To be published in *Behavioral and Brain Sciences* (in press)

Subtracting "Ought" From "Is": Descriptivism Versus Normativism in the Study of the Human Thinking

Shira Elqayam
Division of Psychology, School of Applied Social Sciences, Faculty of Health and Life Sciences, De Montfort University, The Gateway, Leicester, LE1 9BH, UK.
selqayam@dmu.ac.uk
http://www.psy.dmu.ac.uk/elqayam

Jonathan St.B.T. Evans
School of Psychology, Faculty of Science, University of Plymouth, Drake Circus, Plymouth, PL4 8AA, UK.
jevans@plymouth.ac.uk
http://www.plymouth.ac.uk/staff/jevans

Abstract: We propose a critique of normativism, defined as the idea that human thinking reflects a normative system against which it should be measured and judged. We analyze the methodological problems associated with normativism, proposing that it invites the controversial is-ought inference, much contested in the philosophical literature. This problem is triggered when there are competing normative accounts (the arbitration problem*)*, as empirical evidence can help arbitrate between descriptive theories, but not between normative systems. Drawing on linguistics as a model, we propose that clear distinction between normative systems and competence theories is essential, arguing that equating them invites an 'is-ought' inference; to wit, supporting normative 'ought' theories with empirical 'is' evidence. We analyze in detail two research programs with normativist features, Oaksford and Chater's rational analysis, and Stanovich and West's individual differences approach, demonstrating how in each case equating norm and competence leads to an is-ought inference. Normativism triggers a host of research biases in psychology of reasoning and decision making: focusing on untrained participants and novel problems, analyzing psychological processes in terms of their normative correlates, and neglecting philosophically significant paradigms when they do not supply clear standards for normative judgment. For example, in a dual-process framework, normativism can lead to a fallacious 'ought-is' inference, in which normative responses are taken as diagnostic of analytic reasoning. We propose that little can be gained from normativism that cannot be achieved by descriptivist computational-level analysis, illustrating our position with Hypothetical Thinking Theory and the theory of the suppositional conditional. We conclude that descriptivism is a viable option, and that theories of higher mental processing would be better off freed from normative considerations.

Keywords: Bayesianism; competence; computational level analysis; descriptivism; is-ought inference; logicism; normative systems; normativism; rational analysis; rationality; research bias; understanding / acceptance principle

'Would you tell me, please, which way I ought to go from here?'

'That depends a good deal on where you want to get to,' said the Cat.

'I don't much care where –' said Alice.

'Then it doesn't matter which way you go,' said the Cat.

Lewis Carroll, *Alice's adventures in Wonderland*

1. Logicism and normativism and their discontents

In everyday life, we are thoroughly accustomed to normative dictates wherever we turn. When we play chess, we conform to the rules of the game; when we drive, we try to heed traffic laws and know we would be sanctioned if we disobeyed them. In some countries, language is normatively regulated - *L'Académie française* is a prominent example. Voluntary or governmental bodies, like the Advertising Standards Authority in the UK, impose normative constraints on advertisements. And occasionally, normative issues find their way into scientific theories as well.

The research literature in higher mental processing – reasoning, judgment and decision making– is rife with normative considerations. In the study of human reasoning, these have traditionally taken the form of *logicism*– the idea that thinking (1) reflects some internalized form of extensional, classical logic and (2) should be measured against classical logic as a normative system (Evans, 2002) and *ought* in some clearly evaluative sense to conform with it (see Appendix for terminological clarifications). We will dub these two distinct meanings *empirical* versus *prescriptive* logicism respectively. Inhelder and Piaget hold the dubious title of prototype logicists: in their monograph on the formal operations stage of cognitive development (Inhelder & Piaget, 1958), they argued that normal adolescents and ultimately adults attain the ability to reason according to the rules of formal classical logic.

Half a century on, logicism in both its forms is not nearly as dominant in reasoning research as it used to be. Peter Wason's seminal work in the 1960's and 70's was motivated by an attack on empirical logicism in a period dominated by Piagetian theory. In support of this he devised several ingenious reasoning problems including the '2 4'6' task (Wason, 1960), the much researched selection task (e.g., Wason, 1966) and the THOG problem (Wason & Brooks, 1979). However, Wason never seemed to doubt that human reasoning *should* conform to classical logic (i.e., prescriptive logicism), so that his interpretation of the many logical errors observed on his tasks was that people are illogical and *therefore irrational* (see Evans, 2002, for a detailed account). Following the critique of Cohen (1981), however, later researchers began to question whether logic was the right normative system against which to judge the rationality of people's reasoning, so that prescriptive logicism also came under attack. Some researchers have proposed that we should adopt alternative normative systems such as those based on information, probability or decision theory (e.g., Oaksford & Chater, 1991; 1998; 2007), while others proposed that at least some forms of rationality need not necessarily require a normative system at all (e.g., Evans, 1993; 2002; Evans & Over, 1996; Gigerenzer & Selten, 2001). By this position, organisms are rational if they act in such a manner as to achieve personal goals, and that such rationality need not involve any normative rule following.

Our concern here is not with logicism *per se*; in our view, logicism is but a special case of a more general attitude. Consider the empirical and prescriptive tenets of logicism. We could easily substitute for the word 'logic' a name of another normative system, such as Bayesian probability:

Empirical logicism: Thinking reflects logic.

Prescriptive logicism: Rational thinking should be measured against logic as a normative.

Empirical Bayesianism: Thinking reflects Bayesian probability.

Prescriptive Bayesianism: Rational thinking should be measured against Bayesian probability as a normative system.

Our own take on this is that both logicism and Bayesianism are special cases of the same paradigm. We call this paradigm 'normativism'; analogous to what Stein (1996) calls 'the standard picture', it can be formulated in terms closely related to the ones we have already examined. These are:

Empirical normativism: Thinking reflects S.

Prescriptive normativism: Rational thinking should be measured against S as a normative system, and ought to conform to it.

where S is a formal normative system such as logic (classical or otherwise), Bayesian probability, or decision theory. Note a formal theory is not necessarily a normative theory unless taken as such by a specific normativist account. For example, extensional logic can be conceived as a useful computational tool rather than a normative standard for human reasoning. The notable exception is Subjective Expected Utility (SEU), which was developed with a normative goal in the first place (Savage, 1954, p. 19; von Neumann & Morgenstern, 1947, pp. 8-9). Taken in this sense, widely diverse research programs can be said to be normativist. For example, much of the judgment and decision making (JDM) literature is normativist in the *prescriptive* (albeit not in the empirical) sense, with SEU playing the role of the normative system. Even the most famous (and Nobel Prize winning) descriptive theory of risky decision making – the prospect theory of Kahneman and Tversky (1979) – was framed as a demonstration that the standard normative account provided by decision and probability theory failed to accurately describe human economic behavior.

The prescriptive and empirical tenets of normativism can be considered as vectors defining a two-dimensional space, which makes normativism (and its subordinate paradigms, such as logicism and Bayesianism) a matter of varying degrees. Figure 1 maps out the normative space defined by the empirical and prescriptive normativism respectively, on which we have placed a number of leading authors for illustrative purposes. We realize that some readers would wish to debate the exact co-ordinates assigned but that is not important for our purposes here. Note that the right-hand side – the high prescriptive side – of Figure 1 is much more crowded. This is hardly surprising. Historically, positions of prescriptive normativism tend to survive longer than positions of empirical normativism, because – for reasons we will explore later – they are much more difficult to eliminate. A notable example is, as already indicated, Wason's rejection of empirical logicism while continuing to uphold prescriptive logicism (Evans, 2002), and the heuristics and biases program of Tversky and Kahneman (e.g., Kahneman & Tversky, 1979). But a consequence of such a line of argument is that one must conclude people to be irrational. That is why those who feel people should be rational have proposed alternative normative systems (e.g. Cohen, 1981).

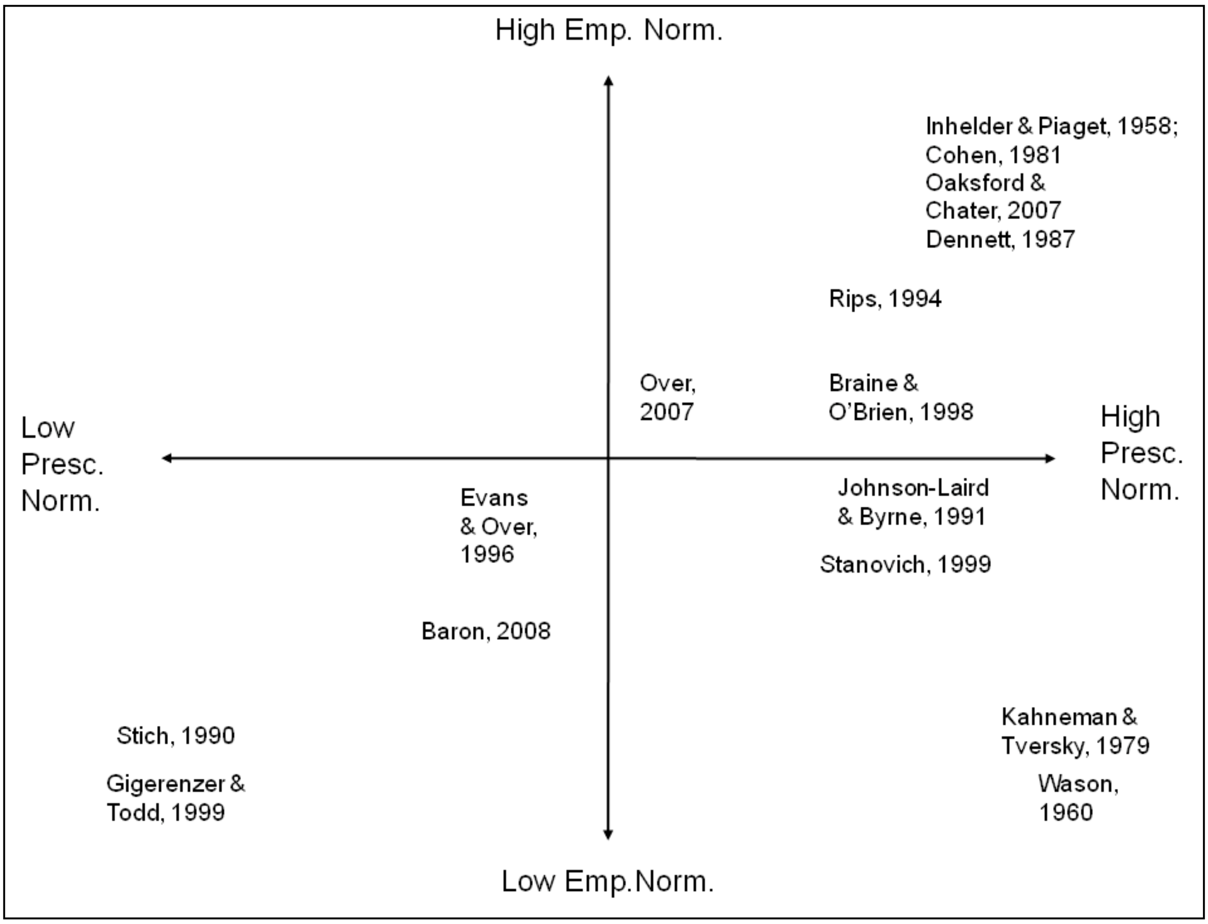

Figure 1: Two-vector normativist space with sample references.

Note: Emp. Norm.: *Empirical normativism:* Thinking reflects normative system S. Pres. Norm.: *Prescriptive normativism*: Thinking should be measured against S as a normative system and ought to conform to it. (For reasons of space, each research program is identified in the Figure by a single sample reference.)

Note, too, that the upper left quadrant of the normative space mapped out in Figure 1 is empty, highlighting that there is no coherent way of proposing high empirical normativism with low prescriptive normativism. In other words, the existence of a normative system is a necessary (albeit not sufficient) condition for the empirical facts of satisfying this system.

Empirical normativism can vary from hard-core positions which consider thought processes to be isomorphic to the normative system (e.g., Inhelder and Piaget's formal operations), to positions which reject the normative system entirely. For example, Gigerenzer (e.g., Gigerenzer & Todd, 1999) famously repudiates any form of normative systems, arguing that heuristic rules of the thumb outperform normative computations. In between are positions which might be termed 'soft logicism', and which postulate that some logical principles, such as non-contradiction, might underlie some thinking, but only to a limited extent (e.g., Over, 2007). Prescriptive normativism can vary according to factors such the *a priori* status of the normative system, the position famously advocated by the philosophers Jonathan Cohen (1981) and Daniel Dennett (1987). Psychologists tend more to regard selecting the appropriate normative system as an empirical issue, a view shared by authors leading such diverse research programs as Oaksford and Chater's rational analysis, which focuses on modal responses as a source for normative evaluations (Oaksford & Chater, 1998; 2007) and the earlier phase of Stanovich's individual differences program (Stanovich, 1999), which focused on the normatively superior performance of cognitively able participants. We will discuss the problems with this approach in section 4, and examine these research programs in further detail in section 5.

Another factor that may vary is whether conformity to a normative system is considered both necessary and sufficient for rationality or only necessary – the latter seems to be more common. Positions high on prescriptive normativism are also typically universalist, explicitly or implicitly taking the view that there is just one 'right', all-encompassing normative system, and all the others are 'wrong'. However, this can still vary to some extent, with some authors (Oaksford & Chater, 1998; 2007) advocating one normative system across the board, while others are willing to accept a different normative solution for each specific task (Stanovich, 1999; 2004; Stanovich & West, 2000). A relativist position of the sort famously advocated by Stich (1990), and to some extent by Baron (2008) and by Stenning and van Lambalgen (2008), would place these programs lower on prescriptive normativism.

In what follows, our main concern is with the prescriptive tenet of normativism – the belief that people ought to conform to a normative standard - although we will have a few things to say about empirical normativism as well. Our thesis is that prescriptive normativism is both problematic and unnecessary in scientific studies of human thinking. We will start by outlining what we mean by normativism in reasoning and decision making research, and how it differs from other forms of rationality. We then examine the possible relations between normative systems and psychological evidence, focusing in particular on the thorny problem of arbitration – cases of conflict between competing normative systems. It will become clear that we have no quarrel with the use of formal theories per se, provided that they are used in a descriptive rather than normative manner.

We shall discuss several problems that normativist thinking has created. First, research programs have been used to derive normative claims from empirical evidence, relying on the controversial inference from *is* to *ought*. We illustrate this with discussion of two leading research programs in the study of the human thinking. Next, we argue that normativism has systematically and harmfully biased the scientific study of thinking, affecting what is studied, how it is studied and how findings are interpreted. We illustrate these further by discussion of the particular problems that normativist thinking has created in the study of dual processes in higher cognition. Finally, we argue that normativism is unnecessary: a descriptive approach aided by computational analysis can address all the relevant scientific questions, ridding the field of the research biases we discuss. We hence conclude that theories of reasoning, judgment and decision making would be better off liberated from normative goals, and augmented by a descriptivist agenda.

2. Normativism, rationality, and the three senses of 'ought'

Normative rationality is not the only type. Here are some of the other concepts of rationality to be found in the literature (also see Nickerson, 2008, for a recent review):

- *Instrumental rationality*. Behaving in such a way as to achieve one's personal goals
- *Bounded rationality*. Behavior that is adaptive within the constraints of cognitive and biological capacity
- *Ecological rationality*: Behavior that is adapted to the environment in which the organism is operating
- *Evolutionary rationality*: Behavior that has been shaped by evolution and which serves the purpose of the genes

What seems to set apart normative rationality from other types of rationality is the 'oughtness' involved in normativism. Bounded rationality, for example, is not bounded because it *ought* to be so. There are just biological limits to how large brains can grow and how much information and how many computational algorithms they can store and execute. There is no 'oughtness' to the Darwinian and Skinnerian algorithms that shape ecological rationality either. Adaptation to the environment is an 'is', not an 'ought'. Darwinian principles are like Newton's laws of mechanics. Unsupported objects fall to the earth not because they ought to, but because that is what the laws of physics dictate. In the same way, there appears to be no scientific justification for 'intelligent design' in evolution. Organisms develop adaptations in accordance with the laws of natural and sexual selection in much the same way as apples fall off trees in compliance with the law of gravity.

A possible argument here is that oughtness is part of what biological function is about; that the idea of function is basically a normative one.[1] Often this argument is couched in adaptationist terms; for example, that the heart has a ‘proper function’ (in the terminology suggested by Millikan; e.g., 1984; 1995; 1996) to pump the blood, which is what it was selected for and therefore what it ‘ought’ to do (although cf. Fodor, 2008; and for response Dennett, 2008). One could even take this argument further and maintain that, by losing ‘oughtness’, we lose our ability to talk about function at all, biological, economic, or otherwise.[2] However, our point is that functional ‘ought’ is a different type of ‘ought’ than the one involved in normativism. ‘Ought’, and its close relations ‘should’ and ‘must’, can take at least three different meanings. Consider:

1. Poverty should not exist.
2. You must take the second exit from the roundabout.
3. Ron should be able to catch the 4.25 to Birmingham.

1 and 2 are deontic: they express evaluation and obligation; 3, on the other hand, is epistemic, expressing belief or probability. In addition, there is a difference between the deontic function of 1, which is evaluative, and of 2, which is to direct a specific course of action. (Schurz (1997) makes a related distinction between what he terms ‘normative’ and ‘valuative’, roughly equivalent to our directive and evaluative sense respectively.) In everyday discourse the directive and evaluative oughts are often combined, as in ‘I ought to donate to Oxfam’. However, as 2 demonstrates, these two deontic senses can be distinguished. Directive oughts are generally instrumental[3] – we need to take that second exit because it would bring us to our destination.

The ‘ought’ of normativism is evaluative: it resembles 1. In contrast, the ‘ought’ of selection-for and of functional analysis in general is directive, as in 2. There was no normative obligation for nature to select hearts for pumping. Natural selection can be said to contain a directive ought in the sense that function constrains (at least to some extent) evolution; what it does not have is the evaluative ‘ought’. With this caveat in place, we have no argument with the rational analysis approach of Anderson (1990), whose main thesis is that rationality is best understood by formal task analysis. We do not need to take a position in the debate over the role of adaptations in evolution (see, e.g., Gould & Lewontin, 1979; and then Fodor, 2008; and Dennett, 2008, respectively), to be wary of normativism. Within bounds, behavior is likely to be adaptive, so that analysis of the task and its environment may well be helpful in developing a formal account of human behavior. Insofar as a research program asks, as Oaksford and Chater's does in their adaptation of Anderson, which of several formal systems is most helpful to achieve one's goals, this falls under our definition of directive ought. It is only when formal systems are regarded as having *a priori*, unconditional value, that the 'ought' becomes an evaluative one. This is a very different argument than the one that leads from rational analysis to normative theory, and that, too, is part of Oaksford and Chater's research program (see section 5).

With this distinction in mind, we can now rephrase some of the debate over instrumental rationality. The separation proposed by Evans and Over (1996) between instrumental and normative rationality (i.e., achieving one's goals versus obeying a normative system respectively) has been contested by various authors. Oaksford and Chater (1998; 2007) objected on the grounds that instrumental rationality needs to be justified, and that this justification should be normative, hence obliterating the boundaries between normative and instrumental rationality. In the terminology proposed here, Oaksford and Chater see the directive 'ought' as inseparable from the evaluative 'ought', while we argue that these two senses are best kept apart.

So it appears to us that normativism is neither necessary nor helpful in discussions of function, adaptation, and ecological and instrumental rationality. Our task as scientists is to observe what people do and to construct and test theories of how they do it. That behavior is *typically* well adapted and that people *typically* achieve their personal goals (with many exceptions, of course) can be described in some terms as 'rational' but without recourse to any normative theory of what people ought to be doing. It is an observation to be accounted for, rather than obligation to be fulfilled.

3. Normative systems and the problem of arbitration

For a normativist position to be coherent, in particular the prescriptive tenet, it has to have a selective notion of norm: there is a 'right', or 'appropriate' normative system for a paradigm (or even all across the board), and there are 'wrong' ones. Nozick (1993), for example, argues for 'nomic universals' – scientific law-like statements, suggesting that norms cannot be particular. In some areas of cognition, deciding on the appropriate norm – and the closely related notion of 'error' – does not seem to pose a practical problem. In most memory paradigms, for example, an error is when one falsely identifies a new stimulus as previously presented, or fails to identify an old stimulus – a practice that goes back to Ebbinghaus and the earliest days of experimental psychology. For psychologists, the problem becomes acute when one tries to adopt this sort of ‘signal detection’ paradigm to reasoning and decision making research, and this is where consensus on the normative system conspicuously fails. But without a clear-cut norm, normativism becomes far shakier.

Normativism thus faces a problem when more than one normative system seems to fit the bill; Evans (1993) calls this the ‘normative system problem’; Stanovich (1999) – ‘the inappropriate norm argument’ (see also Cohen, 1981; Cohen, 1982; Gigerenzer, 1991; Lopes, 1991). Deciding on an appropriate normative system for any set of experimental findings is, more often than not, far from obvious. Indeed, unlike memory, one is hard put to find an experimental paradigm in reasoning and decision making that has just one obvious norm to compare against and no competing alternative norms. In the following, we propose a typology of three normative situations, based on the nature and number of competing normative accounts of a particular experimental paradigm. Of the three types, two involve normative conflict and one involves no conflict. Table 1 summarizes them.

Table 1: The three types of normative conflict

Type	**Conflict / No conflict**	**No. / type of norms involved**	**Examples**
Single	No Conflict	One	Conditional elimination inference
Alternative	Conflict	One Standard + at least one alternative	Conditional introduction inference Wason selection task
Multiple	Conflict	Several, equally standard	Metadeduction

With one established norm and no conflict, single norm paradigms seem to offer the prototypical normativist situation; and, in other cognitive paradigms, they do. Thus, in memory, in signal detection, in most theory of mind paradigms, what is ‘right’ and what is ‘wrong’ is ordinarily beyond dispute. Either there is (for example) a visual signal or there isn’t: the experimental environment is constructed so as to obviate the question. Not so, however, in reasoning and decision making, where single norm paradigms are increasingly rare. One of the few remaining single norm paradigms in reasoning seems to be conditional inference, or, more specifically, conditional *elimination* inferences. Such inferences are typically comprised of a ‘major’, conditional premise of the form ‘if p, then q’, and a categorical premise, e.g., ‘p’. The conclusion is categorical, *eliminating* the conditional form. There are also conditional introduction inferences, in which the conditional form is the conclusion of the inference rather than (one of) its premise(s). Table 2 presents several types of conditional inference.

Table 2: Types of conditional inference

	Inference type	**Form**	**Example**
Conditional elimination inference	Modus Ponens (MP)	*If p then q* *p* *Therefore, q*	If it snows the path will be icy It snows Therefore, the path is icy
	Denial of the Antecedent (DA)	*If p then q* *Not p* *Therefore, not q*	If it snows the path will be icy It does not snow Therefore, path is not icy
	Affirmation of the Consequent (AC)	*If p then q* *q* *therefore, p*	If it snows the path will be icy The path is icy Therefore, it snows
	Modus Tollens (MT)	*If p then q* *Not q* *Therefore, not p*	If it snows the path will be icy The path is not icy Therefore, it does not snow
Paradoxes of material implication (conditional introduction inference)	Paradox 1	*q* *Therefore, if p then q*	The path is icy Therefore, if it snows the path will be icy
	Paradox 2	*Not p* *Therefore, if p then q*	It does not snow Therefore, if it snows the path will be icy

The conditional elimination inferences constitute a single norm paradigm: regardless of one's theoretical position, MP and MT are generally considered valid types of inference, while AC and DA are invalid. Although this validity can and has been contested under specific conditions (e.g., McGee, 1985), experimental paradigms are generally constructed to avoid these conditions. However, this is only half the story. When conditional inference is viewed as a whole, normative considerations are by no means uncontroversial. For example, the paradoxes of material implication (see Table 2) are the subject of some intensive dispute, considered valid in mental model theory (Johnson-Laird & Byrne, 2002; Schroyens, 2010) but deemed invalid in probabilistic approaches (e.g., Evans, Over, & Handley, 2005; Oaksford & Chater, 2007). Hence, when participants judge the paradoxes as invalid (Pfeifer & Kleiter, 2011), mental model theory judges their judgment erroneous, while probabilistic approaches regards it perfectly normative (and see Over, Evans, & Elqayam, 2010, for discussion of another type of conditional introduction inference with conflicting normative judgments).

While the conditional elimination inferences can be considered single norm, conditional introduction inferences, then, are subject to dispute. We call these 'alternative norm' paradigms. Extensively covered by Stanovich (1999), alternative norm paradigms are far more prevalent in psychology of reasoning and JDM (judgment and decision making). In a typical debate of this type, a standard account of a particular observation competes with another, alternative account (or accounts), making an observed behavior normatively rational according to the latter but not according to the former (and vice versa). Examples of alternative norm paradigms are legion (Stanovich, 1999, reviews some classic alternative norm paradigms; Hahn & Warren, 2009, reviews some recent such developments in JDM); the longer a paradigm is studied, the more it tends to have alternative normative systems proposed.

The classic case is probably the Wason selection task (Wason, 1966): a hypothesis testing task designed to test understanding of the logic of conditionals. In the abstract version of this famous task, participants are presented with four cards bearing values such as A, G, 3 and 7, and given a conditional rule of the general form 'if p then q', such as 'if there is an A on one side of the card then there is a 3 on the other side'. Their task is to turn over all the cards, and only the cards, that need to be examined in order to decide whether the rule is true or false. The task is notoriously difficult in its standard, abstract form (for a recent review see Evans & Over, 2004) with only about ten percent of participants (of higher IQ) typically finding the standard normative solution, p and not-q – in this case, A and 7 (A because a not-3 number on the other side would disprove the rule and 7 because an A on the other side would do the same). However, Wason's normative departure point was logicist: the material conditional of the propositional calculus, according to which a conditional statement, *If p then q*, is true whenever q is true or p is false. When measured against alternative normative systems, such as decision theory (Manktelow & Over, 1991), Bayesian probability or information theory (Oaksford & Chater, 1994; 1996), or default logics (Stenning & van Lambalgen, 2008), the prevalent choices can be argued to be rational. For example, Oaksford and Chater (e.g., 1994; 2007) argue that participants select the optimal information in order to decide whether q depends on p or not, and are therefore normatively rational in terms of gaining information.

Clearly, alternative norm paradigms pose a major challenge for normativism. If there is just one 'correct' normative system, what are the mechanisms to arbitrate between the competing accounts? Another problem with the alternative norm paradigm is that what makes one account ‘standard’ and the other ‘alternative’, is often hard to determine. Why, for example, should classical logic be considered ‘standard’ in the case of the selection task, and information theory considered ‘alternative’? Because classical logic was the first proposed or has been around the longest? Oaksford and Chater (2007) have recently argued that Bayesian probability is becoming the dominant paradigm in cognitive science. If this is true, then the current Kuhnian paradigm for the selection task is probabilistic, but the original normative system – and the one that has been around longest – is deductive. So which should we view as the standard and which the alternative?

The problem becomes even more striking when we consider multiple norm paradigms, in which there are several normative systems available but none that appear to be standard. For example, consider the reasoning literature on metadeduction (e.g., Byrne & Handley, 1997; Byrne, Handley, & Johnson-Laird, 1995; Elqayam, 2006; Rips, 1989; Schroyens, Schaeken, & d'Ydewalle, 1999). In this paradigm, reasoners are presented with the Island of Knights and Knaves, in which the inhabitants are either knaves (liars) or knights (truth-tellers). The task is to identify the speakers based on their statements. It is generally assumed in the metadeduction literature that statements can be assigned truth-value based on partial information; for example, that one false conjunct is sufficient to make a conjunction false (so its speaker can be identified as a knave). But consider: ‘I am a knave and snow is black’, described by most participants as *indeterminate* (Elqayam, 2006). Is such a response erroneous, then? The difficulty is that the statement ‘I am a knave’ is paradoxical: it is a version of the Liar paradox (e.g., Martin, 1984). The issue now becomes evaluation of sentences with paradoxical constituents – which brings us to many-valued logics. As Elqayam (2003) argued, given the plethora of many-valued logics (for reviews see Gottwald, 2001; Rescher, 1969), there is little ground for preferring one type of system over the other.

The (increasing) scarcity of single norm paradigms in reasoning and decision making poses a major problem for normativism, since the latter depends on an agreed norm for assessment. Psychologists can, of course, and do get involved in argument about which one is right: perhaps an odd activity for empirical scientists. In fact, the temptation to which they often succumb is to try to resolve the issue empirically. But this leads them into the questionable form of argumentation that involves *is-ought inference,* discussed below. First, we clarify the status of formal theories and their role in empirical science.

4. The computational, the competent and the normative

From the discussion above, the reader might have formed the impression that we reject formal systems entirely in favor of purely processing accounts. However, our objection is not to formal systems *per se*, but to their use as normative systems; it is the deontic, evaluative 'ought' that we caution against. We have no problem with formal systems as competence or computational level systems. Indeed, each of us separately has previously used formal systems as major source of inspiration to construct a psychological theory, albeit on the computational rather than normative level. For example, Evans and Over (2004) utilized the suppositional conditional (Edgington, 1995; 2001); Elqayam (2006) used Kripke's (1975) theory of truth. We did so in much the same way that Chomskyan grammar provided and still provides inspiration to psycholinguistic and neurolinguistic research. (We will take this up again in more detail in section 7.)

This distinction between competence theory and normative theory is paramount to our argument. To illustrate it, we will start with linguistics, where a tradition going back to De Saussure (1966 [original publication 1916]) clearly separates descriptive from normative accounts in favor of the former. Here is a classic example. Consider double negation, as in 'I don't know nothing'. Countless primary school teachers have lectured countless generations that double negation is not Good English. However, double negation is part of the grammar in some variants of English, such as African American Vernacular English (AAVE): A theory seeking to describe the linguistic competence of AAVE speakers would have to include it. Double negation, then, is part of a competence theory of AAVE, although it falls outside normative grammar. While descriptive competence theories aim to describe the rules of language actually used by speakers, normative approaches aim to regulate speech in particular ways, sometimes motivated by a social, educational or political agenda that has little to do with the way human language works. For example, Mustafa Kemal Atatürk's reform of the Turkish language, 'purging' it from centuries of Arabic influence (Lewis, 1999), was grounded in nationalist normativism.

There are quite a few categorizations of levels of enquiry in the cognitive literature (for a review see Stanovich, 1999), but Chomsky's and Marr's are probably the most influential ones in cognitive science, so we will limit ourselves to these two. We use the term 'competence' here in the Chomskyan sense, which is to say, a structural description of abstract knowledge which is quite value-free (although cf. Harris, 1980; 1981). 'Competence' is not intended to be contrasted with 'incompetence', but rather with 'performance', i.e., the instantiation of linguistic competence in actual speech. The

Chomskyan notion of competence is parallel to Marr's (1982) conception of the computational level of analysis – the level that describes *what* is being computed and *why* (e.g., the rules of arithmetic). Marr himself noted the analogue to Chomsky; what Marr's conception adds is the notion of function, which will become important to our discussion later. Additionally, Marr outlined an algorithmic level of analysis, which describes *how* the function is being computed (e.g. the calculator's chip). This is roughly analogous to the Chomskyan 'performance' (although the latter is more heterogeneous; see Jackendoff, 2002). This computational / algorithmic (or competence / performance) distinction is akin to the veteran product / process distinction respectively: the structural description of the output ('product') function is featured on the computational or competence level, while the actual processes involved in a specific task are on the algorithmic or performance level. [Marr also introduced a third level, the 'implementational' (hardware / wetware) level, but this is not relevant to our discussion here.]

The essence of the difference between normative and computational or competence theories is in their respective research questions. As Marr noted, an algorithmic theory asks 'how is…' questions; for example, how a decision is made in various frame contexts. A descriptive competence theory asks 'what is…' questions, for example what the relation is between the negative particle and the verb phrase in AAVE. A normative theory asks evaluative 'ought' questions: 'What ought to be the good use of negation in language?' A normative approach contains an element of evaluation, a sense of 'goodness' and 'badness', 'right' and 'wrong', that is absent from a purely competence account. In short: normative theories are 'ought' type theories; computational theories are 'is' type theories. Note that competence theories and performance theories are both descriptive – what they share is the 'is'.

In conclusion, our position is that the normative and the descriptive functions of competence theories are best kept strictly separate, as they are in mainstream linguistics. At the very least, it is not obvious that norm and competence are one and the same, and we suggest that the burden of proof is on anyone contesting the distinction. We therefore conceptualize competence level explanations – alongside algorithmic level explanations – as descriptive, 'is' type theories, rather than normative, 'ought' type theories. We will argue that failing to distinguish between 'is' and 'ought' inevitably invites a highly controversial type of inference. We now turn to examine this inference and its consequences.

5. Inferring 'ought' from 'is'

Differentiating between normative and competence accounts might not have mattered all that much, were it not for the problem of arbitrating between competing normative accounts. As noted above, normativism has to be selective: where there are alternative systems, only one of them is 'appropriate' (what Stanovich, 1999, calls 'the inappropriate norm argument'). However, with alternative norm and multiple norm paradigms, arbitrating between competing normative systems is both crucial and far from easy. This is where the difference between normative and competence theories becomes critical. Competence theories are descriptive and can hence be supported by descriptive evidence. In contrast, can one support normative theory with descriptive evidence? Can one infer the 'ought' from the 'is'?

The short answer is 'no'. Inferring an 'ought' type conclusion from 'is' type premises is highly controversial, and considered by many authors to be a logical fallacy. First identified by Hume (2000 [originally published 1739/1740]; although cf. MacIntyre, 1959), is-ought inference is made whenever we attempt to derive normative or evaluative conclusion from descriptive premises (although cf. Frankena, 1939; Searle, 1964; Williams, 1985). For example:

Human beings have natural fear of heights.
Therefore, we should not fly in airplanes.

Since the premise has no normative value, inferring a normative conclusion is argued to be fallacious. Is-ought inference is closely related to what is called 'the naturalistic fallacy' (Moore, 1903): deriving ethical norms from natural phenomena, for example deriving ethics from evolution. The term is sometimes extended to any sort of evaluative norm derived from natural observation, and in that sense it overlaps to a great extent with is-ought inference. Our airplane example is problematic both in the is-ought sense and in the naturalistic sense. Note that one can argue that there is an *implicit normative premise*: the belief that we should act according to our natural emotions, including fear. With the premise made explicit, the normative term is included in the premises, and the argument no longer a fallacy. However, indentifying – and justifying – the implicit 'ought' premise can be rather tricky.

We should clarify at this stage, that the is-ought question is a highly polemical one; whether it is always a fallacy is much contested in the philosophical literature (for reviews see Hudson, 1969; Schurz, 1997). However, none of the proposed solutions suggest that is-ought inference is *universally* valid; they typically specify a set of conditions under which it is valid (for example, for constitutive rules only; Searle, 1964). Cases that fall outside these conditions are indisputably invalid. Whether these conditions apply in the case of normativism is moot (Elqayam, 2011), and we propose that the burden of proof is on normativism. We therefore submit that it is preferable to avoid such inference entirely. To do so, we must confine ourselves to competence and not normative theories. In what follows, we will look in detail into two examples of is-ought inference, both made by prominent normativist research programs: Oaksford and Chater's (1998; 2007) rational analysis program, and Stanovich's (1999) early application of the understanding / acceptance principle. We have chosen to focus on these two examples as they are high profile and well respected in the literature. Indeed, we ourselves admire both of these programs in many respects. However, we also contend that each involves evaluative normativist thinking and a form of is-ought inference.

5.1. Oaksford and Chater's Bayesian rational analysis

Since the early nineteen-nineties (Oaksford & Chater, 1991), and culminating in their recent book, *Bayesian rationality* (Oaksford & Chater, 2007), Oaksford and Chater have pioneered a research program that strongly rejects logicism in both its forms, empirical and prescriptive, and endeavors to replace it with another normativist framework, namely Bayesianism. Throughout this period, Oaksford and Chater have advocated in no uncertain terms both empirical and prescriptive Bayesianism – that is to say, the idea that human thinking is both grounded in Bayesian probability and normatively justified by it. Paradoxically, the very rejection of logicism puts Oaksford and Chater at a rather high level of prescriptive normativism. They leave little doubt that their research agenda is fully committed to normativism in its Bayesian form. Adopting Anderson's (1990; 1991) framework of rational analysis, which opts for computational-level task analysis in preference to processing account, they maintain that the evolutionary success of human behavior has to be explained by a computationally adequate normative theory, the basic principles of which are probabilistic. Oaksford and Chater also maintain that the computational level 'must be normatively justified' (Oaksford & Chater, 1998, p. 6). Their argument can be simplified as follows:

Premise 1: People behave in a way that approximates Bayesian probability ('is')
Premise 2: This behavior is successfully adaptive ('is')
Conclusion: Therefore, Bayesian probability is the appropriate normative system ('ought')

In what seems to be a classic is-ought inference, ‘is’ type evidence is brought to bear on ‘ought’ type conclusion (also see Schroyens, 2009). Indeed, Oaksford and Chater (2007) are quite explicit about this: ‘[...] the empirical approach to rationality aims to interpret people’s reasoning behavior so that their reasoning makes sense [...] the formal standards of rationality appropriate for explaining some particular cognitive process or aspect of behavior are not prior to, but rather developed as part of, the explanation of empirical data.’ (p.31). They make a clear distinction between ‘formal’ and ‘everyday’ rationality. While everyday rationality is instrumentally defined by ‘people’s beliefs and actions in specific circumstances’ (2007, p. 19), formal rationality is normatively defined by ‘formal principles of good reasoning’ (2007, p. 21):

> ‘[In] addition to this informal, everyday sense of rationality, […] the concept of rationality also has another root, linked not to human behavior, but to mathematical theories of good reasoning, such as logic and probability. According to these calculi, rationality is defined, in the first instance, in terms of conformity with specific formal principle, rather than in terms of successful behavior in the everyday world.’ (2007, p. 21).

Note how formal rationality is defined in evaluative terms (‘good reasoning’) and contrasted with successful behavior. This seems to be the missing evaluative ‘ought’ link. The evaluative position is then even more clearly laid out in the following:

> ‘[I]f everyday rationality is viewed as basic, assessing rationality appears to be down to intuition. There is a danger here of losing any normative force to the notion of rationality – if rationality is merely conformity to each other's predominant intuitions, then being rational is like a musician being in tune. On this view, rationality has no absolute significance […]. But there is a strong intuition that rationality is not like this at all – that *there is some absolute sense in which some reasoning or decision-making is good, and other reasoning and decision-making is bad*.’ (Oaksford & Chater, 2007, pp. 24-25; italics ours).

With this statement, Oaksford and Chater inject a strong note of evaluation into the debate; they make it quite clear that their normative agenda is evaluative, and that evaluative cannot be boiled down to instrumental. A little further on, they explicitly reject a purely instrumental account of rationality:

> ‘An alternative normative grounding for rationality seems intuitively appealing: good everyday reasoning and decision-making should lead to *successful action*. For example, from an evolutionary perspective, we might define success as inclusive fitness, and argue

that behavior is rational to the degree that it tends to increase inclusive fitness. But now the notion of rationality seems to collapse into a more general notion of adaptiveness.' (ibid, p. 26; italics in original text).

Finally, Oaksford and Chater make a point of arguing that any adaptively rational behavior should be justified in terms of some normative system (1998, pp. 291-297; 2007, pp. 30-31); otherwise, they maintain, its rationality is meaningless.

It seems, then, that what Oaksford and Chater propose is a circle of normativity, in which formal rationality normatively *justifies* everyday rationality (evaluative 'ought'), while everyday rationality provides *empirical evidence* for formal rationality (epistemic 'ought'). With this dual mechanism in place, there seems to be no is-ought inference involved. We have already noted that what appears to be an is-ought inference can be simply enthymatic; if the implicit 'ought' premise is *a priori* filled in, the inference is inarguably valid. This is the route that Oaksford and Chater seem to take. However, whether is-ought inference is indeed avoided is moot. As we have noted earlier (section 1), a normative system is one that is taken as an evaluative 'ought' for human rationality. *A priori* analysis can only show that a theory is well-formed, but, given the multiplicity of well-formed systems and the ensuing arbitration problem, normativism still needs a move from well-formedness to normative status. The latter is not given as a premise; to complete it, Oaksford and Chater use empirical data. Hence, it can still be argued that they draw is-ought inference.

Before concluding this section, we should clarify that our reservations are not with rational analysis as a research program, only with its evaluative ought. Oaksford and Chater's thesis is a complex one, mixing several senses of 'ought'. A significant part of their argument is what we called the directive, or instrumental sense of ought: the thesis that, given specific goals, some computational systems are more useful than others, and that empirical data can help clarify which. As this aspect of their approach is descriptive, we have no argument with it at all.

5.2. The individual differences program of Stanovich and West

Another highly influential research program with emphasis on normative and evaluative concerns is Stanovich and West's dual-system theory, based on systematic analysis of individual differences (Stanovich, 1999; 2004; Stanovich & West, 2000). Their theory is of a type termed 'default-interventionist' by Evans (2008) as indeed is JE's own dual-process theory (Evans, 2006; 2007). Thus we can broadly agree with Stanovich and West's assertion that System 1, the heuristic system, triggers contextualized, belief-laden responses that can be intervened on and altered by System 2, the analytic system, and accept their findings that both the likelihood and nature of such interventions are affected by the cognitive ability of the participants.. Where the difficulty arises is in the interpretation of these findings. In these earlier studies (summarized by Stanovich, 1999) higher ability participants mostly gave more ‘correct’ answers on these tasks, according to the standard norm applied. Thus it appeared that ‘correct’ reasoning required a high probability of intervention and / or a higher quality of reasoning, both associated with high cognitive capacity. In more recent writings, Stanovich has added a number of other pre-conditions for rational reasoning (Stanovich, 2009a; 2009b). He and West have also demonstrated recently that a number of decision biases – as a result - are *not* affected by cognitive ability (Stanovich & West, 2008).

In the earlier work, however, Stanovich directly connected normative theory with computational level analysis, albeit in cautious terms. Prefacing his 1999 book with an extensive review of various theories that depict different levels of analysis, he argued: ‘It is at the intentional level that issues of rationality arise’ (Stanovich, 1999, p. 12). Note that Stanovich merely traced rationality to the intentional level, rather than calling for normative justification of this level in the way Oaksford and Chater do. However, an is-ought inference was still involved in this earlier writing. Its basis was an application of Slovic and Tversky’s (1974) understanding / acceptance principle: the empirical normativism idea that the more one understands the normative principles involved in a specific task, the likelier is one to accept these principles. Hence, cognitively gifted reasoners are likely to endorse the appropriate normative system involved in a specific task. Stanovich also added the converse, prescriptive normativism principle: responses of the more able participants provide the decisive clue for arbitrating between normative systems: whatever they endorse is the appropriate system for a particular task. ‘The direction that performance moves in response to increased understanding provides an empirical clue as to what is the normative model to be applied’ (Stanovich, 1999, p. 63). For example, when higher ability participants provide what is traditionally viewed as the ‘correct’ answer to the Wason selection task (Stanovich & West, 1998), this was taken to imply that deductive logic rather than information gain should be accepted as the appropriate normative system for this problem.

A form of is-ought inference was apparent at this stage, although to some extent moderated by the restricted applicability to elite reasoners (Elqayam, 2003). The ‘is’ evidence was performance by higher ability participants; the ‘ought’ conclusion was the choice of a particular normative system as appropriate. Stanovich actually acknowledged an inherent naturalistic fallacy (Stanovich, 1999, pp. 59-60), although he maintained that a worse version of the same fallacy is made by the camp which regards behavior as *a priori* rational (we concur). He also argued that 'if the theorists discussed so far are actually committing the naturalistic fallacy, then many of the best minds in cognitive science seem to be doing so' (Stanovich, 1999, p. 60). Here, too, we concur – but would point out this did not solve the problem. Indeed, perhaps this *is* the problem.

It is important to note that Stanovich and West themselves no longer use this arbitration strategy, and have discontinued even the use of the term 'normative rationality' (Stanovich & West, 2003). However, the is-ought strategy in Stanovich (1999) still has current influence over the research community. For example, it has recently been extended to support sensitivity to diversity as a normative account for category based induction (Feeney, 2007). In a later phase in the development of Stanovich and West's theory, the focus is on instrumental rationality in the traditional sense of achieving one's goals, and epistemic rationality in the sense of holding well-calibrated beliefs (Stanovich, 2004, 2009b; Stanovich & West, 2003), with which we have no quarrel. However, there are still clear evaluative elements in their approach. While the term 'normative' has been dropped, the term 'error' has not: a recent book (Stanovich, 2009b) presents an extensive discussion of the source of reasoning and decision making errors, implying norms.

It is important to note that we have no argument with Stanovich's (and others') position when examined from the angle of *applied* science. If your object is to improve thinking (rather than to understand it), then you must have criteria for distinguishing good thinking from bad (more on this in section 8).

5.3. Evaluative ought vs directive ought

Having described the is-ought inference in Oaksford and Chater's rational analysis, and in the (still influential) earlier formulation of Stanovich and West's approach, we come now to a crucial test: comparing them. Recall how the arbitration problem poses a major challenge to normativism; it is particularly striking here. While both approaches share an evolutionary agenda, each starts from a completely different evaluative position and draws completely different normative conclusions. Oaksford and Chater's rational analysis, with its adaptationist leanings, starts with the presupposition that evolution optimizes, and that gene-dictated behavior is by definition rational. In contrast, Stanovich and West adopt a view of rationality that is self-described as Meliorist (Stanovich, 1999; 2004; 2009b). That is, they do not believe that people are invariably rational, but rather that they are capable of being so and that this capability can be improved by education and training.

Individual differences in reasoning pose major difficulties for the optimization stance of Oaksford and Chater. Consider the case of the abstract Wason selection task. The early Stanovich and West (e.g., 2000) have argued for logic as the correct normative system because those of *highest ability* solve it in these terms. But these are only about 10-20% of those tested. By contrast, Oaksford and Chater argue that information theory is the correct normative theory of the task because it can account for the *majority* of responses to the problem. So 'is-ought' theorists are in dispute as to what is the 'is' from which to infer the 'ought'.

This is not a chance outcome; we submit that the very nature of is-ought type of research program is bound to lead to these differences. Adaptations *per se* can only provide us with epistemic or at most directive 'oughts'. What happens when two directives clash? This is the case that Stanovich highlights. In a dual-system approach, Systems 1 and 2 may pursue different goals by different mechanisms (Stanovich, 2004; see also Evans, 2010b). We cannot describe a unique standard even for instrumental rationality. When directive 'oughts' conflict, it seems to be evaluative 'oughts' that drive the evaluation for the theoretician. While Oaksford and Chater do not seem to acknowledge that there might be a clash, Stanovich and West do, and their solution is not determined by the empirical data but by evaluative considerations, i.e., the idea that rationality is determined at the *individual* level, giving preference not only to System 2 but its application by those of high intelligence. System 2, for example, is portrayed as an intelligent 'robot' that can and should rebel against the tyranny of the genes which created it (Stanovich, 2004).

6. Normativist research biases

It may seem to some readers, as it did to a referee of an earlier draft of this paper, that we are objecting only to the *style* of research and writing about human thinking, and that our comments have few implications for the substance of such research programs. This is far from the case. In fact, we wish to argue the opposite: that normativism has seriously biased and distorted the ways in which psychologists go about studying thinking, reasoning and decision making (see Table 3). It makes a very substantial difference to how we practice our craft, and to the research questions we ask both on the processing and computational level. A descriptivist approach may free the psychology of reasoning and JDM from these research biases.

Table 3: Normativist research biases in psychology of reasoning and JDM.

Normativist research bias	**What it means**	**Level of analysis**	**Research practice**
Prior rules bias	People have built-in normative systems	Computational / processing	Exclude trained participants; exclude helpful knowledge
Interpretation bias	Responses are presented in terms of normative correctness	Processing	Report responses in terms of their normative correlates; assume normative status equals processing
Clear norms bias	Look for unambiguous norms	Computational	Exclude multiple norm paradigms from psychological enquiry

Let us start with the case of logic and the deduction paradigm (Evans, 2002). The standard practice in the psychology of reasoning, at least until the past decade or so, was as follows. You draw a sample of participants, specifically excluding any who have had formal training in logic. You present them with problems which are either abstract, or designed in such a way that any influence of real world beliefs can only be interpreted as a bias, since it is orthogonal to the logical structure. You then instruct participants to assume that all the information given is true, to base their reasoning only on the information given and to draw only necessary conclusions. However, normally you do not provide them with any kind of instruction or training in logical principles, including that of necessary inference. You then assess performance against standard logical solutions to the problems, and also count as a cognitive bias any source or variance other than the logical structure. This describes the predominant practice in the psychology reasoning over the past 50 years or so, and explains the origins of such terms as belief bias (e.g., Evans, Barston, & Pollard, 1983; and see section 6.1 below) and matching bias (Evans, 1972); the balance has only started to shift in recent years.

We ask readers to reflect on whether the deduction paradigm could have developed this way without the logicist and normativist thinking that preceded it. The argument encouraged by Inhelder and Piaget and numerous philosophers is essentially this: (classical) logic provides the laws of rational thought in all contexts. People are rational. Therefore, logic must be built into people's heads in some innate and *a priori* manner. We call this the *prior rules bias.* This is basically an empirical normativism approach, the idea that thinking reflects a normative system, and it has implications for computational level analysis as well as processing accounts. From this, everything about the deduction paradigm follows, including the use of participants untrained in logic and contexts lacking helpful pragmatic cues. If people are rational, then they should still give the logical answers. Researchers then seem to be astonished when participants get them wrong, in contrast with the remarkable achievements of the human species in many specific fields of endeavor that require advanced reasoning.

Without logicism, the study of rationality in reasoning might have been entirely different. Why on earth, for example, should our notion of rationality exclude *learning*? Why not start with the observation that people can become expert reasoners in law, medicine, science, engineering and so forth, noting that in every case they spend many years in specialized training to achieve this level of expertise? Why not focus on expert reasoning and how it is acquired? But no, we have spent the past half a century instead studying naïve participants with novel problems, resulting in a Kuhnian crisis as the field struggled to throw off the shackles of logicism (Evans, 2002; 2010b; Oaksford & Chater, 1998; 2007). The new paradigm that is emerging utilizes a wide variation of methods which a focus on uncertainty, belief and pragmatic influences on reasoning. However, merely discarding logicism will not resolve the problem. As befits the topic of this paper, there is an as yet unresolved debate about whether the new paradigm requires an alternative normative theory, such as Bayesianism (Evans, 2010a). The prior rules bias is still active – only the proposed rules have changed.

If we examine the study of judgment and decision making, we find that normativism has dictated research strategy in very similar ways. Here too, researchers predominantly assess rationality by testing naïve participants on novel problems, carefully avoiding any instruction in the rules with which they need to reason. The prior rules bias is evident, for example, in the study of Bayesian reasoning. According to Bayes' theorem, posterior probability judgments should reflect a multiplicative function of prior probabilities (i.e., base rates) and diagnostic evidence. Since the pioneering work of Kahneman and Tversky (1972) there has been much concern, and a very large number of research papers, about the finding that people neglect or underweight base rates in making posterior probability judgments (for review see Barbey & Sloman, 2007). But think what is required to get the problem right. The participants must either know or somehow derive from first principles Bayes' theorem (or at least some sort of approximation), since this is never given by the experimenter. They must also perform the mental arithmetic required to multiply the relevant probabilities. Very few are able to do this, except when the information is presented in transparent 'nested sets' clearly demonstrating the relations between superordinate and subordinate sets (Barbey & Sloman, 2007; Cosmides & Tooby, 1996; Evans, Handley, Perham, Over, & Thompson, 2000; Gigerenzer & Hoffrage, 1995). This facilitates a relatively simple mental representation of the problem, enabling its solution by general reasoning.

What kind of test of rationality do standard tests of Bayesian reasoning provide? Why should we expect people to reason with rules they do not possess, lacking the 'mindware' for the task (Stanovich, 2010)? Granted, many studies have shown base rate neglect in expert groups (Koehler, 1996), with the obvious implication that such groups (doctors, lawyers etc) require training in this kind of reasoning. But this is where the enquiry should have begun. Some kinds of statistical reasoning can be learned by general experience, even if it remains domain specific (Nisbett, Krantz, Jepson, & Kunda, 1983), but others cannot, and require rule-based training. But the question of whether people can be rational statistical reasoners must be assessed when appropriate training has been provided. Evolution may have provided 'fast and frugal heuristics', labor-saving rules of the thumb that work well in some circumstances (Gigerenzer & Todd, 1999) but it certainly cannot provide us with the ability to be lawyers, engineers and rocket scientists *without training*.

Normativism has not just affected the methodology of the psychology of reasoning but also the way in which findings are reported and interpreted on the processing level. We will call this the *interpretation bias*. In JE's first book on the topic (Evans, 1982), he argued that we should desist from the practice of reporting logical accuracy in reasoning tasks and instead report what people actually did. This is particularly critical in the study of conditional inference. The standard paradigm focuses on the elimination inferences and tests whether people will endorse each of the four inferences MP, DA, AC and MT (see Table 2). The traditional practice, still quite common in the developmental literature (e.g., Barrouillet, Markovits, & Quinn, 2001), is to score the number of logically correct inferences endorsed, which means adding *yes* answers to MP and MT (valid inferences) to *no* answers for DA and MT (invalid inferences). But this practice is highly interpretative and misleading. From a cognitive point of view, an inference is either drawn or it is not. The interpretation bias leads researchers to equate endorsing one kind of inference with refusing another as though these were similar rather than opposite cognitive processes. Logicist thinking has even led leading advocates of the mental logic theory of reasoning to propose *entirely different mechanisms* to explain the drawing of valid and invalid conditional inferences (e.g., Braine & O'Brien, 1998), the former based on mental rules and the latter on pragmatic implicatures. But the experimental evidence supports no such distinction. For example, DA (invalid) and MT (valid) inferences are prone to exactly the same form of negative conclusion bias or double negation effect (Evans, Newstead, & Byrne, 1993). How could this be if different mechanisms are involved?

The field of judgment and decision making is, if anything, even more prone to interpretation bias than the psychology of reasoning. Encouraged by the discipline of economics, from which the study of rational decision making derived, studies of JDM have focused again and again on conformity to or deviations from normative theory to the exclusion of psychological accounts of what people are actually doing. This may be why dual-process accounts of JDM were, until recently, mostly proposed by those working predominantly in the psychology of reasoning (Evans, 2007a; Evans & Over, 1996a; Stanovich, 1999b; 2010). Fortunately, following the explicit adoption of the theory by Kahneman and Frederick (2002) this is starting to change. However, JDM still lags behind the new paradigm psychology of reasoning in the use of process tracing methods such as protocol analysis, response times, eye-movement tracking and neural imaging, but again we are pleased to see this is now changing. But why has it taken so long for researchers to get focused on the cognitive processes underlying judgment and decision tasks? In a word: normativism.

Even at the purely computational level, normativist research biases may affect the very research puzzles that psychologists select to study. Recall our classification of conflict between formal systems in section 3: single norm paradigms, where there is no conflict; alternative norm paradigms, where an alternative system competes with the standard one; and multiple norm paradigms, where there is multiplicity of formal systems, none of which can be said to have any precedence. Historically, the psychology of reasoning and decision making tended to be biased towards asking research questions drawing on single norm paradigms, although they have a tendency to mutate into alternative norm paradigms as researchers discover or invent alternative norms. The expectation that there will be a single or at least standard normative system is a natural consequence of empirical normativism: the belief that human thought follows a normative system. It is also crucial for prescriptive normativism, since a normative system has to be clearly identified for prescriptive normativism to make any sense. We call this the *clear norms bias.*

When an alternative norm is proposed, heated debate tends to follow, as normativism requires a clear standard. Moreover, the motivation for proposing alternative norms may be the observation that empirical normativism fails with the existing standard. Oaksford and Chater's (1996) account of selection task choices in terms of expected information gain, and the spate of critical notes that followed it, illustrate both aspects. In JDM, Hahn and Warren

(2009) have similarly taken on normative analysis of lay perception of randomness, long perceived as normatively incorrect in the JDM literature. Arguing that such perceptions are normatively correct when one takes into account the ‘finite attentional window’ through which random strings are typically available to working memory, they also partially exonerated the ‘gambler’s fallacy’ (the conviction that consecutive random draws should be balanced), again with the foreseeable flurry of critical notes.

As to multiple norm paradigms, where no agreed normative standard exists, there is corresponding little experimental work. Examples include embedded conditional statements, where ‘no theory has an intuitively adequate account’ (Edgington, 2008), and conditional introduction, as opposed to the much studied elimination inferences (but see Over et al., 2010 for recent discussion). These are by no means trivial paradigms: both have generated a great deal of discussion in philosophical logic (for review see Edgington, 2008). Issues that psychology of reasoning overlooked despite patent philosophical value tend to be multiple norm paradigms. We find this suggestive to say the least.

6.1. Normativist research biases and dual processing: The ‘ought-is’ fallacy

Dual-process and dual-system theories of higher cognition have become increasingly popular in both cognitive and social psychology (Evans, 2003; 2008; Evans & Frankish, 2009; Kahneman & Frederick, 2002; Lieberman, 2007; Sloman, 1996; Smith & DeCoster, 2000; Stanovich, 1999; 2004). We discuss them here to illustrate how normativism has biased and hindered this particular research program.

Dual-process theories postulate two types of processes – heuristic, rapid, parallel preconscious processes, versus analytic, effortful, sequential processes that correlate with general ability. Dual-system theories add the stronger postulate that these processes are anchored in distinct cognitive systems (Evans, 2003), which Stanovich (1999) dubbed ‘System 1’ and ‘System 2’ respectively. Dual-process and dual-system theories can at most be empirically normative to a moderate extent, since the two process cue different responses. Historically, dual-process theories of reasoning and decision making have been used to explain conflict between normatively correct responding and cognitive biases. Evans (1982)’s early two-factor theory of reasoning, for example, proposed that logical and non-logical processes combined in determining behavior. A classic example is the belief bias paradigm in syllogistic reasoning, in which participants have to judge the validity of arguments which are either logically valid or invalid, and have either believable or unbelievable conclusions. Evans, Barston, and Pollard (1983) established that people will both prefer logically valid conclusions (with belief constant) and believable conclusions (with logic constant), which they characterized at the time as a within-participant conflict between *logic* and belief. Stanovich's (1999) earlier research program on individual differences in cognitive performance associated normatively correct responding with high cognitive ability and belief biases with low cognitive ability, with a theoretical account in terms of greater ability for System 2 reasoning in those of high cognitive capacity (although cf. Stanovich, 2009a; 2009b). Similar appeals to System 2 as means of avoiding biases and achieving normatively correct solutions are to be found in other major contributions in the field (e.g. Kahneman & Frederick, 2002; Sloman, 1996). All of which might combine to give the (unfortunate) impression that System 2 is an empirically normativist system - an impeccable mental logic that delivers reliably normative reasoning.

Another source that might have contributed to this impression is Evans and Over (1996)'s commonly cited distinction between two forms of rationality:

Instrumental rationality (*Rationality$_1$*): Thinking, speaking, reasoning, making a decision, or acting in such a way that is generally reliable and efficient for achieving one's goals.

Normative rationality (*Rationality $_2$*): Thinking, speaking, reasoning, making a decision, or acting when one has a reason for what one does sanctioned by a normative theory.

Subsequently in the book, Evans and Over developed a dual-process theory in which they distinguished between implicit and explicit processes. In presenting a dual theory of rationality and dual-process theory of thinking within the same work, Evans and Over (1996) provided a temptation for some readers to confuse the two, even though they explicitly cautioned against making such a direct equation (p.147). Given that their definition of normative rationality involved explicit rule following it follows, of course, that Type 2 processing is *necessary* to achieve it. But nothing in their account implies that it is sufficient. With the dominance of the normativist thinking, however, it is all too easy to substitute sufficient for necessary, and hence to assign a one-to-one relation between Type 2 processing and normative solutions.

The equation of System 1 with bias and System 2 with normatively correct reasoning is in fact a dangerous fallacy. The temptation is to treat correct responses as being *diagnostic* of System 2 processing, and biased responses as diagnostic of System 1 processing, an inference to be found throughout the dual-processing literatures. Note that this fallacy is a special case of interpretation bias (see section 6), and hence of empirical normativism, as it presupposes that System 2 corresponds to a normative system. Since the fallacy involves inferring 'is' (System 2 involvement) from 'ought' (normative responses), we will dub it 'the ought-is fallacy'. The ought-is fallacy is particularly hazardous in paradigms where there are just two alternative answers: one considered correct and one considered a bias, as when base rate and diagnostic information are put into conflict in Bayesian reasoning (see, for example, De Neys & Glumicic, 2008; see also Ferreira, Garcia-Marques, Sherman, & Sherman, 2006 on the limitations of this paradigm in dual-processing research). Using a forced-choice paradigm in this way is open to a number of other possible interpretations, including erroneous (by normativist lights) System 2 reasoning, normatively aligned heuristics, guessing and random error.

The claim that heuristics can lead to effective responding rather than the cognitive biases emphasized in the Tversky and Kahneman tradition (Gilovich, Griffin, & Kahneman, 2002) has been well argued by advocates of fast and frugal heuristics (Gigerenzer, 2007; Gigerenzer & Todd, 1999). In situations where participants have had opportunity for relevant experiential learning, they may also make decisions that are more satisfactory and instrumentally efficient using intuition than when allowed to engage in reflective thinking (Dijksterhuis, Bos, Nordgren, & von Baaren, 2006; Klein, 1999; Reyna, 2004; Wilson et al., 1993). On the other hand, it is not hard to see either that System 2, rule based reasoning, can lead to normative errors. For example, we may have learnt (normatively) bad rules such as the 'law of averages' that people draw upon to justify irrational gambling behavior (Wagenaar, 1988). We may also have good rules (from a normative viewpoint), but process them badly. In a recent review of a range of hypothetical thinking tasks Evans (2007) actually attributes the cognitive biases observed in these tasks equally to heuristic (Type 1) and to analytic (Type 2) processes. Stanovich (2009a; 2009b) has also identified recently a form of System 2 processing, which he calls 'serial associative cognition' which may lead to errors and biases.

What this discussion illustrates is that while dual-process research may appear to assume or even benefit from a form of empirical normativism, in which System 2 (but not System 1) is assumed to generate normatively correct responses, this is far from the case. In fact, dual-process research *suffers* from this form of normativist reasoning. It leads researchers to think that they have an easy shortcut method to identify the type of process from the correctness of the response, when none such is in fact available. This has been recognized implicitly in some recent dual-process accounts of how beliefs influence reasoning (Evans & Handley, 1999; Klauer, Musch, & Naumer, 2000; Verschueren, Schaeken, & d'Ydewalle, 2005). These theories propose *both* that (a) beliefs may influence responding directly through heuristic cues and (b) that beliefs may bias the direction and focus of explicit analytic reasoning. Such theoretical developments would not be possible with the System 2 = normative system mindset.

In reasoning theories, ought-is fallacy seems empirically dissociated from is-ought inference in the sense expounded in section 5 – that is, it is different authors that tend to make the two types of inference. In particular, Stanovich's recent research program emphasizes the System 2 sources of biases, thus not only avoiding ought-is fallacy but explicitly precluding it. This is hardly surprising, for an approach that we have already characterized as relatively low on empirical normativism. While is-ought inference is a special case of prescriptive normativism (thinking should be measured against a normative system), ought-is fallacy, with its assumption that System 2 equals mental logic, is a special case of empirical normativism (thinking reflects a normative system). As we commented earlier, prescriptive normativism is necessary for empirical normativism but by no means sufficient.

In conclusion, normative research biases affect what is studied in psychology of reasoning and JDM, how it is studied and how findings are reported and interpreted, both on the processing and the computational level of explanation. *Prior rules bias has* affected research practice by providing undue focus on untrained participants reasoning with novel problems; *interpretation bias* and its close associate *ought-is fallacy* prompted researchers to analyze psychological processes in terms of their normative correlates; and *clear norms bias* has focused attention on single norm paradigms, or on normative rather empirical arguments when they change into alternative norm paradigms, and arguably sentenced multiple norm paradigms to unwarranted neglect. These biases are highly prevalent and afflict much of the field. While it might be possible to patch them up *ad hoc*, we contend they can be most parsimoniously eliminated with a descriptivist approach, focusing on observing and explaining the thinking and reasoning that people *do*, without the prior concerns about what they *ought* to do.

7. Can we manage without a normative theory?

The previous three sections have reviewed the pitfalls of normativism. First, we have argued that in a quest to solve the thorny arbitration problem, theorists have fallen into the practice of dubious is-ought inference, which in the worst case can lead to circular reasoning: people ought to do whatever it is they actually do! Next, we showed how normativist thinking has biased and constrained the relevant research programs. Illustrating the problem with the case of research on dual processes, we also identified a specific bias which we term 'ought-is' fallacy: the belief that System 2 is responsible for normative responding (and System 1 for errors and biases).

We now seem to be faced with a dilemma. On the one hand, the problems we have identified with normativism make it highly questionable as a meta-theoretical framework for the psychology of reasoning and judgment and decision making (JDM). On the other hand, the long and productive history of normative approaches in reasoning and JDM should give one pause before throwing them overboard. Formal systems such as logic and Bayesianism have provided major incentives and inspiration to countless research paradigms. Popper's logicist philosophy of science was the main motivation behind Wason's selection task and the 2-4-6 task; decision theory motivated Tversky and Kahneman's heuristics and biases program. Can we make do in reasoning and JDM without normative theories altogether?

Evaluative normative considerations are just one way in which formal theories can be useful for psychological theorizing. There is a wide range of possible relations between formal systems and psychological theory, depicted in Figure 2. Formal theories can also *constrain* psychological theorizing; that is, psychological theory can be formed in a way that takes computational level theories into account, and can provide a useful formal language.[4] Formal theories can also *inspire* psychological theory, which can be seen as a special case of weak constraining: a single idea or principle is taken from the formal theory, leaving a wide margin for psychological principles to be developed semi-independently. Psychological theorizing and data can also reflect back on formal theories: to *arbitrate* between formal accounts, either normatively or descriptively; and to judge the psychological *validity* of formal accounts.

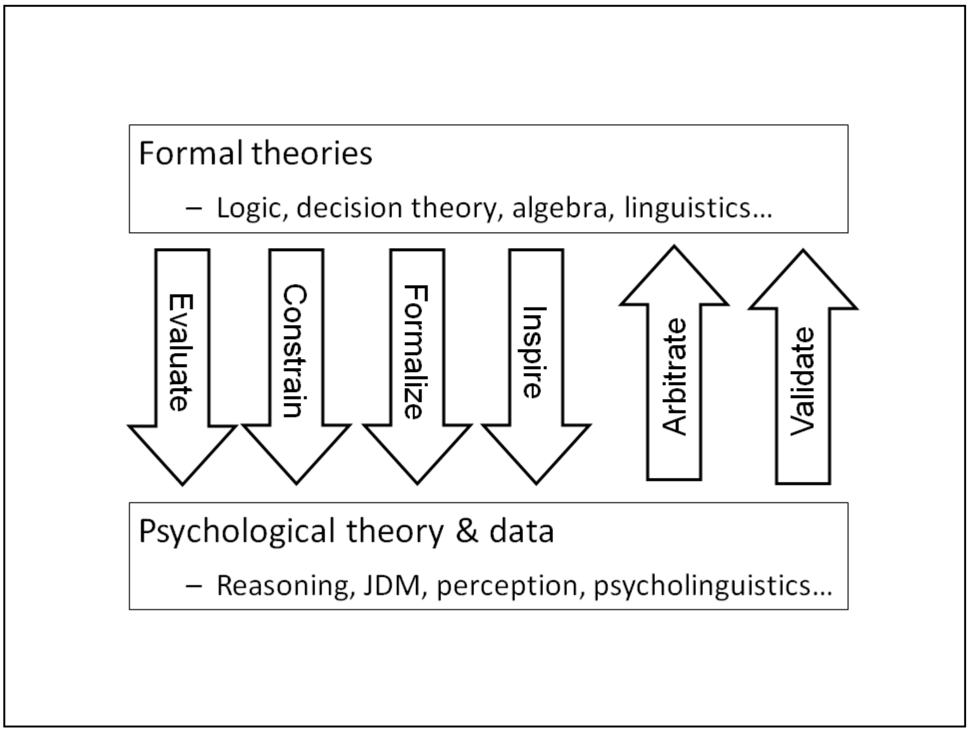

Figure 2: *Models of interaction between formal and psychological theories*

Normativism potentially utilizes almost all the relations shown in Figure 2 except validation: formal theory both inspires and constrains psychological theory. In contrast, with descriptivism, there is considerable variety. Although no theoretical approach seems to be explicitly committed to what we called 'descriptivism', some theoretical approaches can be characterized this way post-hoc. One such approach has been adopted by Gigerenzer and his research group (e.g., Gigerenzer, 2007; Gigerenzer & Selten, 2001; Gigerenzer & Todd, 1999). Gigerenzer and Todd (1999), in introducing their research program on 'fast and frugal' heuristics, appeal first to bounded rationality (Simon, 1982) then to ecological rationality and finally to evolutionary rationality. What they specifically *exclude*, however, is normative rationality. The concept of normative rationality is replaced with the concept of ecological rationality, the rationality of adaptive behavior. In noting the probabilistic revolution that has undermined logicism, they comment that their approach embraces its emphasis on uncertainty '...without sharing its focus on probability theory, either as a description or an attainable norm of human behavior'. Paradoxically perhaps, it is this very rejection of normativism has led Stanovich (1999, pp. 57-58) to ascribe the naturalistic fallacy (see section 5) to Gigerenzer and Todd, maintaining that they link the normative to the descriptive, even if only to reject it.

We have no argument with adaptive rationality *per se*; indeed, we keep to the position expressed by Over and Evans (1997, p. 255-256), which regarded any sort of instrumental rationality (adaptive rationality included) as primary. However, we agree with Stanovich (1999) that Gigerenzer's approach invites a strong is-ought inference: it seems that merely opting for a position of adaptive rationality does not inoculate a theory against it. Furthermore, given the heuristic value of formal theories, we are not convinced that a descriptivist theory should do without formal theories entirely; we argue that many of the relations depicted in Figure 2 can be maintained in a descriptivist approach.

The closest model we have is our preferred theoretical framework of hypothetical thinking theory (HTT), a dual-process framework of human thinking (Evans, 2006; 2007). According to HTT, the main characteristic of analytic (or Type 2) processes is the ability to simulate hypothetical possibilities. HTT proposes that Type 2 processes utilize epistemic mental models. Unlike the semantic mental models of mental model theory (Johnson-Laird & Byrne, 1991; 2002), which represent states of the world, epistemic mental models represent what we *believe* about the world. For this reason, they can include subjective probabilities, causal relations and so on which must be excluded from truth-verifiable semantic models. This emphasis on subjective belief makes HTT explicitly Bayesian (Evans, 2007, p. 33). Here is an interesting point of comparison with Oaksford and Chater's rational analysis program. While both programs are Bayesian (and decision-theoretic) in their treatment of deductive competence, in HTT, Bayesian theory is drawn on for its psychological features: subjectivity, belief, uncertainty. Indeed, Evans (2007) comments that Bayesian philosophy is a 'more credible *descriptive* model' of scientific thinking than the Popperian approach (p. 33, italics ours).

A similar case is the theory of the suppositional conditional, a special case of hypothetical thinking theory (Evans & Over, 2004; see also: Evans, Handley, & Over, 2003; Evans et al., 2005). Strongly influenced by Edgington's philosophical work (e.g., Edgington, 1995; 2003; 2008), it draws on a famous philosophical tenet, the Ramsey test, which argues that, to judge a conditional sentence of the form 'if p then q', we add 'p' hypothetically to our stock of knowledge and evaluate 'q' in this context (Ramsey, 1990 [originally published 1931]). In HTT, The philosophical Ramsey test is translated into a psychological principle of hypothetical thought and linked to Evans's (e.g., 1989) earlier hypothesis that 'if' triggers focus on the conditional's antecedent (i.e., the 'p' part of 'if p then q'). Evans and Over (2004, p. vi) explicitly state that they do not try to answer normative questions. The philosophical foundation is there, but the concerns of the psychological theory are different: for example, it is strongly committed to a psychological dual-processing framework, while remaining uncommitted to any specific philosophical version of the Ramsey test (Evans & Over, 2004, chapter 2; for review of the philosophical literature cf. Edgington, 2008). Evans et al. (2003, p. 323) also refer to the Ramsey test as a source of *inspiration*.

More generally, HTT opts for an (implicitly) descriptivist framework by emphasizing the primacy of psychological processes. Evans (2007) specifically critiques the use of normative rationality in psychological inquiry. He discusses the issue in some detail, commenting that violations of normative theory are not an informative source for theorizing on cognitive processes (p. 108); and creating a clear distinction between normative rationality and analytic thinking (p. 159-161). He also addressed the issue of normative rationality explicitly (p.161), with the comments: 'Normative rationality is essentially a philosophical and not a psychological concept. Analytic reasoning may (or may not) involve following explicit rules as some theorists argue […] but the relation of those rules to formal normative theories cannot form part of our psychological definition of System 2 thinking.' HTT, then, utilizes formal theories as a source of inspiration and even as weakly constraining psychological theorizing, without, however, accepting their normative role. While descriptivism can take different forms, we advocate the type that makes use of formal systems as fully as possible, with HTT as the closest existing realization of it.

In conclusion, we believe that descriptivism is a viable alternative to normativism. It can offer as much as normativism in terms of the heuristic value of formal theories, without the problematic inferences and attendant research biases.

8. Final thoughts and conclusions

It is not our purpose to exclude normativism entirely from scientific endeavor. There is a need for research in education, planning, policy development and so on, in all of which norms play a crucial role. The Meliorist position is a strong case in point, both the version advocated so powerfully by the individual differences research program of Stanovich and West (2000; Stanovich, 1999; 2004; 2009b), and the version put forward by Baron (e.g., 2008). Such authors wish to find ways improve people’s reasoning and decision-making and therefore require some standard definition of what it means to be rational. This is an entirely different enterprise from the scientific investigation of the cognitive processes. Take the case of gambling. Gambling on games of chance that guarantee expected losses is commonplace in Western cultures and as many as 1-2% of the population may become pathological gamblers (Raylu & Oei, 2002). Such behavior appears to be neither normatively nor instrumentally rational (although cf. Hahn & Warren, 2009, on the gambler's fallacy), if you define rationality as behaving in such a way as to achieve financial gain.

Gambling is interesting in part because of its apparent irrationality. However, this is not because we (scientists) regard gamblers as immoral people or sadly lacking in the theory of Bayesian decision making. What is striking is the apparent lack of *instrumental* rationality – people persist in behaviors that bring them more losses than gains. Among many interesting psychological findings in this field is that people hold false theories about chance and probability that reinforce their gambling habits (Wagenaar, 1988), although it is unclear whether these are causal or confabulatory (Evans & Coventry, 2006). In our view, normative theory has no role to play in the study of gambling behavior, except perhaps to motivate interest in the topic. We need to understand what people are doing and why, rather than discussing what they are *not* doing. By contrast, normative theory has a clear role in the *treatment* of pathological gambling. Disabusing people of their false beliefs and teaching the normative theory of probability has been shown to be an effective form of cognitive therapy for problem gamblers (Raylu & Oei, 2002). The researchers who discovered this were clearly judgmental: they believed that pathological gambling was a bad thing, a severe problem for the individual who should therefore be helped to stop doing it. We find this approach suitable – necessary, even – in applied research, but we see it as wholly inappropriate in basic theoretical research.

Meliorism does not have to rely on normative theories. As the gambling case shows, it is sufficient to focus on instrumental irrationality. While the earlier Meliorist research program of Stanovich and West did employ normative theories, we have no argument with norms used to facilitate performance, so long as their evaluative aspect is acknowledged to precede research rather than follow from it. We note also that Baron's Meliorist position is explicitly instrumental: 'The best kind of thinking, which we shall call *rational thinking,* is whatever kind of thinking best helps people achieve their goals' (2008, p. 61); 'rational decisions can be defined relative to a person at a given time, with a given set of beliefs and goals' (ibid., p. 63).

For decades, the normativist agenda has reigned supreme in the psychology of human thinking – deductive reasoning and decision making – and it is still pretty much a dominant paradigm. However, its tenets and practices can no longer be taken for granted. The controversial inferences we have pointed out –is-ought and ought-is respectively – are both a result of combining the descriptive with the prescriptive, as are several biases that affect the conduct and interpretation of research work. We do not deny that, as with gambling, normative thinking can attract interest in a phenomenon. In general, with a combination of evolutionary programming and instrumental learning we would expect people to achieve most of their goals, most of their time. If we observe, for example, that many people buy high and sell low in the stock markets, then we are more likely to study their behavior than if it were the other way around. Instrumentalism is the default behavior and easily explained. The converse behavior is more interesting, we agree. It may also be important to understand it from an applied perspective.

In conclusion, our argument is that psychology of human thinking would be better off with a descriptivist agenda. Normativism has played out its role in the history of the research on human thinking. The descriptivist approach views theoretical research on reasoning and decision making research as descriptive rather than evaluative. The object is not to judge such behavior, but to understand and predict it, using all relevant theoretical and methodological tools. Formal and computational models have an important role to play in this without being evaluative or being used to justify is-ought inference. We contend that evaluative considerations need only be invoked in educational and other applied research where the object is to improve human thinking and performance. A shift away from normativism and towards descriptivism has played a crucial role in the development of linguistics as a mature science. A similar change of direction may prove just as beneficial for the study of human reasoning and decision making.

Acknowledgements

This work was partially supported by research travel funding from De Montfort University to the first author. A preliminary sketch of some the arguments in this paper has been presented in the Cambridge Meeting on Dual Processing, July 2006, and the European Cognitive Science conference in Delphi, May 2007. We are grateful to Paul Bloom, Nick Chater, Keith Frankish, Ulrike Hahn, Mike Oaksford, David Over, and two anonymous reviewers, for valuable comments on earlier drafts; John Clibbens, Vinod Goel, Ulrike Hahn, Philip Kargopoulos, Gernot Kleiter, Niki Pfeifer, Keith Stanovich, Rich West and members of the Plymouth Thinking and Reasoning group for useful discussions.

Appendix: Some terminological notes

It is difficult to tread anywhere in psychology, philosophy and cognitive science without disturbing the ghosts of previous terminological usage. Rather than encumber the reader with a host of terminological footnotes, we pool them together here.

Logic and logicism. Throughout this paper the term ‘logicism’ is used in its meaning in psychology of reasoning, i.e., the idea that classical logic is both a descriptor and a normative standard of human thinking. This is distinct from the logicism of philosophical logic, i.e. the view that mathematics can be reduced to logic (e.g., Whitehead & Russell, 1962). The common denominator is that logic is conceived as primary. Note, too, that for the sake of simplicity, we sometimes resort to ‘logic’ as short for ‘classical, extensional, bivalent, monotonic logic’ (unless stated otherwise, as we do in section 3).

Bayesianism. The term ‘Bayesianism’ already has a prevalent use in statistics and the psychology of judgment and decision making, referring to the subjectivist approach to probability, in contrast to the conventional or frequentist paradigm (Howson & Urbach, 1993). This approach can be applied normatively or descriptively, but most of our discussion touches on Bayesianism as a norm, so when we refer to ‘Bayesianism’ it should be taken in the normative sense unless noted otherwise. We discuss the difference between normative and descriptive approaches to Bayesianism in section 7.

Descriptive. We use the term ‘descriptive’ as the contrastive of ‘normative’. The term has another sense in the psychological literature, which can best be rendered as ‘a-theoretical’ –

that is, description of observational phenomena without attempt for theoretical analysis. This is not the sense we mean for ‘descriptive’; we refer to descriptive *theories*. For example, prospect theory (Kahneman & Tversky, 1979) is explicitly descriptive in the sense of being non-normative, but certainly not in the sense of being non-theoretical.

Descriptivism. The term ‘descriptivism’ has been used by Hare (1969; 1993) to denote more or less the opposite of what we mean by descriptivism. Hare follows Austin (1961), who coined the term ‘descriptive fallacy’ to denote the error of assuming that all language is descriptive or truth-functional; descriptivism for Hare means the systematic application of the descriptive fallacy. Of course, this is not what we mean by descriptivism. For us, descriptivism means, among other things, avoiding the descriptive fallacy rather than falling prey to it. We do so by identifying which terms are descriptive and which are deontic, and concentrating on the former in psychology of human thought.

Notes

[1] We owe this point to Keith Frankish.

[2] We owe this point to Mike Oaksford.

[3] We owe this point to David Over.

[4] We owe this point to Gernot Kleiter and Niki Pfeifer.

References

Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, N.J.: Erlbaum.

Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences, 14,* 471-517.

Austin, J. L. (1961). Other minds. In *Philosophical papers* (pp. 44-84). Oxford, UK: Oxford University Press.

Barbey, A. K. & Sloman, S. A. (2007). Base-rate respect: From ecological rationality to dual processes. *Behavioral and Brain Sciences,* 241-254.

Baron, J. (2008). *Thinking and Deciding (4th Edition)*. Cambridge: Cambridge University Press.

Barrouillet, P., Markovits, H., & Quinn, S. (2001). Developmental and content effects in reasoning with causal conditionals. *Journal of Experimental Child Psychology, 81,* 235-248.

Braine, M. D. S. & O'Brien, D. P. (1998). The theory of mental-propositional logic: Description and illustration. In M.D.S.Braine & D. P. O'Brien (Eds.), *Mental logic* (pp. 79-89). Mahway, NJ: Lawrence Erlbaum Associates.

Byrne, R. M. J. & Handley, S. H. (1997). Reasoning strategies for suppositional deductions. *Cognition, 62,* 49.

Byrne, R. M. J., Handley, S. J., & Johnson-Laird, P. N. (1995). Reasoning from suppositions. *The Quarterly Journal of Experimental Psychology, 48A,* 915-944.

Cohen, L. J. (1981). Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences, 4,* 317-370.

Cohen, L. J. (1982). Are people programmed to commit fallacies? Further thought about the interpretation of data on judgement. *Journal for the Theory of Social Behaviour, 12,* 251-274.

De Neys, W. & Glumicic, T. (2008). Conflict monitoring in dual process theories of thinking. *Cognition, 106,* 1248-1299.

De Saussure, F. (1966). *Course in general linguistics*. (original publication 1916). New York, NY: McGraw-Hill.

Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: MIT Press.

Dennett, D. C. (2008). Fun and Games in Fantasyland. *Mind and Language, 23,* 25-31.

Edgington, D. (1995). On conditionals. *Mind, 104,* 235-329.

Edgington, D. (2001). Conditionals. In E.N.Zalta (Ed.), *Stanford Encyclopedia of Philosophy (http://plato.stanford/edu/entries/conditionals/)* (Stanford: Stanford University.

Edgington, D. (2003). What if? Questions about conditionals. *Mind & Language, 18,* 380-401.

Edgington, D. (2008). Conditionals. In E.N.Zalta (Ed.), *Stanford Encyclopedia of Philosophy (http://plato.stanford.edu/archives/win2008/entries/conditionals/)* (pp. xx). Stanford: Stanford University.

Elqayam, S. (2003). Norm, error and the structure of rationality: The case study of the knight-knave paradigm. *Semiotica, 147,* 265-289.

Elqayam, S. (2006). The collapse illusion effect: a pragmatic-semantic illusion of truth and paradox. *Thinking and Reasoning, 12,* 180.

Elqayam, S. (2011). Grounded rationality: A relativist framework for normative rationality. In K.I.Manktelow, D. E. Over, & S. Elqayam (Eds.), *The Science of Reason: A Festschrift in Honour of Jonathan St.B.T. Evans* (pp. 397-420). Hove, UK: Psychology Press.

Evans, J. St. B. T. (1972). Interpretation and matching bias in a reasoning task. *Quarterly Journal of Experimental Psychology, 24,* 193-199.

Evans, J. St. B. T. (1982). *The psychology of deductive reasoning*. London: Routledge.

Evans, J. St. B. T. (1989). *Bias in Human Reasoning: Causes and Consequences*. Brighton: Erlbaum.

Evans, J. St. B. T. (1993). Bias and rationality. In K.I.Manktelow & D. E. Over (Eds.), *Rationality: Psychological and Philosophical Perspectives* (pp. 6-30). London: Routledge.

Evans, J. St. B. T. (2002). Logic and human reasoning: An assessment of the deduction paradigm. *Psychological Bulletin, 128,* 978-996.

Evans, J. St. B. T. (2003). In two minds: Dual process accounts of reasoning. *Trends in Cognitive Sciences, 7,* 454-459.

Evans, J. St. B. T. (2006). The heuristic-analytic theory of reasoning: Extension and evaluation. *Psychonomic Bulletin & Review, 13,* 378-395.

Evans, J. St. B. T. (2007). *Hypothetical thinking: Dual processes in reasoning and judgement*. Hove: Psychology Press.

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. (2010a). Reasoning. In D.Reisberg (Ed.), *The oxford handbook of cognitive psychology* (pp. xx). New Yrok: Oxford University Press.

Evans, J. St. B. T. (2010b). *Thinking twice: Two minds in one brain*. Oxford: Oxford University Press.

Evans, J. St. B. T., Barston, J. L., & Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory and Cognition, 11,* 295-306.

Evans, J. St. B. T. & Coventry, K. (2006). A Dual-Process Approach to Behavioral Addiction: The Case of Gambling. In R.W.Wiers, A. W. Stacy, R. W. Wiers, & A. W. Stacy (Eds.), *Handbook of implicit cognition and addiction* (pp. 29-43). Thousand Oaks, CA US: Sage Publications, Inc.

Evans, J. St. B. T. & Frankish, K. (2009). *(Eds.) In two minds: Dual processes and beyond*. Oxford, UK: Oxford University Press.

Evans, J. St. B. T., Handley, S. H., & Over, D. E. (2003). Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory and Cognition, 29,* 321-355.

Evans, J. St. B. T., Newstead, S. E., & Byrne, R. M. J. (1993). *Human Reasoning: The Psychology of Deduction*. Hove & London: Erlbaum.

Evans, J. St. B. T. & Over, D. E. (1996). *Rationality and Reasoning*. Hove: Psychology Press.

Evans, J. St. B. T. & Over, D. E. (2004). *If*. Oxford: Oxford University Press.

Evans, J. St. B. T., Over, D. E., & Handley, S. J. (2005). Suppositions, extensionality, and conditionals: A critique of the mental model theory of Johnson-Laird & Byrne (2002). *Psychological Review, 112,* 1040-1052.

Feeney, A. (2007). Individual differences, dual processes, and induction. In A.Feeney & E. Heit (Eds.), *Inductive reasoning* (pp. 302-327). Cambridge: Cambridge University Press.

Ferreira, M. B., Garcia-Marques, L., Sherman, S. J., & Sherman, J. W. (2006). Automatic and controlled components of judgment and decision making. *Journal of Personality and Social Psychology, 91,* 797-813.

Fodor, J. A. (2008). Against Darwinism. *Mind and Language, 23,* 1-24.

Frankena, W. (1939). The naturalistic fallacy. *Mind, 48,* 464-477.

Gigerenzer, G. (1991). How to make cognitive illusions disappear: Beyond 'heuristics and biases'. In W.Stroebe & M. Hewstone (Eds.), *European review of social psychology* (pp. 83-115). Chichester: Wiley.

Gigerenzer, G. (2007). *Gut feelings: The intelligence of the unconscious*. New York, NY: Viking.

Gigerenzer, G. & Selten, R. (2001). *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.

Gigerenzer, G. & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York & Oxford: Oxford University Press.

Gilovich, T., Griffin, D., & Kahneman, D. (2002). *Heuristics and biases: The psychology of intuitive judgement*. Cambridge: Cambridge University Press.

Gottwald, S. (2001). *A Treatise on Many-Valued Logics (Studies in Logic and Computation, vol. 9)*. Baldock: Research Studies Press Ltd.

Gould, S. J. & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Scoiety of London, Series B, 205,* 581-598.

Hahn, U. & Warren, P. A. (2009). Perceptions of randomness: Why three heads are better than four. *Psychological Review, 116,* 454-461.

Hare, R. M. (1969). Descriptivism. In W.D.Hudson (Ed.), *The is-ought question: A collection of papers on the central problem in moral philosophy* (pp. 240-258). London: Macmillan.

Hare, R. M. (1993). *Essays in ethical theory*. New York, NY: Oxford University Press.

Harris, R. (1980). *The language-makers*. London: Duckworth.

Harris, R. (1981). *The language myth*. London: Duckworth.

Howson, C. & Urbach, P. (1993). *Scientific reasoning*. (2nd edition). Chicago: Open Court.

Hudson, W. D. (1969). *(Ed.) The is-ought question: A collection of papers on the central problem in moral philosophy*. London: Macmillan.

Hume, D. (2000). *A treatise on human nature (Original publication date 1739-1740)*. Oxford: Clarendon Press.

Inhelder, B. & Piaget, J. (1958). *The growth of logical thinking*. New York: Basic Books.

Jackendoff, R. (2002). *Foundations of language: Brain, meaning, grammar, evolution*. Oxford; New York: Oxford University Press.

Johnson-Laird, P. N. & Byrne, R. M. J. (1991). *Deduction*. Hove & London: Erlbaum.

Johnson-Laird, P. N. & Byrne, R. M. J. (2002). Conditionals: a theory of meaning, pragmatics and inference. *Psychological Review, 109,* 646-678.

Kahneman, D. & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgement. In T.Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgement* (pp. 49-81). Cambridge: Cambridge University Press.

Kahneman, D. & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology, 3,* 430-454.

Kahneman, D. & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263-291.

Koehler, J. J. (1996). The base rate fallacy reconsidered: Descriptive, normative and methodological challenges. *Behavioral and Brain Sciences, 19,* 1-53.

Kripke, S. (1975). Outline of a Theory of Truth. *Journal of Philosophy,* ***72****,* 690-716.

Lewis, G. L. (1999). *The Turkish language reform: A catastrophic success*. Oxford University Press, USA.

Lieberman, M. D. (2007). Social Cognitive Neuroscience: A Review of Core Processes. *Annual Review of Psychology, 58,* 259-289.

Lopes, L. L. (1991). The rhetoric of irrationality. *Theory and Psychology, 1,* 65-82.

MacIntyre, A. C. (1959). Hume on 'is' and 'ought'. *Philosophical Review, 68,* 451-468.

Manktelow, K. I. & Over, D. E. (1991). Social roles and utilities in reasoning with deontic conditionals. *Cognition, 39,* 85-105.

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: Freeman.

Martin, R. L. (1984). *(Ed.) Recent essays on truth and the Liar paradox*. Oxford, UK: Oxford Univerity Press.

McGee, V. (1985). A counterexample to modus ponens. *Journal of Philosophy, 82,* 462-471.

Millikan, R. G. (1984). *Language, thought, and other biological categories : new foundations for realism*. Cambridge, Mass.: MIT Press.

Millikan, R. G. (1995). *White Queen psychology and other essays for Alice*. Cambridge, Mass.: MIT Press.

Millikan, R. G. (1996). Pushmi-pullyu representations. In L.May, L. Friedman, & A. Clark (Eds.), *Minds and morals* (Cambridge, MA: MIT Press.

Moore, G. E. (1903). *Principia ethica*. New York, NY: Cambridge University Press.

Nickerson, R. S. (2008). *Aspects of rationality: Reflections on what it means to be rational and whether we are*. New York: Psychology Press.

Nisbett, R. E., Krantz, D. H., Jepson, D. H., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review, 90,* 339-363.

Nozick, R. (1993). *The nature of rationality*. Princeton, N.J.: Princeton University Press.

Oaksford, M. & Chater, N. (1991). Against logicist cognitive science. *Mind & Language, 6,* 1-38.

Oaksford, M. & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review, 101,* 608-631.

Oaksford, M. & Chater, N. (1996). Rational explanation of the selection task. *Psychological Review, 103,* 381-391.

Oaksford, M. & Chater, N. (1998). *Rationality in an Uncertain World*. Hove, UK: Psychology Press.

Oaksford, M. & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford: Oxford University Press.

Over, D. E. (2007). Content-independent conditional inference. In M.J.Roberts (Ed.), *Integrating the mind: Domain general versus domain specific processes in higher cognition* (Hove: Psychology Press.

Over, D. E., Evans, J. S. T., & Elqayam, S. (2010). Conditionals and non-constructive reasoning. In M.Oaksford & N. Chater (Eds.), *Cognition and conditionals: Probability and logic in human thinking* (pp. 135-151). Oxford, UK: Oxford University Press.

Over, D. E. & Evans, J. St. B. T. (1997). Two cheers for deductive competence. *Current Psychology of Cognition, 16,* 255-278.

Pfeifer, N. & Kleiter, G. D. (2011). Uncertain deductive reasoning. In K.I.Manktelow, D. E. Over, & S. Elqayam (Eds.), *The Science of Reason: A Festschrift in Honour of Jonathan St.B.T. Evans* (pp. 145-166). Hove, UK: Psychology Press.

Ramsey, F. P. (1990). General propositions and causality (original publication, 1931). In D.H.Mellor (Ed.), *Philosophical papers* (pp. 145-163). Cambridge: Cambridge University Press.

Raylu, N. & Oei, T. P. S. (2002). Pathological gambling:: A comprehensive review. *Clinical psychology review, 22,* 1009-1061.

Rescher, N. (1969). *Many-valued logics*. New York, NY: McGraw-Hill.

Rips, L. J. (1989). The psychology of knights and knaves. *Cognition, 31,* 85-116.

Savage, L. J. (1954). *The Foundations of Statistics*. New York: Wiley.

Schroyens, W. (2009). On is an ought: Levels of analysis and the descriptive versus normative analysis of human reasoning. *Behavioral and Brain Sciences, 32,* 101-102.

Schroyens, W. (2010). Logic and/in psychology: The paradoxes of material implication and psychologism in the cognitive science of human reasoning. In M.Oaksford & N. Chater (Eds.), *Cognition and conditionals: Probability and logic in human thinking* (Oxford, UK: Oxford University Press.

Schroyens, W., Schaeken, W., & d'Ydewalle, G. (1999). Error and bias in meta-propositional reasoning: A case of the mental model theory. *Thinking and Reasoning, 5,* 65.

Schurz, G. (1997). *The is-ought problem: An investigation in philosophical logic*. Dordrecht: Kluwer.

Searle, J. R. (1964). How to derive 'ought' from 'is'. *Philosophical Review, 73,* 43-58.

Simon, H. A. (1982). *Models of bounded rationality.* Cambridge, MA.: MIT Press.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119,* 3-22.

Slovic, P. & Tversky, A. (1974). Who accepts Savage's axiom? *Behavioral Science, 19,* 364-371.

Smith, E. R. & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review, 4,* 108-131.

Stanovich, K. E. (1999). *Who is Rational? Studies of Individual Differences in Reasoning*. Mahway, NJ: Lawrence Elrbaum Associates.

Stanovich, K. E. (2004). *The robot's rebellion: Finding meaning in the age of Darwin*. Chicago: Chicago University Press.

Stanovich, K. E. (2009a). Distinguishing the reflective, algorithmic, and autonomous minds: Is it time for a tri-process theory? In J.St.B.T.Evans & K. Frankish (Eds.), *In two minds: Dual processes and beyond* (pp. 55-88). Oxford, UK: Oxford University Press.

Stanovich, K. E. (2009b). *What intelligence tests miss: The psychology of rational thought*. New Haven; London: Yale University Press.

Stanovich, K. E. & West, R. F. (1998). Cognitive ability and variation in selection task performance. *Thinking and Reasoning, 4,* 193-230.

Stanovich, K. E. & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate. *Behavioral and Brain Sciences, 23,* 645-726.

Stanovich, K. E. & West, R. F. (2003). The rationality debate as a progressive research program. *Behavioral and Brain Sciences, 26,* 531-534.

Stanovich, K. E. (2010). *Decision making and rationality in the modern world*. New York: Oxford University Press.

Stein, E. (1996). *Without good reason: The rationality debate in philosophy and cognitive science*. Oxford: Oxford University Press.

Stenning, K. & van Lambalgen, M. (2008). *Human reasoning and cognitive science*. Cambridge, Mass.: MIT Press.

Stich, S. P. (1990). *The fragmentation of reason: Preface to a pragmatic theory of cognitive evaluation*. Cambridge, Mass.: MIT Press.

von Neumann, J. & Morgenstern, O. (1947). *Theory of games and economic behavior*. (2nd ed.) Princeton, N.J.: Princeton University Press.

Wagenaar, W. A. (1988). *Paradoxes of Gambling Behaviour*. Hove & London: Erlbaum.

Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology, 12,* 12-40.

Wason, P. C. (1966). Reasoning. In B.M.Foss (Ed.), *New Horizons in Psychology I* (pp. 106-137). Harmandsworth: Penguin.

Wason, P. C. & Brooks, P. G. (1979). THOG: the anatomy of a problem. *Psychological Research, 41,* 79-90.

Whitehead, A. N. & Russell, B. (1962). *Principia Mathematica (original publication, 1910)*. Cambridge: Cambridge University Press.

Williams, B. (1985). *Ethics and the limits of philosophy*. Cambridge, Mass.: Harvard University Press.