that differ somewhat in how miserly they are. The error might have resulted from a miserly default that was left uncorrected. Alternatively, the error might have resulted from an override failure as defined in the model in [Figure 3](#): the struggle between conflicting intuitive and normative responses in which sustained override was not achieved (C in [Figure 4](#) and [Table 1](#)). Failure of sustained override represents a less miserly tendency than does an unchecked default to the miserly response served up by the autonomous mind.

A parallel ambiguity also complicates the inferences that can be made from correct responding on heuristics and biases tasks. Such responding can eventuate from successfully sustained override or from the automatic triggering of overlearned mindware from System 1, as indicated in [Figure 1](#). This is why it cannot be inferred that all correct responses on such tasks will necessarily be slow.

In summary, normative responding on heuristics and biases tasks is multiply determined and incorrect responding also can result from a variety of information processing defects.⁵ The model presented here identifies three processing defects: inadequately learned mindware; failure to detect the necessity of overriding the miserly response; and failure to sustain the override process once initiated. These defects are intertwined in heuristics and biases tasks, however. For example, mindware overlearning facilitates detection and makes sustained override easier – thus the degree of mindware instantiation is strongly related to the probability of successful detection and override. Treating these tasks as pure indicators of miserly processing is unwarranted. Errors on them cannot be unambiguously attributed to miserly processing – and correct responses are not necessarily the result of computationally expensive cognition.

Acknowledgments

Valuable comments on this manuscript were received from Wim De Neys, Jonathan Evans, Valerie Thompson, Maggie Toplak, and Richard West.

⁵My purpose here has been to clarify some of the interrelationships among the concepts underlying miserly processing. Specifically, the dependence of the processing operations of detection and override on the presence of mindware has been the focus. The implications of mindware presence being a continuous quantity were of particular importance and these were explored in detail. My purpose was not to adjudicate all of the contentious issues surrounding the dual-process meta-theory, which would need a much longer and fleshed-out treatment. Nevertheless, regarding these larger issues, I can anticipate that some readers may see the clarifications outlined here as making dual-process theory less parsimonious. This is a valid worry, but I feel that not all of the points raised here really increase the theoretical degrees of freedom as much as they make the methods of a dual-process-based experiment a little more cumbersome. For example, a clear implication of the present essay is that we must measure the presence of the relevant mindware in order to understand the exact locus of a response error.

Disclosure statement

No potential conflict of interest was reported by the author.

References

- Achtziger, A., & Alos-Ferrer, C. (2014). Fast or rational? A response-times study of Bayesian updating. *Management Science*, *60*, 923–938.
- Ainslie, G. (2001). *Breakdown of will*. Cambridge: Cambridge University Press.
- Arkes, H. R. (2016). A levels of processing interpretation of dual-system theories of judgment and decision making. *Theory & Psychology*, *26*, 459–475.
- Bago, B., & De Neys, W. (2017). Fast logic? Examining the time course assumption of dual process theory. *Cognition*, *158*, 90–109.
- Clark, A. (2001). *Mindware: An introduction to the philosophy of cognitive science*. New York, NY: Oxford University Press.
- Dawes, R. M. (1976). Shallow psychology. In J. S. Carroll & J. W. Payne (Eds.), *Cognition and social behavior* (pp. 3–11). Hillsdale, NJ: Erlbaum.
- Dawkins, R. (1982). *The extended phenotype*. New York, NY: Oxford University Press.
- De Neys, W. (2006a). Automatic-heuristic and executive-analytic processing during reasoning: Chronometric and dual-task considerations. *Quarterly Journal of Experimental Psychology*, *59*, 1070–1100.
- De Neys, W. (2006b). Dual processing in reasoning – Two systems but one reasoner. *Psychological Science*, *17*, 428–433.
- De Neys, W. (2012). Bias and conflict: A case for logical intuitions. *Perspectives on Psychological Science*, *7*, 28–38.
- De Neys, W. (2014). Conflict detection, dual processes, and logical intuitions: Some clarifications. *Thinking & Reasoning*, *20*, 169–187.
- De Neys, W. (Ed.). (2018). *Dual process theory 2.0*. London: Routledge.
- De Neys, W., & Franssens, S. (2009). Belief inhibition during thinking: Not always winning but at least taking part. *Cognition*, *113*, 45–61.
- De Neys, W., & Glumicic, T. (2008). Conflict monitoring in dual process theories of thinking. *Cognition*, *106*, 1248–1299.
- de Sousa, R. (2007). *Why think? Evolution and the rational mind*. Oxford: Oxford University Press.
- Einhorn, H. J., & Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgment and choice. *Annual Review of Psychology*, *32*, 53–88.
- Evans, J. St. B. T. (2007). On the resolution of conflict in dual process theories of reasoning. *Thinking and Reasoning*, *13*, 321–339.
- Evans, J. St. B. T. (2008). Dual-processing accounts of reasoning, judgment and social cognition. *Annual Review of Psychology*, *59*, 255–278.
- Evans, J. St. B. T. (2010). *Thinking twice: Two minds in one brain*. Oxford: Oxford University Press.
- Evans, J. St. B. T. (2014). *Reasoning, rationality and dual processes*. London: Psychology Press.
- Evans, J. St. B. T. (2018). Dual process theory: Perspectives and problems. In W. De Neys (Ed.), *Dual process theory 2.0* (pp. 137–156). London: Routledge.
- Evans, J. St. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, *8*, 223–241.
- Ferreira, M. B., Mata, A., Dokin, C., Sherman, S. J., & Ihmels, M. (2016). Analytic and heuristic processes in the detection and resolution of conflict. *Memory & Cognition*, *44*, 1050–1063.

- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT University Press.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19, 25–42.
- Frey, D., Johnson, E., & De Neys, W. (2018). Individual differences in conflict detection during reasoning. *Quarterly Journal of Experimental Psychology*, 71.
- Gilbert, D. T. (2006). *Stumbling on happiness*. New York, NY: Alfred A. Knopf.
- Hull, D. L. (2001). *Science and selection: Essays on biological evolution and the philosophy of science*. Cambridge: Cambridge University Press.
- Jimenez, N. B., Rodriguez-Lara, I., Tyran, J., & Wengstrom, E. (2018). Thinking fast, thinking badly. *Economics Letters*, 162, 41–44.
- Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice Hall.
- Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58, 697–720.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus & Giroux.
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49–81). New York, NY: Cambridge University Press.
- Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, 64, 515–526.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3, 430–454.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237–251.
- Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavior and Brain Sciences*, 36, 661–679.
- La Berge, D., & Samuels, S. (1974). Toward a theory of automatic information processing in reading. *Cognitive Psychology*, 6, 293–323.
- Li, M., & Chapman, G. B. (2009). “100% of anything looks good”: The appeal of one hundred percent. *Psychonomic Bulletin & Review*, 16, 156–162.
- Li, N., van Vugt, M., & Colarelli, S. (2018). The evolutionary mismatch hypothesis: Implications for psychological science. *Current Directions in Psychological Science*, 27, 38–44.
- Liberali, J. M., Reyna, V. F., Furlan, S., Stein, L. M., & Pardo, S. T. (2012). Individual differences in numeracy and cognitive reflection, with implications for biases and fallacies in probability judgment. *Journal of Behavioral Decision Making*, 25, 361–381.
- Loewenstein, G. F. (1996). Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes*, 65, 272–292.
- Loewenstein, G. F., Read, D., & Baumeister, R. (Eds.). (2003). *Time and decision: Economic and psychological perspectives on intertemporal choice*. New York, NY: Russell Sage.
- Moore, R. (2017). Fast or slow: Sociological implications of measuring dual-process cognition. *Sociological Science*, 4, 196–223.
- Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin*, 132, 297–326.
- Navon, D. (1989). The importance of being visible: On the role of attention in a mind viewed as an anarchic intelligence system. *European Journal of Cognitive Psychology*, 1, 191–238.
- Newman, I. R., Gibb, M., & Thompson, V. A. (2017). Rule-based reasoning is fast and belief-based reasoning can be slow: Challenging current explanations of belief-bias

- and base-rate neglect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43, 1154–1170.
- Nozick, R. (1993). *The nature of rationality*. Princeton, NJ: Princeton University Press.
- Oaksford, M., & Chater, N. (2014). Probabilistic single function dual process theory and logic programming as approaches to non-monotonicity in human vs. artificial reasoning. *Thinking & Reasoning*, 20, 269–295.
- Pennycook, G., Fugelsang, J. A., & Koehler, D. J. (2015). What makes us think? A three-stage dual-process model of analytic engagement. *Cognitive Psychology*, 80, 34–72.
- Perkins, D. N. (1995). *Outsmarting IQ: The emerging science of learnable intelligence*. New York, NY: Free Press.
- Posner, M. I., & Snyder, C. (1975). Facilitation and inhibition in the processing of signals. In P. Rabbitt & S. Dornic (Eds.), *Attention and performance* (Vol. 5, pp. 669–682). London: Academic Press.
- Rand, D. G., Tomlin, D., Bear, A., Ludvig, E. A., & Cohen, J. D. (2018). Cyclical population dynamics of automatic versus controlled processing: An evolutionary pendulum. *Psychological Review*, 124(5), 626–642.
- Reyna, V. F., & Brainerd, C. J. (2011). Dual processes in decision making and developmental neuroscience: A fuzzy-trace model. *Developmental Review*, 31, 180–206.
- Richerson, P. J., & Boyd, R. (2005). *Not by genes alone: How culture transformed human evolution*. Chicago, IL: University of Chicago Press.
- Risen, J. L. (2016). Believing what we do not believe: Acquiescence to superstitious beliefs and other powerful intuitions. *Psychological Review*, 123, 182–207.
- Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, 134, 207–222.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127–190.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69, 99–118.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63, 129–138.
- Sinayev, A., & Peters, E. (2015). Cognitive reflection vs. calculation in decision making. *Frontiers in Psychology*, 6, 532.
- Slooman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3–22.
- Slovic, S., & Slovic, P. (2015). *Numbers and nerves: Information, emotion, and meaning in a world of data*. Corvallis: Oregon State University Press.
- Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning*. Mahwah, NJ: Erlbaum.
- Stanovich, K. E. (2004). *The robot's rebellion: Finding meaning in the age of Darwin*. Chicago, IL: University of Chicago Press.
- Stanovich, K. E. (2010). *Decision making and rationality in the modern world*. New York, NY: Oxford University Press.
- Stanovich, K. E. (2011). *Rationality and the reflective mind*. New York, NY: Oxford University Press.
- Stanovich, K. E., & Toplak, M. E. (2012). Defining features versus incidental correlates of Type 1 and Type 2 processing. *Mind & Society*, 11, 3–13.
- Stanovich, K. E., & West, R. F. (2008). On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology*, 94, 672–695.

- Stanovich, K. E., West, R. F., & Toplak, M. E. (2016). *The Rationality quotient: Toward a test of rational thinking*. Cambridge, MA: MIT Press.
- Stich, S. P. (1990). *The fragmentation of reason*. Cambridge: MIT Press.
- Stupple, E. J., Pitchford, M., Ball, L. J., Hunt, T., & Steel, R. (2017). Slower is not always better: Response-time evidence clarifies the limited role of miserly information processing in the Cognitive Reflection Test. *PLoS One*, *12*(11), e0186404. Retrieved from <https://doi.org/10.1371/journal.pone.0186404>.
- Sun, R. (2015). Interpreting psychological notions: A dual-process computational theory. *Journal of Applied Research in Memory and Cognition*, *4*, 191–196.
- Taylor, S. E. (1981). The interface of cognitive and social psychology. In J. H. Harvey (Ed.), *Cognition, social behavior, and the environment* (pp. 189–211). Hillsdale, NJ: Erlbaum.
- Thaler, R. H. (2015). *Misbehaving: The making of behavioral economics*. New York, NY: Norton.
- Thompson, V. A. (2009). Dual-process theories: A metacognitive perspective. In J. Evans & K. Frankish (Eds.), *In two minds: Dual processes and beyond* (pp. 171–195). Oxford: Oxford University Press.
- Thompson, V. A. (2014). What intuitions are... and are not. In B. Ross (Ed.), *Psychology of learning and motivation* (Vol. 60, pp. 35–75). New York, NY: Elsevier.
- Thompson, V. A., & Johnson, S. C. (2014). Conflict, metacognition, and analytic thinking. *Thinking & Reasoning*, *20*, 215–244.
- Thompson, V. A., & Morsanyi, K. (2012). Analytic thinking: Do you feel like it? *Mind & Society*, *11*, 93–105.
- Thompson, V. A., Prowse Turner, J. A., & Pennycook, G. (2011). Intuition, reason, and metacognition. *Cognitive Psychology*, *63*, 107–140.
- Tomlin, D., Rand, D. G., Ludvig, E. A., & Cohen, J. D. (2015). The evolution and devolution of cognitive control: The costs of deliberation in a competitive world. *Scientific Reports*, *5*, 11002.
- Toplak, M. E., West, R. F., & Stanovich, K. E. (2011). The Cognitive Reflection Test as a predictor of performance on heuristics and biases tasks. *Memory & Cognition*, *39*, 1275–1289.
- Toplak, M. E., West, R. F., & Stanovich, K. E. (2014). Assessing miserly processing: An expansion of the Cognitive Reflection Test. *Thinking & Reasoning*, *20*, 147–168.
- Trippas, D., Handley, S. J., Verde, M. F., & Morsanyi, K. (2016). Logic brightens my day: Evidence for implicit sensitivity to logical validity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*, 1448–1457.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.
- Wang, L. (2009). Money and fame: Vividness effects in the National Basketball Association. *Journal of Behavioral Decision Making*, *22*, 20–44.
- Westbrook, A., & Braver, T. S. (2015). Cognitive effort: A neuroeconomic approach. *Cognitive Affective and Behavioral Neuroscience*, *15*, 395–415.