

Subtracting “ought” from “is”: Descriptivism versus normativism in the study of 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, United Kingdom

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, United Kingdom

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 a 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 programmes 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 the 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 judgement. 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, such as the Advertising Standards Authority in the United Kingdom,

SHIRA ELQAYAM is a Senior Lecturer in Cognitive Psychology in De Montfort University, United Kingdom. She has published theoretical as well as experimental work in human reasoning and rationality, and has held several research grants to study these topics. She is currently working on a psychological theory of inference from “is” to “ought.”

JONATHAN EVANS is Professor Emeritus of Cognitive Psychology at the University of Plymouth, United Kingdom. He is author or co-author of eight books and more than 150 scientific publications on the psychology of thinking, reasoning, and decision making. His recent work has focused on (a) the psychology of conditionals and (b) the development, review, and critical discussion of dual-processing models of higher cognition. His most recent book is *Thinking Twice: Two Minds in one Brain*, published in 2010 by Oxford University Press.

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, judgement, and decision making – is rife with normative considerations. In the study of human reasoning, these considerations 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 1960s and ’70s 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; 1998a; 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 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 system.

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.

Here *S* is a formal normative system such as logic (classical or otherwise), Bayesian probability, or decision theory. Note that 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 programmes can be said to be normativist. For example, much of the judgement 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 behaviour.

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 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 coordinates 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. Two notable examples are, as already indicated, Wason’s rejection of empirical logicism while continuing to uphold prescriptive logicism (Evans 2002), and the heuristics and biases programme 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 that people should be rational have proposed alternative normative systems (e.g., Cohen 1981).

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 hardcore positions which consider thought processes to be isomorphic to the normative system (e.g., Inhelder and Piaget’s formal

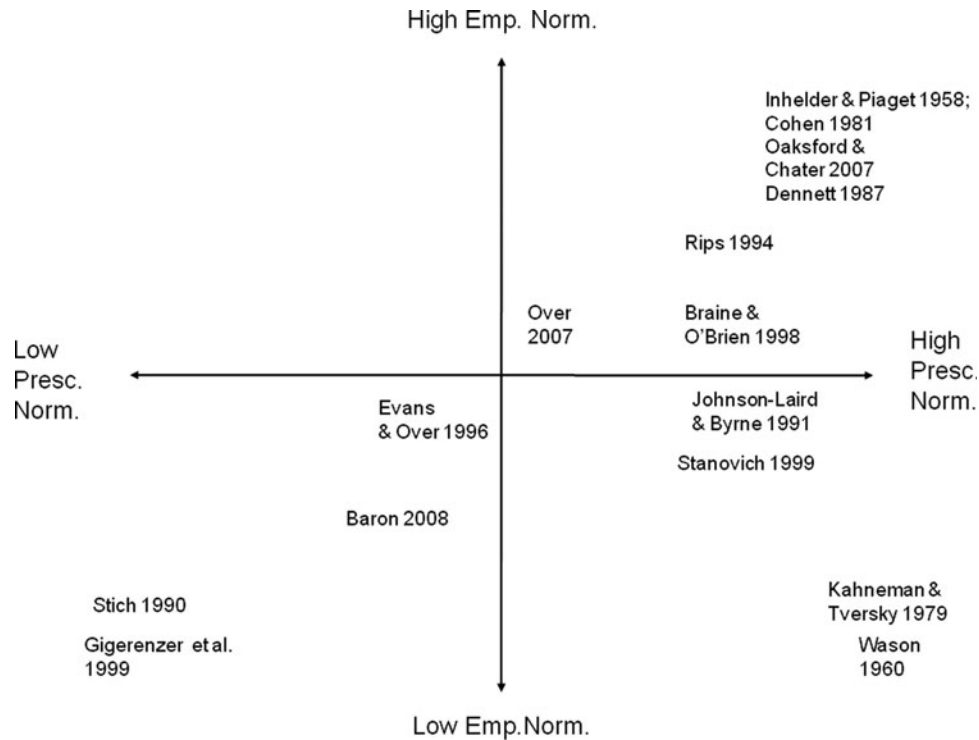


Figure 1. Two-vector normativist space with sample references.

Note: Emp. Norm. = *Empirical normativism*: Thinking reflects normative system, *S*.

Presc. 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 programme is identified in the figure by a single sample reference.)

operations), to positions which reject the normative system entirely. For example, Gigerenzer (e.g., Gigerenzer et al. 1999) famously repudiates any form of normative system, arguing that heuristic rules of 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 as 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 programmes as Oaksford and Chater’s rational analysis, which focuses on modal responses as a source for normative evaluations (Oaksford & Chater 1998a; 2007), and the earlier phase of Stanovich’s individual differences programme (Stanovich 1999), which focused on the normatively superior performance of cognitively able participants. We discuss the problems with this approach in section 4, and examine these research programmes 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 1998a; 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 2000b). 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 programmes 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 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 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 – that is, 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 programmes 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 programmes in the study of 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, judgement, 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: Behaviour that is adaptive within the constraints of cognitive and biological capacity.

Ecological rationality: Behaviour that is adapted to the environment in which the organism is operating.

Evolutionary rationality: Behaviour 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. Instead, 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 Ruth Millikan; e.g., Millikan 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 *must* and *should*, 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.

Meanings (1) and (2) are deontic: they express evaluation and obligation; meaning (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 instrumental³ – for example, 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, behaviour 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 behaviour. Insofar as a research programme 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 programme (see sect. 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 (1998a; 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, whereas 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 behaviour 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; 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, in contrast to 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.

With one established norm and no conflict, *single-norm paradigms* seem to offer the prototypical normativist situation. In other cognitive research domains, single-norm paradigms are indeed both basic and ubiquitous. Thus, in memory, in signal detection, and 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.

The conditional elimination inferences constitute a single-norm paradigm: Regardless of one’s theoretical position, MP (Modus Ponens) and MT (Modus Tollens) are generally considered valid types of inference, whereas DA (Denial of the Antecedent) and AC (Affirmation of the Consequent) 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 (again, 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 et al. 2005; Oaksford & Chater 2007). Hence, when participants judge the paradoxes as invalid (Pfeifer & Kleiter 2011), mental model theory judges their response as erroneous, whereas probabilistic approaches regard the same response as perfectly normative (and see Over et al. [2010] for discussion of another type of conditional introduction inference with conflicting normative judgements).

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 the psychology of reasoning and JDM (judgement 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 behaviour 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] review some recent such developments in JDM). The longer a paradigm is studied, the more it tends to have alternative normative systems proposed.

Table 1. *The three types of normative conflict*

Type	Conflict/No conflict	Number/Type of norms involved	Example(s)
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	Metaduction

Table 2. *Types of conditional inference*

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

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 are 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). Only about 10% of participants (of higher IQ) typically find the standard normative solution: *p* and *not-q* (which in this example would be 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 appears to be standard. For example, consider the reasoning literature on metaduction (e.g., Byrne & Handley 1997; Byrne et al. 1995; Elqayam 2006; Rips 1989; Schroyens et al. 1999). In this paradigm, participants are presented with the Island of Knights and Knaves, whose inhabitants are either knaves (liars) or knights (truth-tellers). The task is to identify the speakers based on their statements. It is generally (albeit implicitly) assumed in the metaduction 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 this sentence: “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, as the latter depends on an agreed norm for assessment. Psychologists can, of course, and do get involved in arguments about which norm 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 foregoing discussion, the reader might have formed the impression that we reject formal systems entirely in favour 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 a 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; 2008); 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 (1916/1966) clearly separates descriptive from normative accounts in favour 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 of centuries of Arabic influence (Lewis 1999), was grounded in nationalist normativism.

There are quite a few categorizations of levels of inquiry 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 (e.g., Chomsky 1965), which is to say, a structural description of abstract knowledge that is

quite value-free (although cf. Harris 1980; 1981). “Competence” is not intended to be contrasted with “incompetence,” but rather with *performance*, that is, 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 (as we have seen in sect. 2) has implications for our discussion. 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/competence theories is in their respective research questions. As Marr noted, an algorithmic theory asks “how *is* . . .” questions; for example, how is a decision made in various frame contexts. A descriptive competence theory asks “what *is* . . .” questions; for example, what is the relation 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,” of “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 (1739–1740/2000; although cf. MacIntyre 1959), is-ought inference is made whenever we attempt to derive a 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 implicit premise made explicit as a second premise, the normative term is included in the premises, and the argument no longer a fallacy. However, identifying – 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 suggests that is-ought inference is *universally* valid; solutions typically specify a set of conditions under which it is valid (e.g., 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 programmes: Oaksford and Chater’s (1998a; 2007) rational analysis programme, and Stanovich and West’s (Stanovich 1999; Stanovich & West 2000b) individual differences research programme. We have chosen to focus on these two examples because they are high profile and well respected in the literature. Indeed, we ourselves admire both of these programmes 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 1990s (Oaksford & Chater 1991), and culminating in their recent book, *Bayesian Rationality* (2007),

Oaksford and Chater have pioneered a research programme that strongly rejects logicism in both its forms, empirical and prescriptive, and endeavours 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 behaviour 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 1998a, p. 6). Their argument can be simplified as follows:

Premise 1: People behave in a way that approximates Bayesian probability (“is”).

Premise 2: This behaviour 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 an “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 behaviour so that their reasoning makes sense ... [T]he formal standards of rationality appropriate for explaining some particular cognitive process or aspect of behaviour 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. Whereas 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 behaviour, 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 behaviour in the everyday world. (2007, p. 21)

Note how formal rationality is defined in evaluative terms (“good reasoning”) and contrasted with successful behaviour. 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 behaviour 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. (2007, p. 26; italics in original)

Finally, Oaksford and Chater make a point of arguing that any adaptively rational behaviour should be justified in terms of some normative system (Oaksford & Chater 1998a, pp. 291–97; 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 enthymematic; 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 (sect. 1), a normative system is one that is taken as an evaluative ought for human rationality. A *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 programme, 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 programme of Stanovich and West

Another highly influential research programme 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 2000b). Their theory is of a type termed “default-interventionist” by Evans (2008), as indeed is Jonathan Evans’ own dual-process theory (Evans 2006; 2007). Hence, 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 we can 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 preconditions 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 better 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 provided 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 behaviour 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 that 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 that

they 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 on 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.

Lastly, 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 objective 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 versus 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. Although 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 behaviour 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 people 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., 2000b) have argued for logic as the correct normative system because those who are of *highest ability* solve the problem in these terms. But these participants 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 research programmes of the is-ought type is bound to lead to these differences. Adaptations per se can provide us only 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. Whereas Oaksford and Chater do not seem to acknowledge that there might be a clash, Stanovich and West do, and their solution is determined not by the empirical data but by evaluative considerations; that is, the idea that rationality is determined at the *individual* level, giving preference to not only 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 programmes. 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 on both the processing and

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 inquiry

the computational levels. A descriptivist approach may free the psychology of reasoning and JDM from these research biases.

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 that 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 of variance other than the logical structure. This describes the predominant practice in the psychology of reasoning over the past 50 years or so, and it explains the origins of such terms as *belief bias* (e.g., Evans et al. 1983; and see sect. 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 the answers wrong, in contrast with the remarkable achievements of the human species in many specific fields of endeavour 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 1998a; 2007). The new paradigm that is emerging utilizes a wide variation of methods, with 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

target article, there is an as yet unresolved debate about whether the new paradigm requires an alternative normative theory, such as Bayesianism (Evans, in press b). The prior rules bias is still active – only the proposed rules have changed.

If we examine the study of judgement 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 judgements 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 judgements (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), as 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 et al. 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 2010a)? 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 et al. 1983), but other kinds cannot and require rule-based training. 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,” labour-saving rules of thumb that work well in some circumstances (Gigerenzer et al. 1999), but it certainly cannot provide us with the ability to be lawyers, engineers, and rocket scientists *without training*.

Normativism has affected not just 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 Jonathan Evans' 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 it 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 et al. 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 AC (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. Logician 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 et al. 1993). How could this be if different mechanisms are involved?

The field of judgement and decision making is, if anything, even more prone to interpretation bias than is 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 2007; Evans & Over 1996; Stanovich 1999; 2010a). 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 judgement 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, Hahn and Warren 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 correspondingly 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 the 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 paradigm.

Dual-process theories postulate two types of processes: heuristic, rapid, parallel preconscious processes (“Type 1”) versus analytic, effortful, sequential processes that correlate with general ability (“Type 2”). 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, because the two processes 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) early two-factor theory of reasoning, for example, proposed that logical and non-logical processes combined in determining behaviour. A classic example is the *belief bias paradigm* in syllogistic reasoning, in which participants have to judge the validity of arguments that are either logically valid or invalid and have

either believable or unbelievable conclusions. Evans et al. (1983) established that people will prefer both logically valid conclusions (with belief constant) and believable conclusions (with logic constant), which the authors characterized at the time as a within-participant conflict between *logic* and *belief*. Stanovich’s (1999) earlier research programme on individual differences in cognitive performance associated normatively correct responding with high cognitive ability and associated 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’s (1996) 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. (p. 8)

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 sect. 6), and hence of empirical normativism, as it presupposes that System 2 corresponds to a normative system. Because 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, e.g., De Neys & Glumicic 2008; see also Ferreira et al. [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 et al. 2002) has been well argued by advocates of fast and frugal heuristics (Gigerenzer 2007; Gigerenzer et al. 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 et al. 2006; Klein 1998; Reyna 2004; Wilson & Schooler 1991). 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 behaviour (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 et al. 2000; Verschuere et al. 2005). These theories propose *both* (1) that beliefs may influence responding directly through heuristic cues and (2) 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 who tend to make the two types of inference. In particular, Stanovich’s recent research programme 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. Whereas *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 the psychology of reasoning and JDM, how it is studied, and how findings are reported and interpreted, both on the processing and the computational levels 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* has 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 than empirical arguments when they change into alternative-norm paradigms, and has arguably sentenced multiple-norm paradigms to unwarranted neglect. These biases are highly prevalent and afflict much of the field. Although 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 have shown how normativist thinking has biased and constrained the relevant research programmes. Illustrating the problem with the case of research on dual processes, we have also identified a specific bias which we term the *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 JDM (judgement and decision making). 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 programme. 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

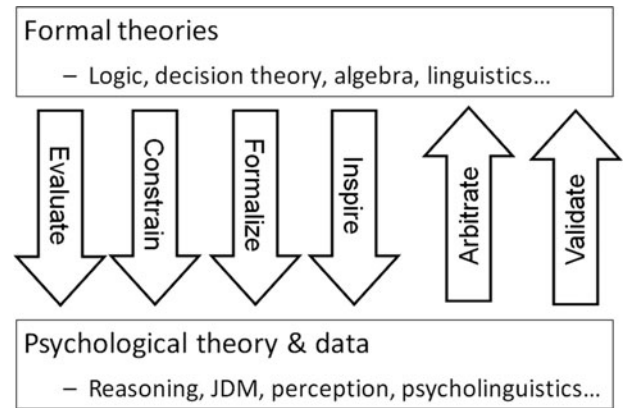


Figure 2. Models of interaction between formal and psychological theories.

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.

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 et al. 1999). Gigerenzer and Todd, in introducing their research programme on “fast and frugal” heuristics (Gigerenzer & Todd 1999), 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 behaviour. In noting the probabilistic revolution that has undermined logicism, Gigerenzer and Todd 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” (Gigerenzer & Todd 1999, p. 6). Paradoxically perhaps, it is this very rejection of normativism that has led Stanovich (1999, pp. 57–58) to ascribe the naturalistic fallacy (see sect. 5) to Gigerenzer (1991), maintaining that he links 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, pp. 255–56), 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 use 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 programme. While both programmes are Bayesian (and decision-theoretic) in their treatment of deductive competence, HTT draws on Bayesian theory for its psychological features: subjectivity, belief, uncertainty. Indeed, Evans (2007) comments that Bayesian philosophy is a “more credible *descriptive* model” of scientific thinking than is the Popperian approach (p. 33, emphasis ours).

A similar case is the theory of the suppositional conditional, a special case of HTT (Evans & Over 2004; see also Evans et al. 2003; 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 1931/1990). 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 a 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, Ch. 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 he creates a clear distinction between normative rationality and analytic thinking (pp. 159–61). Evans also addresses the issue of normative rationality explicitly, with these 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. (Evans 2007, p. 161)

HTT, then, uses 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 does 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 endeavour. 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 programme of Stanovich and West (2000b; also Stanovich 1999; 2004; 2009b) and the version put forward by Baron (e.g., 2008). Such authors wish to find ways to 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 behaviour 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 behaviours 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 behaviour, 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 them 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 judgemental: 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 programme 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 to 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” (Baron 2008, p. 61); “rational decisions can be defined relative to a person at a given time, with a given set of beliefs and goals” (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 the 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 behaviour than if it were the other way around. Instrumentalism is the default behaviour and easily explained. The converse behaviour 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 behaviour, 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.

ACKNOWLEDGMENTS

This work was partially supported by research travel funding from De Montfort University to the first author. A preliminary sketch of some of the arguments in this article 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; to

John Clibbens, Vinod Goel, Ulrike Hahn, Philip Kargopoulos, Gernot Kleiter, Niki Pfeifer, Keith Stanovich, Rich West, and members of the University of Plymouth Thinking and Reasoning Group, for useful discussions.

APPENDIX

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 article the term *logicism* is used in its meaning in psychology of reasoning; that is, 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, namely, the view that mathematics can be reduced to logic (e.g., Whitehead & Russell 1910/1962). 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 judgement 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*; indeed, 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. 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.

Open Peer Commentary

Throwing the normative baby out with the prescriptivist bathwater

doi:10.1017/S0140525X1100046X

Theodora Achourioti^a, Andrew Fugard^b and Keith Stenning^c

^a*Institute for Logic, Language, and Computation/Philosophy Department, University of Amsterdam, 1012 GC Amsterdam, The Netherlands;* ^b*Center for Advanced Studies and Research in Information and Communication Technologies & Society, University of Salzburg, 5020 Salzburg, Austria;*

^c*School of Informatics, The Forum, University of Edinburgh, Edinburgh EH1 1AB, Scotland, United Kingdom.*

t.achourioti@uva.nl andy.fugard@sbg.ac.at

k.stenning@ed.ac.uk

Abstract: It is neither desirable nor possible to eliminate normative concerns from the psychology of reasoning. Norms define the most fundamental psychological questions: What are people trying to do, and how? Even if no one system of reasoning can be the norm, pure descriptivism is as undesirable and unobtainable in the psychology of reasoning as elsewhere in science.

Elqayam & Evans (E&E) construe normativism as the proposal that a unique formal model (logic, Bayesian probability, whatever) is *the* paradigm of human rationality. The authors do recognize that there are programs “lower on prescriptive normativism,” (target article, sect. 1, para. 9), that is, programs that make use of reasoning norms but not as the unique standard of human rationality, either empirically or prescriptively construed; see, for example, Stenning and van Lambalgen (2008). The latter example proposes multiple logics as formal models of reasoning, and that their use necessarily involves normative claims. E&E, in contrast, propose a thoroughgoing descriptivism for the study of human reasoning, but believe that they can preserve the use of formal systems (competence models) in their descriptivism.

We agree that resorting to a single formal system as the standard of rationality is doomed to failure (Stenning & van Lambalgen 2008), but we disagree that formal systems can be preserved in a thoroughgoing descriptivism. Here we explain the role of normativity in formal models by making use of the distinction between constitutive and regulative norms (Kant 1781/1998; Rawls 1955; Searle 1970). And we explain the psychologically crucial role of normativity through the interplay between these two kinds of normativity in the application of formal systems to the modelling of data. It is a corollary that any single formal system will be inadequate for modelling all human reasoning.

E&E equate normativity with the “ought” that is erroneously derived from “is”; what they fail to notice, however, is that without an “ought” there can’t be an “is” in the first place. To help identify the norms that play an important role in human reasoning, one can make use of the distinction between constitutive and regulative norms. Abstractly, norms are *constitutive* of a certain behaviour if they are part of recognising the behaviour as the one that it is identified to be; norms are *regulative* of a certain behaviour if they are responsible for steering a behaviour in a certain direction. The two examples of norms that the authors mention at the very beginning of their article are a good illustration of the distinction: Not conforming to the rules of chess means that one is not playing chess, whereas one is still driving when not heeding traffic laws.

Constitutive norms are *internal* to a certain reasoning system in that they address the question of “*what* the reasoning is,” whereas regulative norms are *external* to a system in that they address the question of “*why* the reasoning is the one that it is.” The distinction may seem analogous to E&E’s empirical–prescriptive distinction, but it is not. Constitutive norms are not simply empirical because

they cannot be described by mere recourse to experience: the rules for playing a certain game are not something one can just observe. Regulative norms are not prescriptive in the sense that they do not prescribe a unique behaviour; rather, they provide those reasons and constraints that make a certain choice possible in the first place (the choice need not be a conscious one).

Syllogistic reasoning provides a laboratory example. We present the 64 argument forms of the syllogism to subjects and want to interpret the results. What should they do? Experimenters assumed for many years that it was obvious that they should obey classical logic – after all this was the original logical fragment, dammit. But Stenning and Yule (1997) have pointed out that there is another important interpretation as *cooperative exposition*, in which the readers’ task is to discover the author’s single intended model. The defeasible logical model of this task is in many ways opposite. These two kinds of logic have different constitutive norms, which mean that they have different regulative norms. Classical logic is partly constituted by its concept of validity – truth of conclusions in all possible models. This means it is regulatively appropriate for use in adversarial argument from closed premises about all interpretations. Defeasible logic is also partly constituted by its own concept of validity – roughly, the truth of conclusion in the intended model – which means it is regulatively appropriate for cooperative reasoning from a database of long-term knowledge plus the present text, to a single interpretation of the text. The first goal of a psychology of reasoning should be the empirical investigation of which of these goals (or others) subjects adopt, and in what contexts. Without knowing how subjects understand the task, we cannot sensibly interpret the data. There are also issues about which goal it is more reasonable to adopt in the experimental context – an important regulative normative question. Subjects themselves often switch interpretations/goals spontaneously during a sequence of problems, sometime uttering expletives when they realise what the experimenter intended – in some ways the most interesting psychological events (see also, Fugard et al. [2011b] for an example in a probabilistic framework).

None of this is available to a thoroughgoing descriptivism, which can tabulate the frequencies of the conclusions drawn for the 64 problems but has no basis for generalising beyond this table, or even justifying the range of options presented. Such thoroughgoing descriptivisms were the goal of positivist theories of science widely discredited by the observation that data is always shot through with theoretical interpretation, knowingly or not (Hanson 1958; Kuhn 1962). It is a corollary that no single formal system can provide a regulative norm for human reasoning, because, at the very least, argument and exposition are two incompatible reasoning goals that people ought to adopt at different times. Bayesian accounts are not a homogenous alternative, because they fail to model the necessary processes of interpretation (Fugard et al. 2011a; Stenning & van Lambalgen 2010). Of course, many other systems of reasoning with other goals are also needed.

It is fashionable in some quarters to claim that subjects frequently have no systematic goal for their reasoning; a suicide note for the psychology of reasoning – subjects with no goals. Psychologists should learn to love normativity, suitably regulated – after all biology is shot through with normativity, and seems to be doing rather well.

Norms for reasoning about decisions

doi:10.1017/S0140525X11000100

Jean-François Bonnefon

CNRS and Université de Toulouse, CLLE, Maison de la Recherche, 31058 Toulouse Cedex 9, France.

bonnefon@univ-tlse2.fr

Abstract: Reasoning research has traditionally focused on the derivation of beliefs from beliefs, but it is increasingly turning to reasoning about

decisions. In the absence of a single, entrenched normative model, the drive toward normativism is weaker in this new field than in its parent fields. The current balance between normativism and descriptivism is illustrated by three approaches to reasoning about decisions.

As suggested by Elqayam & Evans (E&E), normativism faces an arbitration problem when many researchers in the field are familiar with many formalisms that are eligible for normative status. Normativism, obviously, faces even more of a problem when very few researchers in the field have heard of any eligible formalism. Accordingly, normativism is expected to be strong in fields where exactly one formalism is a well-known candidate for normative status. The psychology of reasoning qualifies for this category, and so does the psychology of decision-making. Although there are investigations of formalisms that provide alternative norms for reasoning (Benferhat et al. 2005) or decision-making (Bonnefon et al. 2008), they have garnered considerably less attention than the dominant competitors: logic and Bayesianism (for reasoning), and utility maximization (for decision-making).

An interesting trend is developing, however, with the ongoing convergence of the two fields. Students of reasoning, in particular, are increasingly incorporating decision-theoretic elements in their investigations. Whereas reasoning research has traditionally focused on the derivation of disinterested beliefs from other disinterested beliefs, current research is giving increasing attention to reasoning about preferences and decisions. Under its most general form, the problem is to account for the ways in which people infer the beliefs, preferences, or decisions of other people based on partial knowledge about this whole set of information. Because there is no single, well-entrenched formal model of how this should be done, the drive toward normativism is paradoxically weaker when investigating how people reason about decisions, than when investigating either how they reason in general, or how they make decisions. Even in Artificial Intelligence, models of how agents reason about decisions make a clear distinction between their own decision-making models, the models others may be using, and the extent to which everyone may deviate from these models (Gal & Pfeffer 2008).

Accordingly, when moving from reasoning about facts to reasoning about decisions, E&E's distinction between prescriptive and descriptive normativism unpacks into three distinct questions: (1) Do models of reasoning about decisions consider that people assume other people to be normative decision-makers? (2) Is there an optimal way to reason about decisions, whether or not other people are assumed to be normative decision-makers? (3) If so, does reasoning about decisions reflect this optimal solution? The answers to these three questions can vary to a large extent in different psychological approaches to reasoning about decisions. The computational model offered by Baker et al. (2009), and the descriptive model put forward by Bonnefon (2009) may represent two extremes on that spectrum, while the graphical causal model approach adopted by Hagmayer and Sloman (2009) seems to take an intermediate perspective.

On the normativist side of the spectrum, Baker et al. (2009) offer a model of reasoning about decisions where the response to each of the three questions above is “Yes”. In this model, other people are assumed to be rational decision-makers who choose actions that achieve their desires most efficiently, given their beliefs. By virtue of being computational, the model assumes that there is an optimal solution to the problem of reasoning about decisions (prescriptive normativism), and the normative framework that it uses is Bayesian inverse planning from Markov decision problems. Finally, it is expected that real-life inferences will be better predicted by this optimal model than by simple heuristics (descriptive normativism).

Hagmayer and Sloman (2009, Experiment 4) investigated inferences about the beliefs of other people, based on their preferences and decisions. Their core assumption is that people assume others to conceive of choices as interventions on a causal graphical network. Hagmayer and Sloman do not explicitly

state that the choice-as-intervention (CI) principle is normative, but the reasons they consider for someone not to endorse it have an obvious non-normative flavor (e.g., self-deception, self-handicapping, illusion of control, superstition). Thus, the model appears to feature a normative calculus (causal graphical models) involving the assumption that other people are normative decision-makers (the CI principle). As such, it would qualify for prescriptive normativism, as defined by E&E. Hagmayer and Sloman do not endorse descriptive normativism, however, as they explicitly consider the conditions (e.g., lack of motivation or capacity) under which inferences are not expected to reflect the CI principle and the causal graphical calculus.

Finally, at the descriptive end of the spectrum, Bonnefon (2009) offers a theory of how reasoners infer the decisions and attitudes of agents, based on utility conditional statements such as “if p then q ”, where p and q bear on the utility functions of these and other agents. Rather than assuming people to believe that other agents are fully rational, the theory is based on folk axioms of decision, that is, the naïve beliefs that people have about the way other agents make their decisions. The theory is purely descriptive in the sense that it does not assume the existence of an optimal solution for reasoning about utility conditionals. Its response to each of the three questions above is therefore “No”.

As can be seen from these three examples, the new field of reasoning about decisions is remarkably opened to descriptivist as well as normativist approaches, in contrast to its two parent fields (reasoning and decision-making), which are much more influenced by normativist approaches. In light of the case that E&E make against the dominance of normativist approaches, this flexibility bodes well for the future of the psychology of reasoning about decisions.

The unbearable lightness of “Thinking”: Moving beyond simple concepts of thinking, rationality, and hypothesis testing

doi:10.1017/S0140525X11000598

Gary L. Brase and James Shanteau

Department of Psychology, Kansas State University, Manhattan, KS 66506.

gbrase@ksu.edu shanteau@ksu.edu

http://www.k-state.edu/psych/research/brase_gary.htm

http://www.k-state.edu/psych/research/shanteau_james.htm

Abstract: Three correctives can get researchers out of the trap of constructing unitary theories of “thinking”: (1) Strong inference methods largely avoid problems associated with universal prescriptive normativism; (2) theories must recognize that significant modularity of cognitive processes is antithetical to general accounts of thinking; and (3) consideration of the domain-specificity of rationality render many of the present article's issues moot.

We are happy to agree with Elqayam & Evans' (E&E's) position that prescriptive norms can be damaging to both scientific progress and society to the extent that they repress innovation, creativity, and knowledge. As E&E point out, there is a useful and important distinction to keep in mind between prescriptive normativism (how one ought to behave in order to conform to some abstract ideal) and empirical normativism (how one ought to behave, given a particular model or theory of how the mind works). Somewhere around here, however, we part ways in our views of how the behavioral sciences should deal with this issue. We believe there is little to gain from following E&E down the path of yet another general, and generally vague, approach such as Hypothetical Thinking Theory. Solving the problems they raise will require a better conception of how theories are constructed and evaluated.

Strong inference methods. The method of strong inference (Platt 1964) rejects the traditional null-hypothesis model and instead tests between multiple, viable scientific hypotheses about a phenomenon. One can then devise experiments that adjudicate between these rival hypotheses (by identifying the situations under which the hypotheses made different predictions and conducting research on these situations). This process can be recycled as many times as needed to definitively exclude one or more of the alternate hypotheses and to evaluate subsequent refinements of remaining hypotheses. Using strong inference methods generates multiple empirical norms and would go a very long way towards avoiding the problems associated with prescriptive normativism. Strong inference methods also appear to occupy the top, left corner of E&E’s Figure 1 (see target article), a position which they claim is impossible.

The idea of strong inference is hardly new (Anderson & Shanteau 1977), so why is it not a go-to option for avoiding the pitfalls of prescriptive normativism? One likely reason is the persistence of null-hypothesis statistical testing (NHST) ideology (Krueger 2001; Loftus 1996; Nickerson 2000; S. Sun et al. 2010). This use of NHST is not only entrenched in textbooks and pedagogy, but it also appears to be a product of what comes more easily to the human mind. As F. Scott Fitzgerald (1936) noted, “the test of a first rate intelligence is the ability to hold two opposed ideas in the mind at the same time, and still retain the ability to function” (p.44). We need our science to be inspired by first-rate intelligences, not by second-rate null-hypothesis testing practices.

A theory of “thinking”. The topic of this target article is described as the study of human thinking. “Thinking” is used in the title, in the abstract, throughout the text, and in the name of E&E’s Hypothetical Thinking Theory model. It should not be news to anyone that human cognitive processes, including language, memory, reasoning, decision making, vision, hearing, attention, and problem solving, are not all accomplished by a single thinking system, and therefore are neither expected nor required to follow a single normative standard. It seems odd, therefore, to attempt the construction of a normative system that encompasses such a broad (and vague) realm of mental activities – it is almost sure to be impossible. Just as “living” entails the coordinated activities of many functionally specialized systems (e.g., for breathing, eating, digestion, excretion, etc.), the phenomena of “thinking” entail functionally specialized and coordinated systems such as those listed previously. In fact, thinking as a unitary process is contradicted by the neurosciences (e.g., by localization of function), computer sciences (e.g., in dealing with the frame problem), philosophy (e.g., see issues of indeterminacy), and biology (e.g., based on multiple adaptations designed to address multiple, discrete selection pressures).

Domain specificity. It would be much more productive, we think, to have serious discussions about the domain-specificity of rationality (and associated empirical norms). Issues such as domain-specificity, modularity of the human mind, and the principles by which the domains should be discerned are weighty and contentious (Barrett & Kurzban 2006; Carruthers 2006; Hagen 2005; Samuels 2000). For the purposes of this commentary, however, it is sufficient to note that any degree of functional modularity (e.g., vision versus language versus reasoning) is troublesome for the monolithic E&E account of “thinking.” Even basic views of the functional modularity of the mind (e.g., Fodor 1983) create problems for any completely general model, such as Hypothetical Thinking Theory.

More ambitious views of “massive modularity” (e.g., Pinker 1997; 2002; Tooby & Cosmides 1992) are often based on extensive considerations of evolutionary principles, which are ill-represented by E&E. For example, the work of Oaksford and Chater, which is interesting in many respects, is presented as evolutionarily informed and even adaptationist because they argue “behavior is rational to the degree that it tends to increase inclusive fitness” (Oaksford & Chater 2007, p. 26, cited in E&E’s

target article, sect. 5.1, para. 3). Such use of a vague “inclusive fitness” rationale does not qualify Oaksford and Chater as “adaptationist leaning[.]” (see sect. 5.3, para. 1), just as acknowledging that there are neurological underpinnings to the mind does not make them or anyone else neuroscientists.

In summary, rather than a generic model of thinking (dual process or otherwise), we need models of human cognition that incorporate the numerous functionally distinct domains of human cognition. There are many models of how the modularity of mind is structured, and this situation should be taken advantage of by utilizing strong inference methods.

Competence, reflective equilibrium, and dual-system theories

doi:10.1017/S0140525X11000550

Wesley Buckwalter^a and Stephen Stich^b

^aDepartment of Philosophy, Graduate Center, City University of New York, New York, NY 10016-4309; ^bDepartment of Philosophy, Rutgers University, New Brunswick, NJ 08901-1107.

jbuckwalter@gc.cuny.edu [sstich@rucss.rutgers.edu](mailto:ssstich@rucss.rutgers.edu)

<https://wfs.gc.cuny.edu/JBuckwalter/index.html>

<http://www.rci.rutgers.edu/~stich/>

Abstract: A critique of inferences from “is” to “ought” plays a central role in Elqayam & Evans’ (E&E’s) defense of descriptivism. However, the reflective equilibrium strategy described by Goodman and embraced by Rawls, Cohen, and many others poses an important challenge to that critique. Dual-system theories may help respond to that challenge.

Elqayam & Evans (E&E) propose that the study of human mental processing is best served by keeping normative and descriptive accounts of reasoning competence separate. We enthusiastically endorse E&E’s purely descriptivist approach, and we share their skepticism about normativism. However, we think that E&E have seriously underestimated how hard it is to undermine the view that a competence theory is *also* a normative theory. We begin by explaining why the link between competence and normativity is so deeply embedded in philosophical thinking. We then argue that dual-system theories of reasoning may provide a powerful new way of challenging this link.

Traditionally, the philosophical literature on rationality has focused on the question of when an inference is justified. In one of the most influential passages of twentieth-century philosophy, Nelson Goodman offered the following answer:

The validity of a deduction depends upon . . . conformity to valid rules . . .

But how is the validity of rules to be determined? . . . Principles of deductive inference are justified by their conformity with accepted deductive practice. Their validity depends upon accordance with the particular deductive inferences that we actually make and sanction. If a rule yields unacceptable inferences, we drop it as invalid. Justification of general rules thus derives from judgments rejecting or accepting particular deductive inferences.

This looks flagrantly circular . . . But this circle is a virtuous one. The point is that rules and particular inferences alike are justified by being brought into agreement with each other. *A rule is amended if it yields an inference we are unwilling to accept; an inference is rejected if it violates a rule we are unwilling to amend.* The process of justification is the delicate one of making mutual adjustments between rules and accepted inferences; and in the agreement achieved lies the only justification needed for either. (Goodman 1965, pp. 66–67, emphasis in the original)

In *A Theory of Justice*, John Rawls endorses this “process of mutual adjustment of principles and considered judgments” as the primary method for justifying moral judgments and moral principles (Rawls 1971, p. 20). When this process is successful, we are in what Rawls labels “reflective equilibrium.” Though

both the label and the description of this method are twentieth century products, a case can be made that the method itself was used by Aristotle, Plato, and many other important figures in the history of philosophy (see Rawls 1971, p. 51; cf. Stich 2001).

Rawls famously goes on to note that there is an analogy between the method of reflective equilibrium used in moral philosophy and the Chomskian project of “describing the sense of grammaticalness that we have for the sentences of our native language” (1971, p. 47). For Chomsky and his followers, the project that Rawls has in mind goes by another name: describing a speaker’s grammatical *competence*. A decade later, L. J. Cohen extended Rawls’ analogy to the study of *reasoning competence* (Cohen 1981). To determine our reasoning competence, Cohen maintained, we must assemble lots of information about the inferences that we actually accept and then construct the simplest set of rules that captures most of those inferences. To construct a normative theory of reasoning of the sort that Goodman proposed, we would do exactly the same thing. Thus, Cohen argued, a descriptive account of our reasoning competence and a normative theory of reasoning that is supported by the reflective equilibrium method *must coincide*. Though people may make lots of performance errors in reasoning, their underlying reasoning competence must be normatively impeccable.

According to E&E, “inferring an ‘ought’ type conclusion from ‘is’ type premises is highly controversial, and considered by many authors to be a logical fallacy” (target article, sect. 5, para. 2). They are surely right that *some* inferences of that sort are very controversial in philosophy. However, the inference from the “is” type premise:

(i) This set of (moral or inferential) rules is in reflective equilibrium.

to the “ought” type conclusion:

(ii) This set of rules is justified.

is far less controversial. Indeed, many philosophers would insist that in both ethics and logic this is by far the *best* way of arguing that a set of rules is justified. But for E&E’s critique of normativism to be persuasive, the inference from (i) to (ii) must be challenged.

One of the most effective challenges, we believe, is to emphasize the way in which such inferences can lead to radical relativism. In Stich (1981; 1990) it was noted that it is logically possible that just about any set of inference rules could be in reflective equilibrium for a cognitive agent. So if the inference from (i) to (ii) is accepted, just about any set of inference rules could be justified. This strategy made little headway, however, since there was no evidence that it is *psychologically* possible for people to have significantly different reasoning competences. We are inclined to think that Stanovich’s early work on individual differences goes a long way toward addressing that concern, though we do not endorse his claim that the inferential competence of people with high IQ should be considered normatively privileged (Stanovich 1999).

Recent work on dual-system theories of reasoning poses another powerful challenge to the idea that reasoning competence must be normatively impeccable. If these dual-system theories are on the right track, then reasoning is subserved by two distinct systems, each with its own competence. Moreover, if – as E&E suggest – much of the content of System 2 is *learned* via explicit instruction or acquired from the surrounding culture, then it is to be expected that the inferential rules in System 2 will be different for different individuals and different cultures. If each person’s reasoning is subserved by two quite different systems, and the second system varies significantly among individuals and cultures, the claim that an individual’s reasoning competence must be normatively impeccable is very implausible indeed. As awareness of dual-system theories of reasoning becomes more widespread, it may finally begin to undermine the deeply entrenched philosophical view that a theory of reasoning competence has normative implications.

A role for normativism

doi:10.1017/S0140525X11000471

Igor Douven

Faculty of Philosophy, University of Groningen, 9712 GL Groningen, The Netherlands.

i.e.j.douven@rug.nl

www.rug.nl/staff/i.e.j.douven/index

Abstract: Elqayam & Evans (E&E) argue against prescriptive normativism and in favor of descriptivism. I challenge the assumption, implicit in their article, that there is a choice to be made between the two approaches. While descriptivism may be the right approach for some questions, others call for a normativist approach. To illustrate the point, I briefly discuss two questions of the latter sort.

In their target article, Elqayam & Evans (E&E) argue against prescriptive normativism and in favor of descriptivism. This commentary questions the assumption, implicit in E&E’s article, that in the study of human reasoning and decision making there is a fundamental choice to be made between a normativist and a descriptivist approach. I argue, instead, that there is a role for both approaches. In particular, I argue that some research questions fit most naturally in a normativist framework, while others are better dealt with in a descriptivist framework. E&E’s discussion demonstrates the validity of the latter framework for specific types of studies, but their discussion seems to overlook a number of worthwhile research questions that clearly call for a normativist approach. I would like to draw attention to two such questions.

The first question is suggested by a variety of related arguments that philosophers and mathematicians have proposed in defense of the postulates of Bayesian epistemology, that is, the probability axioms as static norms of rationality, and Bayes’ rule as a dynamic norm of rationality (e.g., Joyce 1998; Rosenkrantz 1992). In sharp contrast with the more familiar Dutch book arguments, which purport to demonstrate that obeying the Bayesian norms has certain *practical* advantages, the newer arguments are explicitly meant to show that obedience to those norms is *epistemically* mandated. For instance, the newer arguments for Bayes’ rule aim to show that by following this rule, in the long run one will be more accurate in one’s assignment of probabilities than one could have been by following any other rule. (Accuracy here is measured by means of a so-called scoring rule.) This claim is then coupled with a (purportedly) a priori argument to the effect that long-run accuracy is what, qua rational epistemic agents, we aim at – which yields the desired epistemic defense of Bayes’ rule. Even if we leave open the question of whether long-run accuracy has been rightly identified as our epistemic goal, the aforementioned arguments, if sound, entail that *if* an agent commits himself or herself to the goal of long-run accuracy, *then* he or she should update her subjective probabilities by dint of Bayes’ rule. Does this hold true of real people? If not, what might explain the discrepancy between the norm and the observed behavior? These are research questions involving a normative “ought,” and nothing E&E say suggests that these questions are not worth investigating. Indeed, a discussion of the a priori arguments of the kind referred to here, and the implications they might have for psychological research, is altogether lacking from their article.

E&E rightly point out that when there are rival normative accounts of some aspect of human reasoning or decision making, experiments will not help us in deciding which of those (if any) is correct. That might suggest that empirical evidence has no normative value indeed. This suggestion appears dubitable, however, when we consider the situation in which we have simply *no* normative account of some aspect of our reasoning, while at the same time it does seem obvious that that aspect is subject to norms of rationality. In such a situation, empirical evidence might help us find those norms. That would go at least some way toward deriving an *ought* from an *is*.

To illustrate the point, consider the issue of updating on conditionals. In a standard Bayesian picture, we ought to update on incoming information by conditionalizing on it (i.e., by using Bayes’ rule). Back in 1980, Brian Skyrms said that “we have no clear conception of what it might be to conditionalize on a conditional” (Skyrms 1980, p. 169). Since then, only very little progress has been made on this matter. That is unfortunate, given that it is uncontroversial that updating on conditionals is governed by *some* norms. Douven and Dietz (2011) and Douven and Romeijn (in press) describe hypothetical cases in which people obtain conditional information. (They learn, for instance, that if it continues to rain, then a specific football match will be cancelled.) These works then point out that we have clear intuitions about these cases as to whether, given the context in which the conditional information is obtained, an update on that information should lead the person receiving the information to raise or, rather, to lower the probability he or she assigns to the conditional’s antecedent, or whether that probability should remain unaltered. However, we are currently clueless as to what the underlying norms might be. Although the past decade has witnessed a surge of interest among psychologists for studying conditionals in a probabilistic setting, so far experimental work on updating on conditionals is still missing. The lacuna is noteworthy in itself, but for present purposes the important point is that getting some understanding of how people *actually* update on conditionals (or perhaps how cognitively superior people update on conditionals) might help to suggest, at least in outline form, the norms that philosophers assume to exist but to this day have been unable to pinpoint. Here, too, I cannot find anything in E&E’s article that indicates that this approach is misguided.

E&E might seem to have provided a very general reason against normativism, to wit, that it has given rise to a number of research biases. I am not sure we must agree that all that they identify as research bias should be acknowledged as such, but some definitely are. Even so, pleading for normativism is not a call for thinking lightly of such problems. E&E certainly have not shown that if one embraces normativism, then one is *bound* to commit those errors. To the contrary, it is reasonable to think that once a bias has been identified as such, it should be possible to steer clear of it. Thus, these research biases give good reason for double- or even triple-checking the design of experiments conducted in a normativist framework, but they do not give good reason for abandoning the normativist approach wholesale.

The historical and philosophical origins of normativism

doi:10.1017/S0140525X11000124

Catarina Dutilh Novaes

Faculty of Philosophy, University of Groningen, 9712 GL Groningen, The Netherlands.

C.DutilhNovaes@uva.nl

Abstract: Elqayam & Evans’ (E&E’s) critique of normativism is related to an inherently philosophical question: Is thinking a normative affair? Should thinking be held accountable towards certain norms? I present the historical and philosophical origins of the view that thinking belongs to the realm of normativity and has a tight connection with logic, stressing the pivotal role of Kant in these developments.

Elqayam & Evans’ (E&E’s) thought-provoking article questions the fruitfulness of a normativistic approach to thinking in experimental psychology, but their considerations are directly related to an inherently philosophical question: Is thinking a normative affair at all? Are the reasoning processes of an agent to be held accountable towards certain norms? It is not obvious what the answer should be. Thinking seems to differ from obviously

non-normative phenomena, such as the behavior of physical objects, but there are also important dissimilarities with obviously normative phenomena, such as human public actions; in particular, while public actions typically have practical consequences, thinking as such is a private affair which only has practical consequences once translated into actions. Now, whether a given class of phenomena, such as thinking processes, does or does not fall within the realm of normativity is a question belonging to meta-ethics (Korsgaard 1998; Wedgwood 2007), as it involves a discussion of the very nature of normativity.

In fact, a negative answer to this philosophical question would be *sufficient* but not *necessary* to establish the methodological thesis defended by the authors: even if thinking turns out to be a normative phenomenon, their plea for a descriptivist approach in psychology may still stand. Just as sociology is descriptive and law is prescriptive concerning human actions, a normative approach to thinking may be more appropriately undertaken elsewhere, not within psychology, as they argue.

In this short commentary, I cannot offer a thorough examination of the philosophical question. Instead, I shall briefly present the historical origins of the views that thinking belongs to the realm of normativity and that logic constitutes the appropriate normative system (logicism). These views are firmly engrained, but the present exercise of conceptual archeology – outlining the substantive and even contentious assumptions behind them – suggests that they are far from uncontroversial. Thus, my analysis can be seen as lending support to E&E’s claims, but in fact it does not settle the issue definitively. Historical analysis shows that the particular conceptual and philosophical background underpinning normativism and logicism as targeted by the authors may be contentious, but it does not offer the final word on the *philosophical* question.

For most of its existence, logic was thought to be primarily the art or science of correct *arguing* and *disputing*, that is, as regulating multi-agent situations taking place in the public sphere. In the Latin Middle Ages, *logica* was synonymous with *dialectica*, and the 16th century author Domingos de Soto writes: “Dialectic is the art or science of disputing” (de Soto 1539–1540, f. iii rb). This is in stark contrast with the now standard definitions of logic as “the branch of philosophy that treats of the forms of thinking in general” (*Oxford English Dictionary*) or as “the science of correct reasoning” (*Webster*). When did logic cease to provide the norms for the public, multi-agent situations of debating, and come to regard the private, mono-agent situations of thinking? The main person responsible for this transformation seems to have been Kant, but the terrain had been well prepared by Descartes (tellingly, the author of the *Rules for the Direction of the Mind*). Indeed, in the preface to *Principles of Philosophy*, commenting on the ideal education for a student, Descartes writes:

After that, he should study logic. I do not mean the logic of the Schools, for this is strictly speaking nothing but a dialectic which teaches ways of expounding to others what one already knows [. . .]. I mean instead the kind of logic which teaches us to direct our reason with a view to discovering the truths of which we are ignorant. (Descartes 1985, p. 186)

But Descartes rejected the traditional conception of logic of his time (scholastic logic) as the appropriate guide for correct thinking; so in E&E’s terms, he was a normativist but not a logicist regarding thinking. It was Kant, in the 18th century, who then laid down the foundations for the still pervasive close association between thought, logic, and normativity by *internalizing* some of the key concepts of the logic of his time.

As discussed by Longuenesse (1998), Kant takes as his starting point the transcendental question, “What are the a priori conditions for the representations of objects in general?”, and reconfigures the logic of his time so as to render it useful for his transcendental project. In particular, he selectively absorbs the notions of “judgment,” “form,” and “categories” as found in the logical textbooks of the time, and puts them to use so as to describe the very conditions of possibility of our thinking and

perceiving. The concept of judgment, for example, traditionally used to refer to linguistic claims made by speakers in the public sphere, is transformed by Kant into the mental act of the understanding involved in the apperception of objects. With Kant, logic no longer primarily concerns argumentation; instead, it concerns the inner mental activities of the lonesome thinking subject. Moreover, rejecting the psychological descriptivism of Wolff, Kant insisted on the *normative* import of the rules of thought as described by logic. According to him, (general) logic deals with “absolutely necessary rules of thought without which there can be no employment whatsoever of the understanding” (Kant 1781/1787/1929, *Critique of Pure Reason*, p. A52/B76).

Crucially, thus, the ideas that thinking belongs to the realm of normative phenomena and that logic provides the canons for correct thinking are essentially *Kantian theses*, which are intimately related to his critical project and to transcendental idealism (as argued in MacFarlane 2000). But if we do not endorse transcendental idealism (and most of us do not), then we have no reason to accept uncritically this particular conception of thinking as a normative affair and the role of logic as the appropriate normative system. The historical connection between Kant and the tradition in experimental psychology in question is in fact rather straightforward: it goes in any case through Piaget, who was clearly under strong Kantian influence (Hergenhahn 2009, p. 624). We can thus conclude that the history of normativism and logicism offers no obvious reasons to endorse these positions, in particular with respect to psychological research.

Just the facts, and only the facts, about human rationality?

doi:10.1017/S0140525X11000136

Jeffrey Foss

Department of Philosophy, University of Victoria, Victoria, BC V8W 3P4, Canada.

jefffoss@uvic.ca

<http://web.uvic.ca/philosophy/people/foss/index.php>

Abstract: Elqayam & Evans’ (E&E’s) laudable program to keep the scientific investigation of human reasoning norm-free and focused on the facts alone is an essential part of a long tradition in the philosophy of science – but it faces deeper difficulties than the authors seem to realize, since reasoning is a competence, and the very concept of competence is normative.

The core virtue of empirical science, and the wellspring of its magnificent success in recent human history, is that it formulates questions to nature that can be definitively answered by observation. Scientific questions lie at the opposite end of the spectrum from questions of value, which concern not what is, but what ought to be, and which have the following logical form: Which *possibilities* should we pursue? Observing the *actual* world cannot answer such questions, and so non-terminating philosophical debates are our only recourse. The scientific spirit is philosophically repelled by such endless debate: its interest is just the facts and only the facts. Animated by this essential scientific goal, Elqayam & Evans (E&E) attempt to draw the fine line beyond which true science, and in particular, the science of human thinking (“higher mental processing – reasoning, judgement, and decision making”; target article, sect. 1, para. 2), must not venture. It *should* embrace “descriptivism” and reject “normativism” – which, as they show in considerable detail, has tainted much of the work in the field.

Put baldly, E&E’s initiative is crucially important if psychology is ever to become a true natural science (like, paradigmatically, physics). Let us pass but quickly over the quasi-paradox that E&E’s essay is *philosophical* rather than *scientific*: Rather than merely observing what psychologists in *fact* do, they make a case for what psychologists *should* do. Let us simply accept that science is a goal-directed activity, and therefore that its

philosophical foundation must define that goal, hence say what scientists *should* do (or must do in order to be scientists, etc.). Let us simply recognize that roughly four centuries ago, scientists self-consciously decided they would accept only *efficient causes* (whatever in the past or present makes things happen), and reject *final causes* (whatever future end-states, or goals, make things happen), the vexatious (because not empirically determinable) “causes” of Aristotelian science. In this way, science would concern just the facts, and only the facts.

But psychological science now wants to consider how it is that human beings pursue and achieve their goals – in other words, how a future goal affects present behavior. Can this be done without venturing beyond the facts and into values? In theory, yes. E&E therefore stand squarely in the mainstream of science when they issue their two-part philosophical prescription for psychological science (see sect. 7): (1) It may employ normativism, but only in the supporting roles of inspiring, formalizing, constraining, or evaluating purely descriptive theories. (2) Psychological theory must be distinct from psycho-therapy.

The paradigm of a science that shuns normativism and is purely descriptive is physics. It is therefore instructive to observe a crucial disparity between psychology and physics, namely that *therapy* has a role in psychology, but not in physics. All processes obey the laws of physics, whether they are the processes of, for example, a watch that keeps perfect time or one that loses an hour a day. Human thought processes also obey the laws of physics – and those of biochemistry and biophysics as well – whether they are perfectly sound or completely out of touch with reality. E&E’s proposal, then, is that the “psychology of reasoning and JDM (judgement and decision making)” should *describe* reasoning and JDM, rather than take a normative position on it: “The object is not to judge such behaviour, but to understand and predict it” (sect. 8, para. 5).

Personally, I am in complete agreement with the motivations and goals of this proposal, for I, too, would like to see the growth of a proper, value-neutral, science of human reasoning. But I think the problems faced by descriptivism go deeper than E&E might suspect, for in the end the question becomes: *What* is it this science is supposed to describe? What is referred to by E&E’s phrase “such behaviour,” that this science should understand and predict rather than judge? Humans have a huge variety of behaviors in which they come to some conclusion, and which might be called reasoning, some of which are models of rationality, and some of which are not merely irrational but perverse, bizarre, crazy – so much so, that they may not truly count as reasoning at all. The elephant in the room, completely ignored by E&E as they conduct their excellent review of the literature, is that they rely on a tacit, unstated definition of rationality and JDM. But that means that we are still in the dark when it comes to just what the subject of our science is – what it must observe and where it must find its data.

The problem faced by descriptivism may be brought into focus by means of the following argument:

1. Intelligence is a competence.
2. Competence is a normative concept.
3. Therefore: No account of intelligence can be entirely descriptive, that is, entirely non-normative.

Various processes can be identified in the brain, from fundamental cellular metabolism to information processing. The science of rationality and JDM is about some subset of information processing; namely, those processes that are rational and/or result in decisions. Leaving aside for now the thorny issue of rationality, what makes a process count as information processing? After all, *every* brain process is metabolic. Why do we take some part or aspect of brain metabolism and say that this part or aspect is information processing? The answer is that this part or aspect solves some particular problem faced by the organism: that is, it provides the organism with a *competence*, the ability to act in a smart, not a stupid, way. But competences come in many degrees and forms: whether or not something

counts as a competence always comes down to whether or not it provides some *good* for the organism. But the very concept of good is normative. Therefore, no account of intelligence can be entirely descriptive – for the simple reason that we cannot identify any observable phenomenon as intelligent unless we see it as contributing somehow to the *good* of the organism.

Overselling the case against normativism

doi:10.1017/S0140525X11000458

Tim Fuller and Richard Samuels

Department of Philosophy, The Ohio State University, Columbus, OH 43210.

fuller.193@osu.edu

<http://people.cohums.ohio-state.edu/fuller193/>

samuels.58@osu.edu

<http://people.cohums.ohio-state.edu/samuels58/>

Abstract: Though we are in broad agreement with much of Elqayam & Evans’ (E&E’s) position, we criticize two aspects of their argument. First, rejecting normativism is unlikely to yield the benefits that E&E seek. Second, their conception of rational norms is overly restrictive and, as a consequence, their arguments at most challenge a relatively restrictive version of normativism.

We are in broad agreement with much of what Elqayam & Evans (E&E) have to say in their article. First, normative theories, process theories, and competence theories are genuinely distinct kinds of theory and ought not to be confused with each other. Second, the psychology of human reasoning ought not to be merely concerned with explaining patterns of deviation from, and conformity to, putative normative standards. Third, we agree that it is extraordinarily hard to adjudicate between competing normative theories. Fourth, we agree that the central goals of the psychology of reasoning are descriptive ones: to identify the patterns of inference in which we engage and the psychological states, processes, and mechanisms that are causally responsible for these patterns. Finally, we strongly suspect that success in this descriptive project does not presuppose a prior identification of a correct normative theory of rationality. For all that, there are two points on which we take issue with E&E’s position.

The prudential case against normativism is problematic. Central to E&E’s case against normativism is that it invites various infelicitous inferences: for example, from is-to-ought, and from ought-to-is. E&E’s suggestion is that excising normative theories will remove such temptations. Yet, they also suggest that precisely the kinds of formal theories that have traditionally been treated as normative standards may also be retained in order to play a host of other roles: for example, constraining and inspiring psychological theorizing, and providing a framework for the articulation of competence theories. In short: Use the (presumed normative) formal theories in psychology, just don’t use them *as* normative theories. But we wonder: If researchers really are as prone to conflating “is” and “ought” as E&E suggest, then why should they be any more successful in avoiding a slide back towards normative interpretations of formal theories? Of course, whether or not this would occur is an empirical matter. But the fact that formal theories have a longstanding dual function – descriptive and normative – and the fact that reasoning is an inherently normative phenomenon – perhaps *the* paradigmatic object of rationality attributions – leads us to suspect that excising normative theories from descriptive psychology will not have the desired effect.

Even supposing researchers avoid a slide back into normative interpretations of formal theories, we doubt that banishing normative theories will yield the results that E&E seek. For it seems that infelicities and biases closely analogous to those that E&E trace to normativism may well still arise from the

misapplication of formal theories, *even were those theories not construed normatively*. Consider, for example, E&E’s complaint that dual process theorists sometimes make the ought-is fallacy of concluding that System 2 underlies a response just because the response is normatively correct. We agree that this is a mistake. But an exactly analogous mistake can result from misapplying formal theories in characterizing cognition, even if one denies their normative status. Suppose, for example, that one incorrectly assumes that the competence theory for System 2 is accurately characterized by some formal theory, say probability calculus. Under such circumstances one may (mistakenly) conclude that, because a response conforms to the formal theory, it was produced by System 2. But in this case the problem cannot be the normative status assigned to the formal theory. Rather, the problem is simply that one is inaccurately using a formal theory to characterize the system’s competence. Moreover, we think that the point generalizes. As far as we can tell, much the same is true for many of the other biases and infelicities that E&E attribute to normativism.

Overselling the case against normativism. E&E appear to suppose that genuine normative standards of rationality satisfy the following conditions:

(a) They are unconditional in at least the sense that their normative status depends on neither the goals of agents nor the functions of the mechanisms involved in reasoning.

(b) They are deontological in the sense that they specify *what it is* to reason correctly – what is *constitutive* of good reasoning – in terms of conformity to some appropriate set of rules or principles. Now perhaps E&E introduce these constraints merely to limit the class of rational norms to those associated with the Standard Picture (Stein 1996). Nevertheless, it is worth noting that this is a very narrow – and quite contentious – characterization of normative standards of rationality. Indeed, according to some of the most intensively discussed philosophical theories of rationality, normative standards routinely violate one or both of these constraints. One especially prominent kind of view, which is often called *consequentialism*, maintains that *what it is* to reason correctly – what is *constitutive* of being a rational reasoning process – is being an effective means of achieving some goal or range of goals. So, for example, according to one well-known form of consequentialism – *reliabilism* – a good reasoning process is one that tends to lead to true beliefs and the avoidance of false ones (Goldman 1986; Nozick 1993). But on such a view, normative standards of rationality are neither deontological nor unconditional. Moreover, as far as we can tell, E&E’s worries fail to generalize to such consequentialist theories of rationality.

With this in mind, E&E appear to be guilty of overselling their case against normativism. Contrary to what they appear to suggest, their arguments cannot plausibly be seen as militating in favor of the view that “theories of higher mental processing would be better off freed from normative considerations” (target article, Abstract). Rather, at most their arguments provide grounds for rejecting a specific kind of normativism: one that construes normative standards of rationality in a narrow and not especially plausible fashion.

Undisputed norms and normal errors in human thinking

doi:10.1017/S0140525X11000483

Vittorio Girotto

School of Arts and Design, University IUAV of Venice, 30123 Venice, Italy.

vgirotto@iuav.it

<http://www.iuav.it/Ricerca1/Dipartimen/dADI/Docenti/girotto-vi/index.htm>

Abstract: This commentary questions Elqayam & Evans’ (E&E’s) claims that thinking tasks are doomed to have multiple normative readings and

that only applied research allows normative evaluations. In fact, some tasks have just one undisputed normative reading, and not only pathological gamblers but also normal individuals sometimes need normative guidance. To conclude, normative evaluations are inevitable in the investigation of human thinking.

With some “increasingly rare” exceptions, the tasks used to investigate human thinking have multiple normative readings. Hence, one cannot establish whether they elicit correct or incorrect answers. This is the main argument used by Elqayam & Evans (E&E) against the normative evaluation of thinking performance. One problem with this argument is that some experimental paradigms for investigating human thinking have just one normative reading. E&E examine one class of these paradigms (i.e., the conditional syllogisms) and claim that it is actually subject to normative dispute. Thus, E&E conclude that sooner or later alternative normative systems will be proposed for any thinking paradigm.

There is reason to doubt this conclusion. Consider the typical estimation tasks wherein respondents judge the total or relative frequency of the items of a given class (e.g., Tversky & Kahneman 1973). Since there is an objective yardstick for enumerating the items (e.g., men in a list of people, words that begin with *r*), these tasks have one objectively correct answer. In some cases, respondents do not produce that answer. For example, Tversky and Kahneman’s respondents judged that the class of men in a list was more numerous than the class of women, when in fact the list contained 19 names of men and 20 names of women. Accordingly, one has to conclude that these respondents erred. And such a conclusion will hold until some alternative account proves that judging a class of 19 items greater than a class of 20 items is normatively correct. E&E may be right in claiming that paradigms of this kind are becoming “increasingly rare” (target article, sect. 3, para. 3). The existence of such paradigms, however, shows that investigating human thinking is not destined to use tasks with competing normative readings, and that evaluating the normative status of a given judgment is not always open to dispute.

E&E correctly insist that one should understand what individuals are doing and why, rather than “discussing what they are *not* doing” (sect. 8, para. 2, authors’ emphasis). The latter motivation may have driven some of the studies that have documented biases in human thinking. Other studies of this sort, however, have investigated the limits of human thinking processes in order to understand them. For example, consider Tversky and Kahneman’s studies mentioned above. They were not aimed to simply reveal the faulty nature of respondents’ estimations, but instead tested the hypothesis that the ease of recall of instances affects frequency and probability judgments in tasks for which there is a normative standard (e.g., judging the frequency of men in a list), as well as in tasks for which there is no such standard (e.g., judging the probability that a given depressed patient will commit suicide). Quite ironically, E&E contrast the controversial nature of thinking studies with the undisputed nature of memory studies, and indicate the “acute” problems derived from adopting memory paradigms to thinking research. Yet, as Tversky and Kahneman’s studies show, memory search is the basis of many judgmental activities, and investigating recall processes may shed light on thinking processes.

E&E concede that in some cases one is entitled to evaluate thinking performance: “If your objective is to improve thinking (rather than to understand it), then you must have criteria for distinguishing good thinking from bad” (sect. 5.2, para. 5). E&E refer to individuals who behave against their interests, such as pathological gamblers. They argue that one has to help these individuals by modifying their wrong beliefs and teaching them the rules of probability calculus. According to E&E, such an instrumental approach is necessary in applied research but totally inappropriate in “basic theoretical research” (sect. 8, para. 2). The point is that even basic research has discovered individuals who need some normative help. Consider

respondents who *bet* on the conjunction of events A&B, rather than on event A (Tversky & Kahneman 1983). If you follow E&E’s recommendation, you should refrain from judging respondents’ bets. Yet, these respondents behave against their interests. They miss a chance of winning the bet, the one in which the conjunction of events A¬-B occurs. Therefore, you should inform them that they have made a bad decision. In doing so, you employ a normative standard; that is, you inform respondents that they do not conform to basic norms of probabilistic calculus. In sum, even basic research may force you to evaluate respondents’ performance and to improve it by means of normative guidance.

Besides applied domains, there is an entire domain of basic research, neglected by E&E, wherein evaluating thinking performance is inevitable. According to E&E, when respondents have to evaluate a posterior probability, in order to “get the problem *right*” (sect. 6, para. 5, emphasis added), respondents need to reason about frequencies or to learn Bayes’ rule. This claim is inaccurate, since respondents, including preschoolers, may solve this problem without reasoning about frequencies (Giroto & Gonzalez 2001; 2008). Preschoolers, of course, are not familiar with the rules of probability calculus. Yet, at around the age of five, they solve this sort of problem. Before that age they fail to, and after that age their performance improves. This example is relevant because it concerns a problem for which even E&E accept that there is only one right solution. However, the entire investigation of the development of thinking processes speaks in favor of normative evaluation. If one does not use normative standards, how could one compare the answers produced by children of different ages? More generally, how could one claim that children’s thinking processes improve (or, for that matter, worsen; see Noveck 2001), if one does not have normative standards to assess them?

E&E are probably right in claiming that some psychologists say, “Respondents should not think this,” in the same evaluative sense in which they say, “Poverty should not exist.” Yet, there are cases, like the ones mentioned above, in which psychologists are entitled to say, “Respondents should not think this,” in the same evaluative sense in which they say, “You should not eat this.”

Normative theory in decision making and moral reasoning

doi:10.1017/S0140525X11000495

Natalie Gold,^a Andrew M. Colman,^b and Briony D. Pulford^b

^a*School of Philosophy, Psychology and Language Sciences, The University of Edinburgh, Edinburgh EH8 9AD, Scotland, United Kingdom;* ^b*School of Psychology, University of Leicester, Leicester LE1 7RH, United Kingdom.*

Natalie.Gold@ed.ac.uk amc@le.ac.uk bdp5@le.ac.uk

<http://homepages.ed.ac.uk/ngold>

<http://www2.le.ac.uk/departments/psychology/pp1/amc>

<http://www2.le.ac.uk/departments/psychology/pp1/staff/bdp5>

Abstract: Normative theories can be useful in developing descriptive theories, as when normative subjective expected utility theory is used to develop descriptive rational choice theory and behavioral game theory. “Ought” questions are also the essence of theories of moral reasoning, a domain of higher mental processing that could not survive without normative considerations.

Normative theories may be superfluous in certain specific cases discussed in Elqayam & Evans’ (E&E’s) stimulating and informative target article. But the fact that some people may be tempted by a fallacious is-ought inference is not sufficient reason for abandoning normative theories in all cases.

A widely held position in philosophy of science is that all scientific observations are theory laden (e.g., Kuhn 1962), partly

because investigators’ theories influence what data they consider it worthwhile to collect. Further, when interpreting behavior, we tend to apply Davidson’s (1973) “principle of charity,” originally intended for interpreting sentences but more widely applicable. This involves assuming that people are generally rational and interpreting their behavior in that light. Thus, behavioral researchers implicitly draw on a normative theory.

Normative theories have also been useful in generating powerful descriptive theories, using a style of theorizing that does not fall foul of the is-ought fallacy. For example, subjective expected utility (SEU) theory is evidently normative, specifying what choices rational agents ought to make in order to satisfy their own desires. By appending to SEU a hypothesis of weak rationality, according to which people try to do the best for themselves in any circumstances that arise, we derive the descriptive principle of *methodological individualism* (Weber 1922/1978, Ch. 1), a mainstay of the contemporary social sciences, reflected in rational choice theory and behavioral game theory, both direct descendants of SEU theory (Elster 1989).

It was Savage (1972), not von Neumann and Morgenstern (1947), who introduced a normative interpretation into SEU theory: “One idea now held by me that I think von Neumann and Morgenstern do not explicitly support, and that so far as I know they might not wish to have attributed to them, is the normative interpretation of the theory” (Savage 1972, p. 97). Normative considerations seem quite natural and useful in judgment and decision making research. If you invite people to make a snap choice between 96 × 69 cents and 87 × 78 cents, most will choose 96 × 69 cents; but if you point out that 96 × 69 = 6,624, whereas 87 × 78 = 6,786, they will swiftly change their minds, if allowed to (Binmore 2009, pp. 22–23). This illustrates two important facts: first, people generally try to act rationally in the sense of maximizing their expected utilities; but second, they are limited by bounded rationality and are prone to error.

A domain of higher mental processing within which normative considerations seem quite unavoidable is moral reasoning. Evaluative “ought” questions are the very essence of moral reasoning. We are currently engaged in a research project investigating judgments as to whether it is morally acceptable to sacrifice one life to save five in the following famous Trolley problem (Foot 1967):

A trolley is running out of control down a railway track. In its path are five people who will be killed if it continues on its course. By operating a lever, you can divert the trolley on to a different track, where a solitary man in its path will be killed. Is it morally permissible to operate the lever?

Most people (90%, according to Hauser 2007) say yes; but Thomson’s (1976) closely related Footbridge problem elicits very different responses:

A trolley is running out of control down a railway track. In its path are five people who will be killed if it continues on its course. You are on a footbridge over the tracks next to a large man. The only way to save the five people is to push the man off the bridge, into the path of the trolley, where only he will be killed. Is it morally permissible to push the man off the footbridge?

Most people (90%, according to Hauser 2007) say no. Why do most people consider it morally acceptable to sacrifice one life to save five in one problem but not the other?

From the perspective of cognitive psychology, the differences in responses to the two problems are reminiscent of the classic demonstration of a *framing effect*, in which two different descriptions of a problem involving a certain number of lives at risk elicit difference responses (Tversky & Kahneman 1981). Greene (2007) has argued that the Footbridge problem tends to engage our emotions to a greater extent than the Trolley problem, and that our emotions deflect us from the utilitarian judgment in the Footbridge problem.

Some philosophers have argued that there are morally relevant distinctions between the two problems. Foot (1967) drew attention to the *doctrine of double effect*, first suggested by the medieval scholastic philosopher Thomas Aquinas, according to which harm is acceptable if it occurs as a foreseen but unintended consequence of an action serving a greater good, as in the Trolley problem, but not as a means to an end, as in the Footbridge problem. Quinn (1989) argued that the difference in responses is justified by the *doctrine of doing and allowing*, according to which pushing the man off the bridge is unacceptable because the harm results from intentional action, rather than from an omission, or failure to act. However, some psychologists have argued that the distinction between omission and commission is the result of a psychological bias (e.g. Ritov & Baron 1992; but see DeScioli et al., in press).

Others have proposed a *universal moral grammar* or UMG (Hauser 2007; Mikhail 2007), according to which normative moral principles, such as a prohibition of killing, are arrived at by an unconscious computational model, analogous to Chomsky’s (1995) universal grammar for human languages, this grammar being in accord with the doctrine of double effect.

How could moral problems possibly be freed from normative considerations? Perhaps some theories of higher mental processing can manage without such considerations, but it is hard to see how this could (or why it should) be generalized to all domains of research.

ACKNOWLEDGMENTS

Our research is supported by Grant No. AH/H001158/1 from the Arts and Humanities Research Council, United Kingdom. We are grateful to Ken Hughes for helpful comments.

Why rational norms are indispensable

doi:10.1017/S0140525X11000641

Ulrike Hahn

School of Psychology, Cardiff University, Cardiff CF10 3AT, United Kingdom.

HahnU@cardiff.ac.uk

<http://psych.cf.ac.uk/contactsandpeople/academics/hahn.html>

Abstract: Normative theories provide essential tools for understanding behaviour, not just for reasoning, judgement, and decision-making, but many other areas of cognition as well; and their utility extends to the development of process theories. Furthermore, the way these tools are used has nothing to do with the is-ought fallacy. There therefore seems no basis for the claim that research would be better off without them.

It is uncontroversial that a full understanding of behaviour involves multiple aspects. Psychology seeks to identify lawful regularities: We seek to understand the “what” of behaviour such that we can predict it. We also seek to understand why these regularities obtain. This involves two distinct kinds of causal explanation: (1) an understanding of the mechanisms/processes that give rise to the behaviour, and (2) functional explanation, that is, an understanding of why this behaviour and not others. Finally, psychology considers how such understanding allows performance to be improved in practice.

For all these questions, normative standards, that is, characterizations of how something “ought to be,” seem indispensable. Trivially, performance cannot be improved without knowing what would count as “better.” Likewise, functional explanation will make reference to the fact that a behaviour maximizes some “desirable” criterion – where both “desirability” and “maximization” typically invoke normative considerations. Even the basic task of identifying behavioural regularities cannot afford blindness to normative considerations.

Rational standards provide essential interpretative tools. Human behaviour typically affords many different interpretations.

In daily life we resolve this ambiguity with the help of “the principle of charity” (e.g., Govier 1987). Given multiple possible interpretations of an utterance, we pick, as a default, the interpretation that renders it most sensible. Of course, this interpretation may prove wrong, and further evidence may lead us to abandon it; that there are default orderings over possible interpretations, however, simplifies massively the task of understanding.

Moreover this has consequences not only for the interpretation of a single utterance or act, but also for prediction. How we interpret what someone has said or done has direct implications for what we expect that person to do in other, similar situations. Furthermore, even without specific knowledge of an individual we can often make reasonably accurate predictions just on the basis of what would be “sensible” (though again, there is no guarantee that these predictions will be correct).

The principle of charity has both informal and more formal manifestations within psychological research. Informally, whenever we observe something surprising, we should ask ourselves whether there is an interpretation of participants’ behaviour that renders it sensible and hence predictable. This may lead to re-evaluation of how our task is understood by the participants, to a different interpretation of what they are doing, and to quite different predictions and observations in response to changes in task parameters. Demonstrations of the (unintended) influence of the pragmatics of the experimental situation in giving rise to seeming “errors” and “biases” are a prime example of such re-evaluations (Hilton 1995).

Likewise, there are clear examples where a focus on the functional, computational level gave rise to more detailed predictions than research focused on process level explanations. For example, Oaksford and Chater’s (1994) probabilistic approach to the conditional provided a level of specificity previously unknown in the field of logical reasoning by introducing quantitative predictions (Hahn 2009).

Ultimately, complementing this functional level of explanation with descriptively accurate process theory will provide even more accurate prediction by capturing not just those aspects where behaviour approximates the normative standard, but also its systematic deviations. But again, for identifying these underlying processes rational norms can provide powerful tools. This is illustrated by ideal observer analysis which has had tremendous success in the study of perception (e.g., Geisler 1987). Ideal observer models draw on the formal tools of probability and decision theory to specify a model of optimal performance, given the available input for a task. Actual human performance is compared to the performance of this ideal agent. In a process of iterative refinement, human performance and ideal observer are brought into ever closer correspondence by incorporating into the ideal observer details of the human system. This approach provides a tool for the elucidation of mechanism and process, embedded in an overall account that seeks to understand the system as “doing the best it can do” given the available hardware. In so doing, it inherently links behavioural prediction, mechanistic and functional explanation, and it might be viewed as a methodological formalization of the principle of charity. Similar programmes are found under the header of bounded rationality or bounded optimality within cognitive psychology. Here, it has been stressed how rational norms can aid the disambiguation between competing theories and assist in the identification of underlying cognitive universals above and beyond the demand characteristics of experimental tasks (Howes et al. 2009).

The fact that most of this work is outside the domain of classic reasoning, judgment, and decision-making research, calls into question Elqayam & Evans’ (E&E’s) claim that these latter areas would make more rapid progress by abandoning their focus on normative standards.

It also highlights the weakness in E&E’s central claims about the “is-ought” fallacy. Spelling out what counts as “sensible” is itself a non-trivial task in that there are a number of normative

frameworks – logic, probability theory, decision theory, ecological rationality, and so on. They can differ in the interpretive focus they provide on a given task, because they make salient different aspects of behaviour. However, lawful connections exist between these frameworks; they are not simply “competitors.” Even more importantly, choosing between them in the interpretation of behaviour is not an inference about their normativity. Classical logic, probability theory, and decision theory all have independent, normative justifications (e.g., in the instrumental rationality of “Dutch books”). Appealing to certain aspects of the data in order to claim, as do Oaksford and Chater, that, in standard logical reasoning tasks, participants behave as if considering these tasks to involve probabilistic inference is not a statement about whether probability theory itself is normative, but about whether it provides the standard that renders participants’ behaviour “sensible” and “intelligible.”

There is no is-ought fallacy here. The normative theory is being used to provide interpretation and explanation of the behavioural data. In so doing, particular aspects of it are highlighted as relevant. As always, such interpretation may or may not turn out to be correct and the process may even be bi-directional; mismatch between what we consider to be “sensible” and what we observe may lead us to alter our understanding of “sensible” and what factors a behaviour should take into account. However, elucidating “sensible” will always involve considerations above and beyond the immediate data themselves – only then can it be of any use in interpretation, explanation, and prediction. The data themselves are prompts, not normative justification.

Normative theories are essential tools for understanding behaviour; it would be foolish to set them aside.

Defending normativism

doi:10.1017/S0140525X11000616

Steven Hrotic

Center for the Study of Interdisciplinarity, University of North Texas, Denton, TX 76203-5017.

stevenhrotic@yahoo.co.uk

Abstract: Elqayam & Evans (E&E) argue that evaluative normativism leads to unacceptable research biases, and should be avoided. Though it is stipulated that the particular biases they discuss are cause for concern, this argument should not be generalized. The boundary between evaluative and goal-directed “directive” norms is difficult to define, and normative assumptions are an integral part of academic progress; moreover, the biases that result may have beneficial potential.

As an anthropologist, I find academic culture fascinatingly full of contradictions. Perhaps that is why, although I find this target article convincing in its details, I interpret its significance rather differently.

As I understand it, Elqayam & Evans (E&E) distinguish between a “directive ought,” in which the statement “one should *x*” is conditional upon the desire to achieve a specific goal, and an “evaluative ought.” The authors write: “It is only when formal systems are regarded as having *a priori*, unconditional value, that the ‘ought’ becomes an evaluative one” (sect. 2, para. 5). The absolute, unquestioned form, it seems, can lead one towards the academic sin of bias.

My first concern is this: In practice, the distinction between “oughts” is fuzzy. For example, does the researcher have to be personally conscious of the goal-directed rationale behind a disciplinary norm to avoid an evaluative ought, or is it enough that it exists? For example, in anthropology there is a clear bias against working for the military. Our disciplinary history shows why we would be cautious, but it would seem the bias is more widely

known than the specific reasons for it. Similarly, must the goal be made explicit to avoid the accusation of evaluative normativism?

Many of my immediate colleagues employ evolutionary perspectives. We often casually refer to the “purpose” of a trait, or the “reason” it was selected for. We know evolution is not purposeful – but a third-party observer could be forgiven for perceiving a teleological error on our part. Perhaps we should not interpret too deeply into researchers’ choice of words when they do not specifically address the meaning of “ought” in their work.

For example, E&E hold Gigerenzer and Todd (Gigerenzer et al. 1999) as an example of the proper way to avoid evaluative normativism. Gigerenzer, they say, “famously repudiates normative systems, arguing that heuristic rules of thumb outperform normative computations” (target article, sect. 1, para. 8; cf. Gigerenzer 2007). To my mind, Gigerenzer and Todd do not do away with normativism but rather substitute one norm with another: they clearly think that problem-solving “ought” to be cost-effective (i.e., considering not just accuracy but also effort and the costs of error).

If we stipulate that academia is a problem-solving institution, perhaps we should consider normativism using a similar cost/benefit analyses. How hard would it be to eradicate evaluative normativism? And when it expresses as bias, is it a “bad thing,” or could there be positive effects?

We should consider how deeply ingrained normative judgments are in the process of academic progress – even in the assumption that progress is “a good thing.” It is somewhat ironic that E&E appear to be making a normative claim when they say we should avoid normativism. One could object that their claim is directive – that this “ought” implies a reason: to further the scientific understanding of human behavior. But this merely points towards another a priori assumption: that understanding human behavior, and specifically understanding it in a particular way, is a “good thing.” There would seem to be a strong interaction between our evaluative “oughts” and academic progress (cf. Kuhn 1996). E&E say that normative research biases affect what researchers choose to study (sect. 6, para. 9); one could also say, normative research biases concentrate efforts. Perhaps progress can be described in terms of improved (or, at least, more useful) normative assumptions.

But I would make two stronger claims. First, bias may be a potentially useful part of academic method. Consider peer review, one of the primary evaluatory mechanisms that allows academia to function. Much effort is expended into making peer review “unbiased,” for example, by hiding authors’ identities from reviewers, lest they fall prey to the “Matthew” and “Matilda” effects that bias reviews in favor of high-status (Merton 1968) and male (Rossiter 1993) academics. While double-blinding conforms to the academic norm of objectivity, it also saves us from having to decide when sex and/or status biases can be justified. For example, it may be *useful* to allow a bias that favors women in fields where they are underrepresented (e.g., engineering, philosophy).

Second, our extra-academic *a priori* norms may alert us to potential conflicts between the intellectual merit of a practice within academia and the broader impacts on society. (My reference to the National Science Foundation’s evaluative criteria is deliberate.) For example, one might feel that the President of Harvard ought not to suggest that women may be less represented in the sciences due to innate differences in their mathematical ability (Bombardieri 2005); or find it inappropriate that the excavation of the Manhattan antebellum African American Burial Ground began without significant African American participation (Harrington 1993). In either case, regardless of our evaluation of intellectual merit, our extra-academic norms alert us to possible negative impacts.

In summary, E&E make what seems to be a valid point regarding a particular normative assumption. I would suggest that we not generalize from this, but consider all biases – even ones

that produce gross biases – individually, in light of their impact on academia and more broadly. Do as E&E have done, and make the a priori assumptions explicit, evaluate the biases they produce, and examine available alternatives instead of dismissing even evaluative normativism out of hand.

Cultural and individual differences in the generalization of theories regarding human thinking

doi:10.1017/S0140525X1100063X

Kyungil Kim and Youngjun Park

Department of Psychology, Ajou University, Suwon, 443-749, Korea.

kyungilkim@ajou.ac.kr cyber13@ajou.ac.kr

Abstract: Tests of a universal theory often find significant variability and individual differences between cultures. We propose that descriptivism research should focus more on cultural and individual differences, especially those based on motivational factors. Explaining human thinking by focusing on individual difference factors across cultures could provide a parsimonious paradigm, by uncovering the true causal mechanisms of psychological processes.

Elqayam & Evans (E&E) suggest that it is only within the paradigm of descriptivism that competing theories can be arbitrated. We agree with their argument; but, in order to more parsimoniously and appropriately arbitrate those theories, descriptivism should focus more on the role of studies on cultural and individual differences for the explanation of the human mind.

Normative considerations have critical limitations for understanding the human mind, especially on the computational level, because normativism does not provide insight into the variability in human cognition and behavior. Descriptivism is an alternative approach. For example, the dual-process or dual-system paradigm proposes two types of processes: (a) heuristic, rapid, parallel preconscious processes, and (b) analytic, effortful, sequential processes that correlate with general cognitive ability. The paradigm’s core idea is that research should focus on observing and explaining the thinking and reasoning that people do, without preconceived concerns about what they ought to do.

Nonetheless, the paradigm’s current form is still insufficient for understanding the human mind, because descriptivism research needs to focus more on cultural and individual differences, especially those based on motivational factors. A major limitation of the explanatory approach is that the observation and, thereby, the interpretation of theory construction is often idiosyncratic to a particular culture or group. Research has shown variations from a given prediction, regardless of whether it is normative or descriptive, in many domains and on different levels between cultures.

For example, researchers had thought that persons from East Asian cultures were more risk-averse in most domains, such as social risk, than are persons from Western cultures; but, surprisingly, the former show greater risk-seeking than Westerners do regarding financial risks (Hsee & Weber 1999). Furthermore, such cross-cultural differences are found not only in higher-level cognition but in other levels as well, despite human-cognition researchers’ intuition that the higher the level of cognitive thinking, the more the influence of individual and cultural differences in the resulting behavior. In contrast, observing fundamental differences between different populations on the perceptual level is not difficult. For example, Chua et al. (2005a) measured the eye movements of American and Chinese participants while they viewed photographs comprising a focal object on a complex background. The Americans fixated more often on the focal objects than did the Chinese, and the Americans tended to look at the focal object more quickly. In contrast, Chinese participants

emphasized background in their visual perception. Even in everyday events, Americans tended to focus more on characteristics to do with self and personal agency and intentionality, and less on emotionality than East Asians did (Chua et al. 2005b).

As these cross-cultural examples show, if every culture or group requires its own explanatory story, then there probably isn't any way to find generalizations about human cognition and behavior. Numerous previous studies observed a degree of variability consistent with this possibility. Obviously, this concept lacks parsimony in regard to explaining human thinking, but parsimony should be one of the core aspects of descriptivism. An alternative (and supplementary) approach is to identify psychological variables that differ between cultures or groups, which would help to unify research on cultural and/or group differences and research on other kinds of individual differences. This individual differences approach seeks psychological variables, differing both within and across cultures or groups, that explain previously observed differences in cognitive performance (Kim & Markman 2006; Weber & Hsee 2000).

In line with this idea, Kim and his colleagues (Kim & Markman 2006; Kim et al. 2007; Markman et al. 2009) have suggested that considering cultural difference can provide better generalizations about human cognition, by emphasizing the processes by which individual difference factors lead to behaviors. For example, Kim and Markman (2006) manipulated fear of isolation (FOI) in American college students and observed that the high FOI group showed a greater relative preference for dialectical and holistic proverbs, which reflect collectivist viewpoints, than did the low FOI group. Another typical study of cultural difference observed this pattern with regard to Chinese and American populations, respectively (Peng & Nisbett 1999). This result could establish a causal link between FOI and cognition. Another example is research into self-construal's effect on judgment. In Gardner et al.'s (1999) study, individuals from the United States and Hong Kong, when primed for independent self-construal, preferred individualist values to collectivist values. In contrast, those primed for interdependent self-construal showed the opposite pattern. More importantly, this priming procedure produced the same outcomes found in cultures promoting individualism and collectivism, respectively, while also causally linking self-construal differences and value differences.

Thus, differences in a motivation-based individual difference factor, such as FOI or self-construal, create different cognitive goals; and different cultures' differing cognitive and behavioral outcomes reflect these different goals. Probably this is why researchers observe variability and inconsistencies from/against any universal theory or explanation, regardless of whether its basis is normativism or descriptivism. Nonetheless, individual research has often explored only one or two aspects of the causal mechanisms of individual differences, the differences' influence on cognitive processes, and/or different behavioral outcomes between cultures.

The particular goals a person can have are culturally determined. However, the goal activation's influence on human thinking may be common across individuals. Individual difference factors might differ across cultures and influence behavior. Therefore, taken together, culturally universal cognitive mechanisms could still cause cultural differences in behavior. While the content of people's goals clearly differs across cultures, the motivational system's mechanisms of operation might be universal. Similarly, while cultures may emphasize different personality characteristics (on average), these characteristics' influence on thinking and behavior could be the same among members of different cultures.

In sum, we suggest research should shift away from assessing broad behaviors and toward assessing the psychological characteristics underlying a behavior's processes, by exploring individual difference factors within/between cultures. Taken at face value, an explanation of human thinking focusing on individual difference factors across cultures makes the theory merely complex; but, ultimately, it can provide a parsimonious paradigm, by uncovering the true causal mechanisms of psychological processes.

Norms and high-level cognition: Consequences, trends, and antidotes

doi:10.1017/S0140525X11000501

Simon McNair and Aidan Feeney

School of Psychology, Queen's University Belfast, University Road, Belfast BT7 1NN, United Kingdom.

a.feeney@qub.ac.uk smcnair01@qub.ac.uk

<http://www.psych.qub.ac.uk/Staff/Profiles/feeney/index.aspx>

Abstract: We are neither as pessimistic nor as optimistic as Elqayam & Evans (E&E). The consequences of normativism have not been uniformly disastrous, even among the examples they consider. However, normativism won't be going away any time soon and in the literature on causal Bayes nets new debates about normativism are emerging. Finally, we suggest that to concentrate on expert reasoners as an antidote to normativism may limit the contribution of research on thinking to basic psychological science.

Normative issues have the potential to bedevil our field (the study of thinking) and Elqayam & Evans (E&E) have done us a great service in laying bare many of the problematic consequences of taking normative theories too seriously. Here, we ask whether normativism has been uniformly harmful, whether the end of normativism is really nigh, and whether the antidote proposed by E&E may do more harm than good.

We are not as alarmed about normativism as are E&E, many of whose arguments concern the psychology of deductive reasoning, and conditionals in particular, where the problem of multiple norms seems to be very acute. However, there are other areas in the study of high-level cognition (for summaries, see Feeney & Heit 2007; Murphy 2002) where normativism has the potential to be equally problematic but descriptivism has held sway.

Even in the areas on which E&E focus, normativism has not been uniformly disastrous. We do find it plausible that there are entire literatures which would not exist were it not for normative considerations. For instance, it is unlikely that anything resembling the actual literature on base rate neglect would exist had there not been a preoccupation with Bayesian norms in the 1960s and 1970s (e.g., Kahneman & Tversky 1973; Peterson & Beach 1967). However, inspired by the gap between normative behaviour and what people do in base rate neglect experiments, very important findings have been described about the difficulties people encounter in representing statistical information. For example, we now know the importance of the way the problem is described in facilitating people's recognition of the set relations underlying statistical problems (see Barbey & Sloman 2007; Evans et al. 2000; Girotto & Gonzalez 2001). Extremely interesting claims about the importance of causal models in statistical reasoning have also been made on the basis of experiments using the base rates paradigm (Krynski & Tenenbaum 2007). We know that people tend to use base rate statistics that they have acquired via experience more than those given to them by the experimenter (Gigerenzer et al. 1988), and the study of base rate neglect has greatly increased our understanding of the role of inhibitory control in thinking (De Neys & Glumicic 2008). None of this work seems to have been carried out in an evaluative spirit, although each of the researchers coded their participants' responses in the standard, normatively determined way. Despite this, all of these studies can fairly be described as having contributed to our understanding of psychological processes. So even in the very select range of domains considered by E&E, normativism has had various consequences. These range from literatures almost coming to a standstill – as seems to be the case with the literature on Wason's selection task – to the continued productive use of a paradigm whose invention was rooted in Kahneman and Tversky's goal of showing that a particular normative theory is an inadequate psychological account.

By alluding to areas in the study of high-level cognition, such as inductive reasoning and categorisation, where descriptivism

rules, we do not mean to suggest that normativism does not have the potential to be perilous. Oaksford and Chater (2007), in their Bayesian analysis of reasoning have been concerned with deciding on the most appropriate norm and with the psychological mechanisms that might approximate that norm. Unfortunately, Bayesian analyses in other domains of high-level cognition (for a review, see Jones & Love 2011) have not paid as much attention to mechanism. It is true that some of these analyses are pitched at the descriptive level (see Krynski & Tenenbaum [2007] on causal models and base rate neglect), but many others work at a computational level (e.g., Kemp & Tenenbaum 2009). As Sloman (2007) has pointed out, computational Bayesian models also work as normative models, whether or not they are described in such terms by their creators. This is because implicit in this type of computational model is the claim that there is a single Bayesian account for a particular type of thinking. No doubt inspired by this insight, Fernbach et al. (2011) have recently described a normative model of causal inductive reasoning based on causal Bayes nets and shown that when people reason predictively, from cause to effect, their inferences do not conform to the prescriptions of the model. This is a very important demonstration for those of us who work on inductive reasoning; but it also feels as if history might be beginning to repeat itself, and rather than being at the end of normativism, we may be about to see another battle in a war that seems likely to end no time soon.

Finally, E&E suggest in a number of places in the target article that we should focus on expert reasoning and how it is acquired. We see several problems with this as an agenda for our field. First, the cognitive biases seen in experts (defined, of course, with reference to some normative theory) are the same as those seen in naïve reasoners (see Bornstein & Emler 2001), so there may be very little to be gained from the exclusive study of experts. Of course, one could study how expert reasoners become expert, but then, if experts display the same biases as naïve reasoners, intervention is clearly required, which necessitates debate about norms. It seems to us that this debate will happen even if the goal of a meliorist intervention is instrumental rationality. This is because, in a domain where complex statistical thinking is required, experts may have to be taught how to approximate a norm in order to attain their goals. However, perhaps the most serious problem with the abandonment of naïve individuals by our field is that this would drastically reduce our contribution to basic psychological science. Thinking is central to what it means to be human and if E&E are correct that the old paradigm doesn't work, then we must find ways to usefully study how naïve and expert participants choose, make judgements, and reason.

ACKNOWLEDGMENT

Simon McNair is funded by a Ph.D. award from the Department of Education and Learning, Northern Ireland.

Norms, goals, and the study of thinking

doi:10.1017/S0140525X11000628

Raymond S. Nickerson

Department of Psychology, Tufts University, Medford, MA 02155.

r.nickerson@tufts.edu

<http://ase.tufts.edu/psychology/people/Nickerson.htm>

Abstract: Elqayam & Evans (E&E) argue that the major objective of research on human thinking should be the development of descriptive theories, and they challenge *normativism* – “the belief that people ought to conform to a normative standard” (target article, sect. 1, para. 10). I contend that although their argument for the importance of developing descriptive theories is compelling, normative theories are also important, not only for improving thinking but for investigating and understanding it as well.

Elqayam & Evans (E&E) have served up an impressive collection of thought-provoking ideas. Most importantly, in my view, they raise challenging questions about what the goals of research on thinking should be, and about the types of theories of thinking that researchers should strive to produce. They argue that the primary goal should be to discover the rules that describe how thinking is actually done, as distinct from how it *ought* to be done – that the theories we strive to develop should be descriptive, as distinct from normative. They contend that criteria for evaluating the quality of thinking are essential to the practical objective of improving thinking, but not to the objective of investigating and understanding thinking.

E&E's argument regarding the importance of description as a goal of research strikes me as compelling, but their dismissal of normativism, defined as “the belief that people ought to conform to a normative standard” (sect. 1, para. 10), does not. I find it easy to agree with the emphasis on description, but difficult to accept the idea that appeal to norms for the purpose of investigating and understanding thinking is inapt.

E&E acknowledge that “formal systems such as logic and Bayesianism have provided major incentives and inspiration to countless research paradigms” (sect. 7, para. 2), but they see these systems' pitfalls as outweighing their positives by enough to warrant consideration of proscribing normativism. “Can we,” they ask, “make do in reasoning and JDM [judgment and decision making] without normative theories altogether?” (sect. 7, para. 2).

In my view, the answer is no. Or, better, that we *could* make do without them, but only at unacceptable cost. I take a normative theory of thinking to be a theory that specifies how we should think, *taking our capabilities and limitations into account*. Some refer to this type of theory as *prescriptive*, and reserve *normative* to connote what would be prescriptive for a creature without human limitations. I find this distinction unhelpful; what may be normative for a creature without human limitations would not only not be prescriptive for us, it would not be normative for us either.

We want to know not only how we reason, but how we *should* reason. Maybe the way we reason *is* how we should, and maybe not. My sense is that the truth is somewhere in between – that our reasoning is not entirely consistent with what any reasonable normative model that takes account of human capabilities and limitations would prescribe, but also not quite as bankrupt as some presentations of the biases to which we appear to be prone seem to suggest.

The study of thinking is motivated by a variety of goals, description, and evaluation – the latter of which requires norms – among them. The results of countless studies show that neither traditional logic nor probability theory is descriptive of human reasoning as it is generally done. And in the opinions of many scholars, neither logic nor probability theory constitutes a plausible normative theory of how reasoning by mere mortals *should* be done. Disputes abound among philosophers and logicians regarding what should be considered normative in both deductive and inductive reasoning. This is illustrated, for example, by the long-running debates about the rules of conditional inference: Are conditional assertions truth functional? Should *If A, then C* be interpreted as the material conditional? Is the probability of the conditional the conditional probability? What does it mean for a counterfactual world to be maximally (or sufficiently) similar to the actual world? The literature is rife with arguments and counterarguments about such questions, most of which are never settled decisively.

That the challenge of developing a plausible normative theory of thinking is daunting has long been recognized by philosophers. MacIntyre (1988), who discusses it at length, takes the position that it is not possible to identify principles of rationality that will be universally recognized as valid. Still, we must have norms, even if any norms that we adopt are unlikely to be acceptable to everyone. We need norms, not only to live by, but in terms of which to understand our thinking – how it is good

and how it is not. Even to say that normativism is wrong is to imply the existence of some norm that justifies that assessment.

There are approaches to the study of reasoning and decision making that attempt to bridge the gap between descriptive and normative theories. One example is that of Fox and colleagues (Fox 2003; Fox & Parsons 1998; Fox et al. 2003), whose approach aims to deal realistically with situations in which people have to operate – in which problems often are poorly formed, relevant knowledge may be lacking, the circumstances may be fluid and time limited, but action is required. Perhaps the best chance of developing truly useful normative theories for such cases is via (descriptive) study of how the more effective reasoners and decision makers deal with them. This is the rationale, I believe, that motivates the naturalistic decision-making approach of Klein and colleagues (Klein 1998; Pliske & Klein, 2003; Zsombok & Klein 1997). In both cases, the goal is not only accurate description of how reasoning and decision making are actually done in real-world situations, but an understanding of the processes that can be used for prescriptive purposes as well.

E&E note several ways in which formal systems can be useful for psychological theorizing, and vice versa. This strikes me as an especially important part of their discussion. However, I wonder about the impression conveyed by their Figure 2 that the flow is primarily from formal systems to psychological theory and data rather than the reverse, and about E&E’s exclusion of validation from the ways in which psychological theory and data influence formal theories (see target article, sect. 7). Excluding validation from the ways in which psychology can influence formal theories raises the question of what the basis of the authority of formal theories could be, if not the warrant given by human judgment. One accepts the laws of (some) logic as binding, if one does, because one finds them psychologically compelling; to what else could one turn?

The “is-ought fallacy” fallacy

doi:10.1017/S0140525X11000665

Mike Oaksford^a and Nick Chater^b

^aDepartment of Psychological Sciences, Birkbeck College, University of London, London, WC1E 7HX, United Kingdom; ^bBehavioural Sciences Group, Warwick Business School, University of Warwick, Coventry CV4 7AL, United Kingdom.

mike.oaksford@bbk.ac.uk nick.chater@wbs.ac.uk
<http://www.bbk.ac.uk/psyc/staff/academic/moaksford>
<http://www.wbs.ac.uk/faculty/members/Nick/Chater->

Abstract: Mere facts about how the world *is* cannot determine how we *ought* to think or behave. Elqayam & Evans (E&E) argue that this “is-ought fallacy” undercuts the use of rational analysis in explaining how people reason, by ourselves and with others. But this presumed application of the “is-ought” fallacy is itself fallacious. Rational analysis seeks to explain how people *do* reason, for example in laboratory experiments, not how they *ought* to reason. Thus, no *ought* is derived from an *is*; and rational analysis is unchallenged by E&E’s arguments.

Elqayam & Evans (E&E) outline an argument that, if correct, would fatally undermine the research programme of the rational analysis of reasoning, in which we and colleagues have been closely involved (e.g., Chater & Oaksford 1999; Hahn & Oaksford 2007; Oaksford & Chater 1994; 1998a; 1998b; 2007; Oaksford et al. 2000). But it would apply equally to models in behavioural ecology (Krebs & Davies 1996), adaptive explanation in evolutionary biology (Sober 1993), ideal observer models in perception (Blake et al. 1996), Bayesian cognitive science (Chater et al. 2006), and rational choice theory and the whole of microeconomics (Kreps 1992).

What these explanations have in common is that they harness normative theories to explain descriptive facts. The structure of

the eye is explained because it forms clear images of the environment; the foraging patterns of a bee are explained as maximizing food intake; a person’s “information foraging” is explained as maximizing the amount of information acquired (Nelson 2005; Oaksford & Chater 1994; Pirolli 2007). But it should be immediately clear that E&E’s application of the is-ought fallacy appears itself to be fallacious. Such explanations do not derive *ought* conclusions (which the is-ought fallacy forbids), but derive *is* conclusions: attempting to explain facts *about* wings, bees, or people.

How then do norms enter in to rational explanations? Anderson (1990; 1991) provides an elegant account, in the context of psychological processes (see Oaksford & Chater 2007):

- Step 1. Specify precisely the goals of the cognitive system.
- Step 2. Develop a formal model of the environment to which the system is adapted.
- Step 3. Make minimal assumptions about computational limitations.
- Step 4. Derive the optimal behaviour function, given 1–3 above. (This requires formal analysis using rational norms, such as probability theory and decision theory.)
- Step 5. Examine the empirical evidence to see whether the predictions of the behaviour function are confirmed.
- Step 6. Repeat, iteratively refining the theory.

Note that norms, such as those from probability theory or decision theory, enter only in Step 4: they help derive optimal behaviour, *given the specification of goals, environment, and computational limitations*. No *ought* is hidden here either. Given that Steps 1 to 3 specify a well-defined problem (which is required, or else an optimal solution to the problem will be ill-defined), then the optimal solution (if there is one) is a matter of *fact* not evaluation: an *is*, not an *ought*. Take the familiar example of the travelling salesman problem: *if* the goal is visiting all towns in the shortest possible route; *if* the map is such and such; *then* it is a matter of fact (not evaluation) that the optimal route is thus and so.

Now E&E might object that the choice of norms to solve the problem at Step 4 can be challenged: Are there not competing normative theories? We suggest, by contrast, that Step 4 must always be well-defined (for the rational explanation to be viable); but that the assumptions that go in to Steps 1 to 3 can be challenged. Thus, in explaining behaviour in Wason’s (1968) selection task, accounts differ concerning whether people are optimizing information gain, or some measure of “utility” (i.e., there are differences over Step 1) (Oaksford & Chater 1994); and theories can differ about assumptions about environmental structure (Step 2) (Klauer 1999; Oaksford & Chater 2003). Theories do *not* differ about, for example, the axioms of probability theory.

Perhaps individual rational explanations are free of E&E’s charge; but might the rhetoric of rational analysis of reasoning fall into this trap? E&E suggest that it may, proposing that the following argument (which they attribute to us in their section 5.1, para. 1) commits the “is-ought” fallacy:

- (1) *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”)

This would indeed be a fatally flawed argument but, *pace* E&E, it is one that no proponent of rational analysis, including us, has ever proposed. Our work has been completely explicit that the normative basis of, for example, Bayesian inference is consistency arguments, such as Dutch book theorems (e.g., Chater & Oaksford, in press; Oaksford & Chater 2007).

We suspect that, in considering the above argument, E&E are conflating the conclusion that Bayesian probability is the appropriate *normative* theory, with the conclusion that rational analyses, using Bayesian methods at Step 4, is the best *descriptive*

theory. Rational analyses, like other scientific theories, are chosen by their fit to the data.

Such issues are crucial in the psychology of reasoning: in building a rational analysis of conditional reasoning with verbal materials (i.e., statements such as *if A, then B; not-B*, and so on). Specifying the *goal* of reasoning (Step 1) may, for example, be crucial: Is the aim to pick out conclusions that have a high probability, given the premises? Or is it to pick out only conclusions that are definitely true, given the premises? How do we interpret the materials that constitute the “environment” over which reasoning must occur (Step 2)? Do people interpret *if... then...* as a material condition (as in propositional logic), interpret it as an assertion about conditional probability (Edgington 1995; Ramsey 1931), or adopt one of many other possible interpretations. Different assumptions will lead to different rational analyses – as with any other type of scientific theory, competing accounts must be adjudicated, primarily by their compatibility with the empirical data. Note, crucially, however, that such assumptions are about *facts*: what *is* the goal of a person’s reasoning; how *do* we interpret the conditional. It is *not* about normative evaluations, such as what *should* be the goal of reasoning or how *should* people interpret the conditional. Such questions, while interesting, are not part of the scientific project of rational analysis.

We conclude, overall, that E&E’s injunction that we never infer an *ought* from an *is* is entirely correct; and entirely consistent with the program of rational explanation of cognition, and more generally with optimality explanations across the biological and social sciences.

Systematic rationality norms provide research roadmaps and clarity

doi:10.1017/S0140525X11000513

Niki Pfeifer

Ludwig-Maximilians-Universität München, Fakultät 10, Munich Center for Mathematical Philosophy, D-80539 München, Germany.

Niki.Pfeifer@lrz.uni-muenchen.de www.pfeifer-research.de

Abstract: Normative theories like probability logic provide roadmaps for psychological investigations. They make theorizing precise. Therefore, normative considerations should not be subtracted from psychological research. I explain why conditional elimination inferences involve at least two norm paradigms; why reporting agreement with rationality norms is informative; why alleged asymmetric relations between formal and psychological theories are symmetric; and I discuss the arbitration problem.

Purely psychological principles like limitations of cognitive capacity guide psychological theories. However, *a priori* rationality norms provide powerful roadmaps for the development of psychological theories. They play essential roles in the context of discovery and in the context of justification. In the context of discovery, they guide the research questions, the tasks, and the evaluation of the results. Wason’s selection task, for example, was developed within the normative framework of classical logic. From a probability logical point of view, this task does not distinguish between the conditional event and the material conditional interpretation, as both provide the same psychological predictions. In the context of justification, rationality norms are used to rationally reconstruct reasoning processes and the empirical data. The probability propagation rules of the modus tollens, for example, are formally much more complex than those of the modus ponens, which explains experimental data (Pfeifer & Kleiter 2009). Moreover, a formal theory provides a language that makes psychological theorizing precise. Therefore, I do not believe that “theories of higher mental processing would

be better off freed from normative considerations” (target article, Abstract).

Elqayam & Evans (E&E) argue that conditional elimination inferences are single-norm paradigms. This is true in the framework of classical logic. However, this does not hold in the framework of probability logic. There are at least two norm conflicts in the context of conditional elimination inferences (modus ponens, modus tollens, affirmation of the consequent, and denial of the antecedent). The probability propagation rules depend on the interpretation of the conditional. As an example, consider the probabilistic modus ponens. If the conditional is interpreted as a conditional event, then the probabilistic modus consists of the following inference:

If $P(A) = x$ and $P(B|A) = y$, then $xy \leq P(B) \leq xy + 1 - x$ is coherent.

If the conditional is interpreted as a material conditional, then the probability propagation rule is a different one:

If $P(A) = x$ and $P(A \supset B) = y$, then $\max(0, x + y - 1) \leq P(B) \leq y$ is coherent.

There are similar norm conflicts for the other three conditional syllogisms (Pfeifer & Kleiter 2005). Therefore, these conditional elimination inferences are not examples of single-norm conflicts.

The authors argue, citing Evans (1982), that “we should desist from the practice of reporting logical accuracy in reasoning tasks, and instead report what people actually did” (sect. 6, para. 7). I agree that empirical studies should report what people actually did. However, I argue that reporting (dis)agreement with rationality norms is important and informative. The basic question is the choice of appropriate rationality norms. After about a decade of reasoning research within the normative framework of classical logic, we can safely state that it is high time to consider extensions or alternative normative frameworks.

Coherence-based probability logic is one example (see, e.g., Pfeifer & Kleiter 2002; 2009; 2010). In this framework, the focus is not on whether or how people draw logically correct inferences about the logical validity of certain argument forms. Rather, the tasks instruct the participants to transmit the uncertainty of the premises to the conclusion. Thus, the focus is on the participant’s degree of belief in the conclusion. Probability logic allows for making psychological predictions precise and offers new psychological explanations of the inferences that people draw. The conditional introduction inference from “B” to “If A, then B”, for example, is *not probabilistically informative* under the conditional event interpretation of indicative conditionals ($P(B|A)$) and this is the reason why most people claim that one cannot infer a conditional from its consequent (Pfeifer & Kleiter 2011). Contrary to standard approaches to probability, this even holds in the special case where the premise is given for sure ($P(B) = 1$). Probability logic tells us which argument forms are probabilistically informative and which ones are not (Pfeifer & Kleiter 2006; 2009). “Probabilistic informativeness” is not an empirical term, it is a criterion derived within the normative framework. I argue that patterns of inferences beyond the conditional syllogisms should be investigated. Psychological plausible principles, like the ability of withdrawing conclusions in the light of new evidence and the defeasibility of everyday inferences, should be investigated. I agree that reasoning experiments should not focus on logical validity. Rather, the degrees of belief in the conclusions should be investigated.

E&E suggest there are asymmetric relations between formal theories and psychological theories and data (see their Figure 2). Indeed, formal theories inspire psychological theories. However, psychological data can inspire formal theories as well. Ford (2005), for example, investigates how experimental data informs artificial intelligence systems and has developed a formal system of nonmonotonic reasoning, which is inspired by psychological data (Ford 2004). Moreover, not only psychological

data arbitrate between formal theories: the converse holds as well. In the field of nonmonotonic reasoning, for example, there are many competing systems. However, System P (Kraus et al. 1990) is an example of a common denominator of rationality principles any system of nonmonotonic reasoning should satisfy. It makes sense psychologically to use such a system to arbitrate between formal theories in this field.

Another way of arbitrating between formal theories is to require psychologically plausible but minimal principles for rationality, such as *coherence* (see, e.g., Coletti & Scozzafava 2002). Coherence requires only avoiding bets that lead to sure loss. This criterion is much weaker than, for example, requiring maximizing expected utility. In my opinion, the relation between normative and descriptive components in a psychological theory of reasoning is a genuinely interactive one.

I agree with the authors that *learning* should be included in reasoning research. Moreover, there seems to be a consensus among reasoning researchers that the interpretation of the task material remains the same within participants. However, this is not the case in general. If participants solve several items of the probabilistic truth table task, the number of conditional event responses tend to increase from about 40% at the beginning of the experiment to about 80% at the end. Thus, the participant's responses converge on the competence answers (Fugard et al. 2011b). Assuming appropriate *bridge laws* that connect “is” and “ought” inferences, I avoid committing the is–ought fallacy, if I claim that rational reasoners *should* converge to the conditional event response.

A case for limited prescriptive normativism

doi:10.1017/S0140525X1100015X

Emmanuel M. Pothos^a and Jerome R. Busemeyer^b

^aDepartment of Psychology, Swansea University, Swansea SA2 8PP, United Kingdom; ^bDepartment of Psychological and Brain Sciences, Indiana University, Bloomington, IN 47405.

e.m.pothos@swansea.ac.uk <http://psy.swan.ac.uk/staff/pothos/>

jbusmey@indiana.edu

<http://mypage.iu.edu/~jbusmey/home.html>

Abstract: Understanding cognitive processes with a formal framework necessitates some limited, internal prescriptive normativism. This is because it is not possible to endorse the psychological relevance of some axioms in a formal framework, but reject that of others. The empirical challenge then becomes identifying the remit of different formal frameworks, an objective consistent with the descriptivism Elqayam & Evans (E&E) advocate.

The problem of pursuing normative theories arises because we think the cognitive system is successful. This reasoning can be summarized in something like, to the extent that cognition is so successful, whichever formal principles it is based on must be normative. But the “success” of cognition is extremely context dependent. Cognitive processing seems optimized with respect to certain types of problems (e.g., Shepard 1992), but, equally, it seems less so for other problems. In other words, the success of cognitive processing is context dependent. For example, it is arguable whether cognitive processing is optimal in the Wason selection task. We can use information theory to explain why people behave in the way they do in such problems, as Oaksford and Chater (1994) did, but this does not alter the fact that this is a deductive problem which has a (deductively) correct answer. Indeed, if cognitive processing were optimal across the board, the world should be problem-free (or in any case have far fewer problems than we do now). For example, people should always choose the right mortgage, they should never succumb to gambling addiction, and policy decisions should always be well-thought out and

optimal in terms of their respective objectives. Equally, depending on training, experience, and so on, different observers may approach the same decision-making situation in different ways. A certain apparent cognitive flexibility in reasoning and decision making appears highly adaptive and would seem to go against an assumption of an all-inclusive, prescriptive normativism. Thus, there are some genuine concerns regarding an assumption of a general prescriptive normativism.

So far we have simply shared some of the concerns of Elqayam & Evans (E&E) regarding prescriptive normativism. But then E&E proceed to argue that the pursuit of formal cognitive theories is consistent with a rejection of prescriptive normativism. It is here that we disagree. Let us first define a formal framework as a quantitative theory based on a set of interdependent axioms, such as Bayesianism, information theory, or Quantum probability (QP). Employing a formal framework for the description of cognitive processes basically implies adopting a set of interrelated postulates. Thus, to the extent that there is a belief that postulate A is psychologically relevant, then postulate B *should* be psychologically relevant as well. This is undeniably elegant in the sense that the framework as a whole can be tested. In other words, once a formal framework is adopted, then implied is a claim of internal prescriptive normativism, since it is assumed that *all* aspects of the formal framework have psychological relevance, at least with respect to a particular range of problems. To a large extent, this is exactly what is so appealing with approaches based on formal frameworks, such as logicism, Bayesianism, and information theory.

From such a perspective, as E&E note in their target article (sect. 7), it is not surprising that Gigerenzer and Todd did not justify their heuristics research program (Gigerenzer et al. 1999) on the basis of normative considerations. An individual heuristic, such as “take the best,” while undeniably successful, is just that: an individual heuristic. Its success is measured by its ability to outperform related heuristics. But it is not possible, for example, to justify such a heuristic in terms of related, manifestly true heuristics or computational intuitions. In other words, confidence in one heuristic does not usually imply confidence in another one; from a prior theoretical point of view, individual heuristics are somewhat interchangeable, even when they are highly successful (Gigerenzer et al. 1999).

Thus, overall, it appears that the pursuit of theories based on formal frameworks necessarily implies a limited, internal prescriptive normativism. We think that there is nothing contradictory in assigning a limited prescriptive normativism to Bayesianism, within a particular range of problems, as long as it is remembered that for an alternative range of problems Bayesianism may be a suboptimal framework. For example, some researchers have argued that for a certain range of decision-making problems, human behavior exhibits strong order or context effects and in such cases the QP theory is a more appropriate framework (e.g., Busemeyer et al. 2006; 2011; Pothos & Busemeyer 2009; Trueblood & Busemeyer, in press). QP theory is like classical probability theory, but for the fact that probability assessment is order (and context) dependent. For example, $P(A \wedge B) \neq P(B \wedge A)$. QP theory has been applied very successfully in the case of physical observables, exactly because of these properties. The QP research program in psychology aims to explore its utility in analogous psychological situations (i.e., situations which exhibit order, context dependence).

Our argument implies a piecemeal view of prescriptive normativism, which is far from the general prescriptive normativism E&E argue against. Is piecemeal prescriptive normativism problematic? In a scientific tradition arguably more successful than ours, physics, there are several normative frameworks (in a physical sense), which though very successful in their limited domains, are actually formally mutually exclusive. The most famous example is general relativity, which

assumes that space is curved, and quantum mechanics, which assumes that space is flat. Unfortunately for physicists, general relativity and quantum mechanics are mutually exclusive. While this is indeed the source of quite some frustration to physicists, it does not prevent them from doing extremely successful predictive science.

Therefore, overall, we think that cognitive scientists will continue to pursue theories based on formal frameworks, because of the elegance of building theories based upon a coherent theoretical framework. We also believe that inevitably this will lead to some limited prescriptive normativism, within particular ranges of problems. The scientific objective should then be one of establishing the range of applicability of different theories and indeed assessing the representational and process convenience of employing different theories in different domains (in cases where predictions of conflicting accounts converge; cf. Kuhn 1962).

Epistemic normativity from the reasoner’s viewpoint

doi:10.1017/S0140525X11000112

Joëlle Proust

Institut Jean-Nicod (EHESS-ENS), UMR CNRS 8129, Ecole Normale Supérieure, F-75005 Paris, France.

jproust@ehess.fr <http://joelleproust.hautetfort.com/>

Abstract: Elqayam & Evans (E&E) are focused on the normative judgments used by theorists to characterize subjects’ performances (e.g. in terms of logic or probability theory). They ignore the fact, however, that subjects themselves have an independent ability to evaluate their own reasoning performance, and that this ability plays a major role in controlling their first-order reasoning tasks.

Although theorists may not believe that there is a single right or wrong answer to a problem, reasoners often do. The target article seems to conclude from indeterminacy as to which normative system is being used by a participant, to the irrelevance of epistemic norms in reasoning. On the other hand, the need to accept norms in applied science is recognized. It is unclear what allows theoretical and practical claims (“pure” and “applied” science) to diverge in their most basic concepts.

Variety in norms can be understood in two ways: either in terms of several normative systems “fitting the bill” (target article, sect. 3), which is a problem for reasoning theorists; or in terms of which norm is most appropriate to a given epistemic task, which is a problem for individual reasoners. Note, however, that these two uses of “norm” are not clearly distinguished in the target article. When Elqayam & Evans (E&E) call a system that “fits the bill” “normative,” they mean that this system is appropriate, or optimal, for solving a given task. “Norm” can also be used, however, to refer to success or error *within* a given “normative system” – there are correct and incorrect ways of *using* a norm. The latter distinction can be clarified through the concept of a “constitutive rule,” that is, a rule that makes a particular cognitive task the task it is. Remembering accurately, remembering exhaustively, and checking whether a conclusion derives from a set of premises, each have different constitutive rules: Their outputs count as appropriate if they are, respectively, cases of accurate recall, exhaustive recall, and coherence tracking (Proust, in press).

This distinction between two uses of “norm” has consequences for the issue of normative conflict. By this term, E&E refer to the existence of several alternative ways of interpreting what a reasoner does in a given task. Such cases, however, generate other types of norm conflict that a theorist should not ignore. For

example, in “belief bias” tasks, participants need to be sensitive to norms of deductive coherence rather than to other norms such as fluency or relevance of epistemic content (Evans et al. 1983). The epistemic norm of interest, in this case, does not consist in a set of optimal procedures for solving a problem (a “normative system”), but rather in the informational constraints inherent to the cognitive goal embedded in the task: syllogistic closure rather than believability. Conflict may occur when several cognitive goals compete for saliency (for a task, for a participant). Participants need to draw on their prior experience to build a representation of the task (i.e., of its structure and cognitive goal), which may not be stable, and may not coincide with the experimenter’s.

The latter kind of norm conflict can be studied on the basis of participants’ self-evaluation in a cognitive task covertly involving norm competition (conflict monitoring is one of the main functions of metacognition; Botvinick et al. 2001). Such a study belongs to the metacognition of reasoning rather than to reasoning per se, but the case of metaduction suggests that control and monitoring processes play a considerable role in how a first-order task is processed (Reverberi et al. 2009). Note that sensitivity to a given norm need not be based on a conceptual type of understanding; familiarity with the task brings with it implicit access to the epistemic norm that constitutes it as the cognitive task it is (Proust, in press). A metacognitive study of reasoning, however, does not seem to be threatened by an is-ought fallacy, because the norm of interest is expressed in participants’ spontaneous self-evaluations and subsequent revisions. Metacognitive reasoners are motivated to correct what *they* see as a mistake; when they predict that a task is beyond their competence, they decline it (or wager against it in retrodictive evaluation) if they are allowed to (Koriat & Goldsmith 1996; Smith et al. 2003).

Should a naturalist reject the participants’ normative sense of error as fallacious? On the present construal of error as a violation of a constitutive rule, no appeal to *a priori* or irreducibly normative facts needs to be made. Natural regularities, such as feedback and regulation laws, can account for a subject’s sense of error (Proust 2009). On this view, violating constitutive rules cannot be a matter of individual preference, as some naturalists have claimed (Dretske 2000; Papineau 1999). From the observation that many different cognitive goals can be entertained, it is tempting to conclude that the epistemic norms can be chosen too: forming false beliefs might be a matter of preferences. Instrumental reasons to control one’s cognition (e.g., “I need to remember her name”) must, however, be distinguished from the normative requirements associated with the chosen type of control (remembering a name is adequate if it is correct). This reflects a contrast, as shown by Broome (1999), between a reason to act and a normative requirement. The first is an “ought” so far as it goes: you may or may not be right in thinking you need to remember this name. A normative requirement, in contrast, is “strict, but relative” (relative to your attempt to remember this name, you are strictly required to find the correct answer).

If these observations are correct, there are alternative normative systems when there are various *instrumental* ways of solving a problem. Such inter- and intra-individual differences in strategies can be studied using brain imagery (Goel & Dolan 2003; Houdé & Tzourio-Mazoyer 2003; Osherson et al. 1998). On the other hand, participants’ attempts to solve the problem *within a system* (logic, probability theory, etc.), involve strict constitutive requirements to which subjects need to be sensitive – and do become sensitive over time. A naturalistic explanation of epistemic norms can then be offered, on the basis of how reasoners monitor and control their own epistemic outputs. Eliminating norms from reasoning would amount to throwing out the baby with the bathwater: The reasoners and their motivation to obtain a correct answer should be of theoretical, and not just practical interest.

Naturalizing the normative and the bridges between “is” and “ought”

doi:10.1017/S0140525X11000525

Katinka J. P. Quintelier^a and Daniel M. T. Fessler^b

^aDepartment of Philosophy and Moral Sciences, and Research Unit, “The Moral Brain”, Ghent University, B-9000 Ghent, Belgium; ^bDepartment of Anthropology, and Center for Behavior, Evolution, & Culture, University of California, Los Angeles, Los Angeles, CA 90095-1553.

katinka.quintelier@ugent.be dfessler@anthro.ucla.edu

http://users.ugent.be/~kquintel

http://www.sscnet.ucla.edu/anthro/faculty/fessler

Abstract: Elqayam & Evans (E&E) suggest descriptivism as a way to avoid fallacies and research biases. We argue, first, that descriptive and prescriptive theories might be better off with a closer interaction between “is” and “ought.” Moreover, while we acknowledge the problematic nature of the discussed fallacies and biases, important aspects of research would be lost through a broad application of descriptivism.

Elqayam & Evans (E&E) present descriptivism as an alternative to normativism and as a way to avoid problematic inferences and research biases. Their proposal entails “identifying which terms are descriptive and which are deontic, and concentrating on the former” (see target article, Appendix: “Descriptivism”). They contend that “evaluative considerations need *only* be invoked [...] where the object is to improve human thinking and performance” (sect. 8, para. 5, our emphasis). Moreover, evaluative considerations are fine so long as norms *precede* research rather than follow from it. In line with this, E&E prefer to entirely avoid inferences from descriptive to normative terms (“is-ought” inferences).

While recognizing that E&E principally address the cognitive sciences, broader applications of their perspective could have detrimental consequences. In this regard, we want to call attention to well-supported interactions between normative and descriptive theories: Descriptive theories can evaluate normative theories, including the evaluative terms. In such cases, the focus in psychology is on normative terms while the aim is not to improve human thinking. We give two examples from naturalistic ethics, where scientific methods are employed to evaluate normative theories. We show how these examples impact fallacies and biases, most importantly the “is-ought” inference. Although the practice of naturalizing the normative is not exclusive to ethics, ethics is nonetheless well-suited to discuss the nuances of the “is-ought” fallacy.

Consider the first example: If we presuppose a principle that links normative with descriptive terms, then we can use descriptive data to reject normative theories. Although this might be uncontroversial, it has had, and still has, far-reaching consequences for specific normative theories. For instance, in the moral sciences, “cannot” is often taken to imply “ought not.” Findings from experimental psychology, neuroscience, child development, and psychopathology suggest that it is most likely impossible to perform moral actions without being emotionally motivated (e.g., Prinz 2007). Therefore, philosophers reject theories that hold that good actions are not motivated by emotions (e.g., a purely Kantian morality). In a similar vein, neuropsychological findings about consciousness and decision making (among others), problematize the existence of free will. To the extent that normative moral concepts such as responsibility, blame, and guilt presuppose the existence of free will, normative theories might need improvement.

This interaction has consequences for fallacies and research biases. As E&E illustrate, psychologists use empirical findings to argue against normative theories – an “is-ought” inference – and interpret responses in terms of their normative correlates – an interpretation bias. If a psychologist’s aim is to evaluate a descriptive theory, then these research practices are indeed fallacious and biased. However, evaluating normative theories opens up research questions that are interesting in their own right (i.e., above and beyond meliorism), and this goal may require interpreting responses in terms of normative correctness. Moreover, normative

theories can then be rejected on the basis of empirical data. Importantly, all this requires that researchers state their commitment to a principle that links normative terms with descriptive terms.

Second, depending on one’s epistemological or meta-ethical views, the meaning of a normative term may, at least partly, depend on descriptive data, such as how lay people use or understand the term. For instance, in a classical defense of moral relativism, Harman uses the following argument: “If we learn that a band of cannibals has captured and eaten the sole survivor of a shipwreck, we will speak of the primitive morality of the cannibals and may call them savages, but *we will not say that they ought not to have eaten their captive*” (Harman 1975; our emphasis). Provided we know who “we” is in this quote, we have an empirically testable claim: Do we really not say that they ought not to have eaten their captive? Whether the meaning of “ought” is relative or not may depend on such descriptive facts, constituting another way in which descriptive theories can evaluate normative theories. Indeed, the persistent tendency in philosophy to refer to lay people’s linguistic behavior is now a major impetus for experimental philosophy. Again, this practice provides a bridge from “is” to “ought.” Automatically rejecting such inferences, as E&E suggest, would then preclude naturalized and experimental philosophy.

Then again, one should not conclude that “is-ought” inferences are by default sound. When confronted with an “is-ought” inference, the investigator must sort out and reject wrong inferences. One should check if the authors beforehand explicitly stated a principle or philosophy linking normative with descriptive terms, as illustrated in our examples. Inadvertently inferring an “ought” from an “is” is wrong. However, if a reason is specified, the relevant question becomes whether one can agree with the principle or the philosophical rationale. For instance, 19th century evolutionary ethicists asserted that what is more evolved or developed is also morally better. This is, in the first place, wrong because evolution does not have a direction or purpose. In the second place, Moore (1903) interpreted these principles as attempts to state an analytic definition, entailing that the meaning of “good” (a normative term) was *entirely* cashed out in natural (descriptive) terms. Such attempts are also wrong: Philosophers and lay people use and understand normative terms as meaning that one cannot entirely cash them out in descriptive terms. On the other hand, we see no reason to reject an “is-ought” inference if one agrees with the stated reason, and therefore with the relevance of empirical data for normative theories.

Indeed, at least for naturalized philosophy, we relocate the burden of proof to descriptivism. Contrary to common wisdom, the classical arguments against “is-ought” inferences are not that far-reaching. Neither Moore nor Hume stated that all normative terms are entirely devoid of descriptive influence. Even Moore did not oppose an empirical research program that attempts to find what we ought to do. Hume (1739–1740) uses his famous “is-ought” passage to argue that it is not merely by reason or observation that we discover vice and virtue, but by sentiment. Precluding all “is-ought” inferences is a very common overreaction. In light of this, caution is in order when E&E’s position is read with an eye toward other areas of psychological inquiry.

Truth-conduciveness as the primary epistemic justification of normative systems of reasoning

doi:10.1017/S0140525X11000537

Gerhard Schurz

Department of Philosophy, University of Duesseldorf, D-40225 Duesseldorf, Germany.

schurz@phil.uni-duesseldorf.de

www.phil-fak.uni-duesseldorf.de/philo/personal/thphil/

Abstract: Although I agree with Elqayam & Evans’ (E&E’s) criticisms of *is-ought* and *ought-is* fallacies, I criticize their rejection of normativism on two grounds: (1) Contrary to E&E’s assumption, not every normative system of reasoning consists of formal rules. (2) E&E assume that norms of reasoning are grounded on intuition or authority, whereas in contemporary epistemology they have to be justified, primarily by their truth-conduciveness.

I highly appreciate Elqayam & Evan’s (E&E’s) keen criticism of *is-ought* and *ought-is* fallacies in the psychological literature on reasoning. In this comment I shed critical light on E&E’s account of *normative systems of reasoning* and of *prescriptive normativism*, which is understood by E&E as the view that human (rational) thinking *ought* to conform to a normative system of reasoning, *S*. Differing kinds of prescriptive normativism result from different normative systems *S* – but according to E&E (target article, sect. 1), *S* must always be a *formal* normative system, such as logic, probability theory, or decision theory.

My first objection concerns *formality*. While E&E are right that the normative systems that have dominated the psychology of reasoning have been *formal* systems, not every normative system of reasoning needs to be formal. For example, systems of spatio-temporal, mechanical, or chemical reasoning are not formal. Formality and universality are connected as follows: Formality implies universality, because formal systems abstract from the content of non-logical symbols (Schurz 2005). The other direction need not hold: A non-formal system of reasoning (e.g., reasoning in chemistry) can still be highly universal. For E&E, the reasoning systems of prescriptive normativism are typically universal, although they admit exceptions (sect. 1). In contrast, I want to suggest that the question of prescriptive normativism should be entirely *separated* from the questions of formality and universality. For example, Gigerenzer, who is classified by E&E as an anti-normativist (sect. 1), behaves like a prescriptive normativist when he recommends certain *heuristic rules* for prediction tasks, only that he regards these rules as neither formal nor universal (Gigerenzer et al. 1999).

My second objection concerns E&E’s overly narrow understanding of prescriptive normativism. This understanding rests on their distinction between (what they call) *evaluative* norms and *instrumental* (or directive) norms, and corresponding forms of rationality. While instrumental norms tell what one ought to do to achieve one’s personal goals, evaluative norms tell what one ought to do *simpliciter*. For example, Gigerenzer’s adaptive rationality is for E&E a kind of instrumental rationality, but Oaksford and Chater’s Bayesian rational analysis (see Oaksford & Chater 2007) is a kind of evaluative rationality. I think that E&E’s distinction is misleading because both kinds of rationality, “adaptive” and “Bayesian,” have in turn to be *justified* in the light of the *same superordinate epistemic norms*.

E&E (sect. 2, para. 4) write that their distinction is close to my distinction between the “normative” and the “valuative” (Schurz 1997), but this is a misunderstanding. E&E’s distinction is related, rather, to my distinction between *fundamental* versus *derived* norms, and corresponding values. According to the standard characterization in ethics (Frankena 1963), instrumental norms are norms that are *derived* from other – given and ultimately fundamental – norms, by the *means-end* principle that runs as follows: If (premise 1:) *A* is a (fundamental or given) norm, and (premise 2:) *B* is a necessary (or optimal) means for achieving *A*, then (conclusion:) *B* is a derived norm (Schurz 1997, Ch. 11.2, p. 8). While premise 1 is normative, premise 2 is descriptive-empirical in character. Therefore, the means-end principle explains how the findings of empirical scientists can become practically relevant *without* committing an is-ought fallacy: Empirical findings allow one to derive a multitude of derived norms from a small set of extra-scientifically given (intersubjectively accepted) fundamental norms (Schurz 2010, sect. 6).

E&E think that “instrumental norms” are not value-laden but descriptive (sect. 2), because they assume that the fundamental norms from which they are derived are *personal* goals. But for

many ethically important instrumental norms (such as “reduce carbon dioxide emissions”) the assumed fundamental norm has a high degree of intersubjective acceptance. This means that instrumental norms are also more-or-less value-laden: They *inherit* their value from the value of the fundamental norm or goal from which they are derived.

Concerning “evaluative” norms, E&E seem to assume that they are not in need of further justification by higher-ranking norms. This view is also mistaken. For example, E&E’s evaluative norm “Poverty should not exist” (sect. 2, para. 3) is instrumental for the higher-ranking norm that people should not suffer. Just the same is true for normative systems of reasoning. E&E, or at least the psychological tradition which they criticize, seem to assume that normative reasoning systems are accepted on the basis of mere intuition or authority. This view has its historical roots in the Kantian doctrine that principles of reasoning are *a priori* and *necessary*. This view is no longer considered tenable in contemporary epistemology.

Norms of reasoning have to be justified by *higher-ranking epistemic goals*. The most important fundamental epistemic goal – not mentioned in E&E’s article – is *truth-conduciveness*: Rules of reasoning should be a reliable means to increase true beliefs and avoid false beliefs (Bishop & Trout 2005; Goldman 1986; 1999; Leplin 2009, Ch. 2; Schurz 2009). This *gold-standard of epistemology* is, for example, used to justify the inference rules of classical logic by the provable fact that they lead in all possible worlds from true premises to true conclusions. Given recent findings that the conditionals in human reasoning correspond more to high conditional probability assertions than to exceptionless implications of classical logic (Evans et al. 2003), the same gold-standard has been used to argue that the system of (conditional) reasoning that maximizes truth-conduciveness is not classical logic, but system *P* or the stronger system *Z* of probabilistic default reasoning (Adams 1998; Goldszmidt & Pearl 1996; Leitgeb 2004; Schurz 1998; 2005; 2007). Incidentally, this fact contradicts E&E’s claim (sect. 3, para. 3 and Table 1) that conditional elimination inferences constitute a single-norm paradigm. Truth-conduciveness has also been applied to demonstrate the optimality of *meta-induction* in prediction tasks (Schurz 2008). Other generally accepted epistemic goals (besides truth-conduciveness) are *simplicity* and *efficiency*.

In conclusion, while I agree with E&E’s criticism of certain dogmatic views on normative systems of reasoning, I don’t agree with their rejection of normativism per se. A non-dogmatic and weak version of prescriptive normativism, in the form of normative recommendations regarding which typical human thinking processes are more and which are less truth-conducive and reliable, is still an important part of *practically applied psychology*.

Reason is normative, and should be studied accordingly

doi:10.1017/S0140525X11000549

David Spurrett

School of Philosophy and Ethics, University of KwaZulu-Natal, Durban 4041, South Africa.

spurrett@ukzn.ac.za

Abstract: Reason aims at truth, so normative considerations are a proper part of the study of reasoning. Excluding them means neglecting some of what we know or can discover about reasoning. Also, the normativist position we are asked to reject by Elqayam & Evans (E&E) is defined in attenuated and self-contradictory ways.

Ambrose Bierce, making a satirical joke, defined “Logic” as “The art of thinking and reasoning in strict accordance with the

limitations and incapacities of the human misunderstanding” (Bierce 1971, p. 211). The humour depends partly on how the definition confuses the normative aspect of logic, concerned with principles of *good* reasoning, with a descriptive fact about how people sometimes *do* reason.

The target article defends a related claim, although without humorous intent: that the study of human reasoning is better pursued as a descriptive enterprise, focused on how people actually think, without reference to normative considerations. However, Elqayam & Evans’ (E&E’s) own normative proposal (in this regard) demands too much, and its critical target is insufficiently precise.

Imagine someone who reads the entire target article. Before she does so, she says her aim is to figure out what it says. Afterwards, she announces that she has read a detailed argument for the view that normative considerations are an essential part of the study of human reasoning. I suspect that even the authors of the target article would be inclined, with the rest of us, to say that this person had made a mistake, that somehow she had got things wrong. Reasoning *does* aim at truth, and asking whether it succeeds or fails, or whether a reasoner commits errors, or does as well as he or she could, is always appropriate, if not always easy to answer. The view being defended by the target article requires that these evaluations, and existing discoveries guided by them, be abandoned when we are engaged in scientific study of human reasoning.

Consider, for example, the work of Kuhn (1991), based on interviews with people from various professions and walks of life. All were asked to give reasons for their views on unemployment, education, and crime. Analysing the interviews, Kuhn was generously prepared to count as a genuine reason any claim that was “(a) [...] distinguishable from description of the causal sequence itself and (b) [bore] on its correctness” (Kuhn 1991, p. 45). Only 16% of non-college subjects generated genuine evidence for their beliefs about crime, 28% for unemployment, and 29% for education. College subjects did better, but not perfectly (61% for crime, 53% for unemployment, and 66% for education). If we exclude normative considerations, we can’t regard Kuhn’s work as showing that many people reason poorly. It would, though, be nonsensical to read it descriptively as a discovery about what reasons actually are. Allowing outright irrelevancies and restatements of the conclusion to count as reasons empties “reasoning” of significant content.

E&E appear to recognize this problem when they say, regarding the work of Stanovich and West, that a normativist tendency is “to some extent moderated by the restricted applicability to elite reasoners” (sect. 5.2, para. 3). It is difficult, though, to see how a truly committed descriptivist could think there *were* such things as elite reasoners, which involves admitting that they are very *good* at something. Shortly after the “moderated” remark, the target article admits this, rejecting the notion of a reasoning error, since errors imply norms. The descriptive approach demanded here is blind in the same way as a physical study of two clocks, one fast and one accurate, concluding that the each clock consistently follows the laws of physics.

This objection may seem too strong. Don’t E&E allow functional considerations, and evaluation of reasoning processes by instrumental lights, but distinguish this from the error of normativism? They try, but don’t succeed. They say it is “only when formal systems are regarded as having *a priori*, unconditional value that the ‘ought’ becomes an evaluative one” (sect. 2, para. 5) and hence an instance of normativism. This sets the bar too high – many (perhaps all) norms of reasoning are not *a priori*. Most, maybe all, can be understood as conditional, to be used if there is enough time, unless some other criterion trumps them, as giving a fallible justification, and the like.

Other parts of the justification of “normativism” are similarly frustrating. “Normativism” judges reasoning by the standards of a “system,” but “system” isn’t defined. This judgement involves “measuring against” and “conforming to,” leaving it unclear

whether normativism judges process, or outcome, or both. Normative systems are “formal,” but “formal” isn’t defined. In any event, many plausible norms (“repeating the claim isn’t giving a reason”) are not formal. In conflict with the claim noted above that it is “only” *a priori* norms that are problematic (indicating a requirement), we are elsewhere told that “[p]rescriptive normativism can vary according to factors such as the *a priori* status of the normative system” (sect. 1, para. 8) (suggesting an option). It would help if it was clearer what normativism actually was.

Two points should not need emphasising. First, we do not know whether there is a single ideal theory of reasoning. It seems likely (because individual theories exhibit their own paradoxes and complications) that no one theory could be ideal, and that the best possible reasoning involves a varied toolkit. It does not follow that abstract theories of good reasoning have no role to play in assessing or interpreting actual reasoning. Second, we do not know exactly how our reasoning systems work. It seems beyond doubt that we don’t have a single reasoning process, that we can learn a variety of skills and techniques, and that sometimes one of our reasoning dispositions looks by the lights of one theory to be irrational, yet according to another turns out to be an adaptive shortcut or good move. Our computational limits are finite, but except in a few very vague ways, we don’t know how to say what kinds of reasoning are definitely beyond us.

The upshot of these two points is that it is genuinely difficult to say in some cases what good reasoning amounts to, and to say either how good actual human reasoning is, or how good it could be. Since “ought” implies “can,” we shouldn’t complain of people that they don’t do what they can’t do. The relevance of any particular norm to human reasoning, then, really is constrained by the limits of our abilities. It does not follow that we should stop caring about the difference between good and bad reasoning.

Normative models in psychology are here to stay

doi:10.1017/S0140525X11000161

Keith E. Stanovich

Department of Human Development and Applied Psychology, University of Toronto, Toronto, ON M5S 1V6, Canada.

keith.stanovich@utoronto.ca

http://web.mac.com/kstanovich/Site/Home.html

Abstract: Elqayam & Evans (E&E) drive a wedge between Bayesianism and instrumental rationality that most decision scientists will not recognize. Their analogy from linguistics to judgment and decision making is inapt. Normative models remain extremely useful in the progressive research programs of the judgment and decision making field.

This target article falls into an odd genre in psychological metatheory – a genre where authors try to *stop* the field from doing something. In the present case, the authors want researchers to stop using normative theory in a certain way. Such papers are indeed sometimes highly cited and discussed, but in my experience they rarely succeed in actually shaping ongoing empirical research. The evolution of research programs continues apace, determined more by the yield in new empirical findings than by arguments about metatheory. This is likely to be even more true in the present case, where the authors are trying to constrain a particular practice in one of the most fruitful and influential research programs in all of psychology – the heuristics and biases approach. Using normative theory to study decision making was useful in the past, and it continues to be useful. The present target article is also embarrassingly timed, coming on the heels of the 2008/2009 financial crisis. Empirical and theoretical work in behavioral finance – much of

it using normative approaches – now constantly appears in media articles attempting to explain aspects of the crisis (as well as famous scams such as the Madoff Ponzi).

In their target article, Elqayam & Evans (E&E) drive an odd wedge between instrumental rationality and Bayesianism (the latter defined as so-called normative rationality). The former is seen as appropriate, whereas the latter is seen as the type of approach that should be avoided. Contrary to this bifurcation, most theorists see the Bayesian strictures as being a crucial component of instrumental rationality, not a definition of rationality somehow separate from the instrumental view. Joyce (2004) gives what most decision scientists would see as the canonical view in *The Oxford Handbook of Rationality*:

Bayesianism provides a unified theory of epistemic and practical rationality based on the principle of mathematical expectation. In its epistemic guise it requires believers to obey the laws of probability. In its practical guise it asks agents to maximize subjective expected utility [SEU]. (Joyce 2004, p. 132)

In short, most theorists incorporate Bayesianism, in Joyce’s sense, into the notion of instrumental rationality, so the attempt to separate the two in the target article is a nonmainstream view. If it is only normative theory *outside* of approaches related to instrumental rationality (defined, as Joyce does, as *including* the Bayesian and associated SEU approaches) that is being attacked here, then 90% of the heuristics and biases tradition is spared, and the authors are left with a straw man that may apply to the four-card and a few other strange tasks. The whole attempt to drive a wedge between Bayesianism and instrumental rationality leads to another odd conclusion – that Bayesian norms do not help us as organisms achieve our goals. Precisely the reason that people should want to follow the axioms of utility theory (transitivity, etc.) as normative models is that failure to follow them means that a person is not maximizing utility. They should want to avoid becoming a money pump.

Another finding that is somewhat embarrassing for a view (like the one in the target article) that stresses the multiplicity of norms is that people most often retrospectively endorse the Bayesian and SEU norms that they violate (Shafir 1993; Shafir & Tversky 1995). In introducing the collection of Amos Tversky’s writings, Shafir (2003) stresses this very point: “The research showed that people’s judgments often violate basic normative principles. At the same time, it showed that they exhibit sensitivity to these principles’ normative appeal” (p. x). For example, Koehler and James (2009) found that non-normative “probability matchers rate an alternative strategy (maximizing) as superior when it is described to them” (p. 123).

E&E follow Cohen (1981) in using the grammatical competence/performance distinction borrowed from linguistics as a structural analogue for understanding judgment and decision making. However, the analogy to linguistics here is inapt because it ignores the fact that the cultural evolution of reasoning tools has no analogy in language. Because normative models are tools of rationality and because these tools undergo cultural change and revision, there is no idealized human “rational competence” that has remained fixed throughout history. Thus, the analogy to the linguistic domain is forced and inappropriate, as argued many years ago by Jepson et al. (1983):

The analogy between language and inductive reasoning fails to recognize that there is far more to language than grammatical competence. Important cultural inventions, such as writing or new vocabulary, increase the effectiveness of language use.... We contend that effective inductive reasoning is also a skill; that it depends on cultural innovations, such as probability theory. (Jepson et al. 1983, p. 495)

Psychologists are going to quite naturally study how beliefs are formed. That is inherently one of our topics. For example, the formation of stereotypes is a huge topic in social psychology. Of course, we *could* study beliefs without ever asking whether they were true. But what would be gained by refraining from asking if a particular stereotype studied in social psychology was generally true? When we do start to ask whether a belief is

true – whether it maps to reality – then we have brought in the issue of epistemic rationality and its norms. It is the same with the study of decision making. We can link decisions with hedonic states and utility without ever asking whether a hedonic state or utility was maximized. But I cannot see the point of refraining from this question. Once we don’t refrain from this step, norms of instrumental rationality (axioms of consistency and freedom from context effects) come in. In short, I do not think that psychology will heed any of these admonitions to refrain from normative language. I know I won’t. So while this target article is provocative and full of deep issues, it is not going to change the type of psychology my colleagues and I practice. In our lab, subjects *do* make, what we unabashedly call – errors.

Understanding reasoning: Let’s describe what we really think about

doi:10.1017/S0140525X11000148

Robert J. Sternberg

Office of Academic Affairs and Department of Psychology, Oklahoma State University, Stillwater, OK 74078.

Robert.sternberg@okstate.edu <http://osu.okstate.edu/acadaffr/>

Abstract: I suggest psychologists would more profitably study a totally different area of human reasoning than is discussed in the target article – the inductive reasoning people use in their everyday life that matters in consequential real-life decision making, rather than the deductive reasoning that psychologists have studied meticulously but that has relatively less ecological relevance to people’s lives.

It is really hard for me to concentrate on writing this commentary, because I’m worried about so many other things besides this. It is March 12, 2011, and my worries of the day include: (a) whether the reactors affected by the earthquake and tsunami in Japan will melt down, spewing radiation across that country and, potentially, the world; (b) whether Libya will end up continuing with the Gaddafi dictatorship or fall to the rebels; (c) whether there is any prospect of my getting a raise in the next several years, given the total mess of state governments in the United States; and (d) whether our 2-month-old triplets will get through the night peacefully or whether instead tonight will prove to be a “screamer.” Surely you can understand how these issues can seem more important than writing a commentary! I hate it when radiation leaks from a nuclear reactor all across the world! I also hate not getting a raise for several years on end.

Come to think of it, problems such as these are what frustrate me with the literature on human thinking, and in particular, on reasoning. The field has done a marvelous job of studying the problems that consume us just so little of the time. I’m used to this. My own field is human intelligence, and it too specializes in studying the problems that are easy to quantify (e.g., how should we measure IQ) rather than the problems that are important (e.g., if humans are so smart and their IQs increasing over the generations, why are they destroying the world in which they live?). As a field, psychology has mastered the study of the expedient instead of the consequential.

Although the psychology of reasoning has overwhelmingly concentrated on deductive reasoning, most problems we deal with in our everyday lives are inductive. And even when inductive reasoning has been studied, much of the work (including my own early work) has focused on IQ-test-like problems such as analogies, series completions, and classifications, or on problems that lend themselves to the potential application of quantifiable formulas. But the problems we most often face are not like these at all. What are the kinds of reasoning problems of consequence that people face?

If I continue to smoke, what is the likelihood that smoking will kill me or cause me grievous illness?

If I continue to gamble at the rate I am gambling, what is the likelihood I will drive myself and my family into financial ruin?
If I try meth or cocaine or some illegal drug once, what are the chances I will become addicted?

What are the chances that the sexual encounter I’m planning will end up giving me VD?

If I marry the woman (man) I’m dating, what are the chances that it will end in divorce, and especially, an ugly divorce?

If I keep gaining weight, what are the chances I’ll end up with Type II diabetes?

If I take this job, will I get tenure or end up jobless after a few years?

These are the kinds of problems that weigh heavily on people’s minds and have profound consequences. What is odd is how little psychologists of reasoning have given them any thought, perhaps because they are hard to study and harder, if one does study them, to lead to articles in prestigious journals that are more concerned with tightness of design and analysis than with the real-world consequentiality of the problems being studied.

Elqayam & Evans (E&E) study an important question in the psychology of human reasoning: To what extent are prescriptive, normative models useful in understanding how people think? Their conclusion that prescriptive normative models are not useful seems reasonable. But they do not go far enough. The problem is fundamentally that these models, right or wrong, apply only to a very narrow sliver of the kinds of problems that people face in their everyday lives.

When we study deductive reasoning, the question of whether there are normative solutions to which we should pay attention appears weighty and important. But few problems of consequence in our lives have deductive or even any meaningful kind of “correct” solution. Try to think of three, or even, one! When we look at inductive-reasoning problems such as those listed earlier, in which we try to predict future events that matter to us, the question of whether there is a normative solution seems rather remote. Obviously, there is not. It does not even matter whether we are talking about prescriptive or empirical normative models.

The problem, of course, is not limited to psychologists studying reasoning. Most of the standardized tests used in the United States are of the multiple-choice format or are short essays that have normative or at least empirically scorable “correct” answers. The advantage of such items is that they give the appearance of objectivity, can be easily scored, and yield impressive-looking statistics. The disadvantage is that they are remote from the kinds of problems people face in their lives, which may be why people are so much less than adequate in making real-world decisions. Schooling gives them so little experience in dealing with the kinds of problems of consequence they really will face. If one looks at government leaders, it is truly hard to find many (any?) who seem to be adept at inductive reasoning of the kind they need to govern.

We need a new psychology of human reasoning, one devoted to the world in which people live rather than the world in which academic psychologists live.

Normative benchmarks are useful for studying individual differences in reasoning

doi:10.1017/S0140525X11000562

Edward J. N. Stupple^a and Linden J. Ball^b

^aCentre for Psychological Research, University of Derby, Kedleston Road, Derby, DE22 1GB, United Kingdom; ^bDepartment of Psychology, Lancaster University, Lancaster LA1 4YF, United Kingdom.

e.j.n.stupple@derby.ac.uk l.ball@lancaster.ac.uk

http://psychology.derby.ac.uk/staff/Ed_Stupple.html

<http://www.psych.lancs.ac.uk/people/LindenBall.html>

Abstract: We applaud many aspects of Elqayam & Evans’ (E&E’s) call for a descriptivist research programme in studying reasoning. Nevertheless, we contend that normative benchmarks are vital for understanding individual differences in performance. We argue that the presence of normative responses to particular problems by certain individuals should inspire researchers to look for converging evidence for analytic processing that may have a normative basis.

Elqayam & Evans (E&E) bring a timely focus to the numerous pitfalls that can arise in studying human reasoning through a normativist research programme. While recognising many of these pitfalls, we nevertheless contend that normative accounts provide an invaluable reference point for understanding *individual differences* in performance and need not invoke the prescriptive notion that people “ought” to reason in a particular way. Instead, we argue that normative views can usefully inform theorising about the cognitive processes individuals employ without necessarily limiting the nature or scope of the processes considered. Our recent individual differences research seems to have key points of contact with the descriptivist programme recommended by E&E as a counterpoint to the normativist one. At the same time, there appear to be important points of departure that limit our ability to commit fully to their pure, descriptivist vision. Our commentary aims to explore some key points of alignment and non-alignment.

To begin, we fully agree with E&E that there is a fallacy in inferring that the mere observation of normative responding is diagnostic of underlying analytic processing (i.e., normative responses can be taken neither as necessary nor sufficient in identifying analytic processing). A case in point relates to the phenomenon of belief bias, where, for example, normative responses to syllogisms with valid-believable and invalid-unbelievable conclusions can provide no evidence whatsoever as to whether analytic processing has occurred, because either a heuristic or an analytic process would give rise to an identical (normative) response. Any attempt at identifying analytic strategies within this paradigm would need to focus on invalid-believable and valid-unbelievable conflict problems. With valid-unbelievable items the dominant response is normative conclusion acceptance, despite conclusion unbelievability. This suggests that analytic processing may be occurring with such items, with a key proposal being that reasoners attempt to “disconfirm” unbelievable conclusions by searching for a single, counterexample model (see, e.g., Evans’ Selective Processing Model [SPM]; Evans 2000). Because valid-unbelievable problems have no counterexample models, an analytic process would then lead to normative conclusion acceptance. For invalid-believable problems the dominant response is non-normative conclusion acceptance. The SPM explains this by proposing that reasoners attempt to “confirm” believable conclusions by searching for a single supporting model, which is readily available. Overall, this paradigm suggests a highly nuanced picture whereby normative and non-normative responding may arise from analytic or non-analytic processes.

Despite interpretational difficulties, we suggest that the presence of normative responses with particular problems can provide researchers with a vital stimulus to search for converging evidence for analytic processing, and should not be ignored. For example, according to the SPM, normative responses should never be seen for invalid-believable items; yet they do arise and it seems advisable for researchers to explore their underlying basis. Particularly useful in such research is the acquisition of fine-grained *process-tracing* evidence such as that derived from think-aloud protocols (Evans et al. 1993; Lucas & Ball 2005) and inspection-time and eye-tracking methods (Ball et al. 2006; Stupple & Ball 2008; Stupple & Waterhouse 2009). Our recent process-tracing studies of belief bias reveal that individuals who give logical responses to invalid-believable conflict syllogisms process these problems significantly more slowly compared to non-conflict problems and to individuals who respond non-logically. Coupling such evidence with neuroscientific data (De Neys

et al. 2008; Goel & Dolan 2003) indicating the involvement of the right lateral prefrontal cortex (a region specialised for cognitive monitoring) in normative responding, naturally leads to a view of individuals who respond logically to conflict problems as deploying analytic strategies sensitive to normative considerations. If such evidence additionally correlates with individual-difference measures that are predictive of normative responding (e.g., high working memory capacity, generation of alternatives and need for cognition; Torrens et al. 1999), then we assert that such triangulation points inescapably to an association between normative responses and analytic processing.

Whether or not this latter research process based around methodological triangulation is descriptivist under E&E’s view is, however, unclear, given the guiding role of normative considerations in the whole endeavour – right down to defining responses as “logically correct” or otherwise from the outset. We acknowledge that belief bias research is not immune from issues that pertain to selecting an appropriate normative benchmark. Indeed, Todd and Gigerenzer (2000) have argued that it may not be rational to prefer logic to belief. We also agree that it is problematic to equate the appropriate norm with the response given by the most cognitively gifted reasoners (Stanovich & West 2000b). Nevertheless, we see great merit in examining the cognitive processes employed by such gifted individuals. Moreover, De Neys (2006) has demonstrated that reasoners attempt to adhere to a normative standard irrespective of their cognitive ability, which suggests that these standards hold some genuine descriptive value. Indeed, we have obtained preliminary data indicating that normative training in a matching bias paradigm (as used by Stuppel & Waterhouse 2009) serves to increase logical responding and slow response times. Enhancements in logical responding have also been demonstrated with clarified quantifiers (Schmidt & Thompson 2008) and training in the concept of logical necessity (Prowse-Turner & Thompson 2009). These modifications to performance could be construed as merely bringing responses into closer alignment with an existing norm, but we would argue that there is not only merit in exploring methods of enhancing reasoning performance as measured against normative standards, but also in examining the associated cognitive processes and individual differences arising from such methods.

We therefore seek clarification as to where normativism stops and descriptivism starts. For us the situation seems more like a continuum than a dichotomy. We also admit to confusion over E&E’s claim that normative standards may be legitimate in a meliorist research programme designed to enhance reasoning skills but are inappropriate as a benchmark in research on underlying cognitive processes. This seems to lead to a conundrum if we wish to examine the underlying cognitive processes arising in response to the interventions deployed in meliorist research. If post-intervention responses align with normative ones, then how are we to go about understanding such changes *unless* we accept that they reflect normative principles?

Probability theory and perception of randomness: Bridging “ought” and “is”

doi:10.1017/S0140525X11000653

Yanlong Sun^a and Hongbin Wang^{a,b}

^aSchool of Biomedical Informatics, University of Texas Health Science Center at Houston, Houston, TX 77030; ^bDepartment of Psychology, Tsinghua University, Beijing, 100084 China

Yanlong.Sun@uth.tmc.edu Hongbin.Wang@uth.tmc.edu

Abstract: We argue that approaches adhering to normative systems can be as fruitful as those by descriptive systems. In measuring people’s

perception of randomness, discrepancies between human behavior and normative models could have resulted from unknown properties of the models, and it does not necessarily lead to the conclusion that people are irrational or that the normative system has to be abandoned.

To a great extent, the normatively “ought” inference criticized by Elqayam & Evans (E&E) is due to the discrepancy between the observed human behavior and the predictions by a normative model. When such a discrepancy is found, researchers either draw a conclusion that humans are irrational, because the discrepancy is deemed an error or a bias (e.g., Tversky & Kahneman 1974), or they adopt an “alternative-norm paradigm” to exonerate a certain human fallacy (e.g., Hahn & Warren 2009). In general, we disagree with E&E that “Normativism has played out its role in the history of the research on human thinking” (target article, sect. 8, para. 5). In our opinion, neither the “ought” inference nor any particular normative model should be abandoned. Rather, we take that what has biased and hindered research programs on human thinking is the practice of directly comparing *any* human behavior with *any* particular product of a normative model and being judgmental on human rationality based on partial comparisons.

We elaborate on the topic of “perception of randomness,” which has been “long perceived as normatively incorrect in the JDM [judgment and decision-making] literature” (target article, sect. 6, para. 10). Interestingly, this normatively “ought” inference resulted primarily from comparing human behaviors with a single normative model, the independent and stationary Bernoulli trials – for example, tossing the same coin repeatedly (for a review, see Oskarsson et al. 2009). However, the process of Bernoulli trials is not as simple as it may seem. It has myriad properties, depending on the variable being measured and the underlying parameters, and some of the properties are fairly novel to the researchers in psychology. Therefore, when the model is used to measure human behaviors, it is critical to distinguish between specific products and products with specific properties (e.g., Lopes & Oden 1987; Nickerson 2002).

One particular phenomenon in human perception of randomness is that people pay special attention to streak patterns (e.g., consecutive heads in a row when tossing a fair coin). By contrast, the process of Bernoulli trials predicts that the *frequency* of streak patterns is the same as non-streak patterns of the same length. This discrepancy has motivated numerous studies and continued to yield new findings on two fronts, the nature of human perception of randomness and the appropriate usage of normative models. (To name a few of these studies: Ayton & Fischer 2004; Bar-Eli et al. 2006; Burns 2004; Falk et al. 2009; Falk & Konold 1997; Gilovich et al. 1985; Nickerson & Butler 2009; Oppenheimer & Monin 2009.) In our own research (Y. Sun & Wang 2010a; 2010b), we found that it is critical to distinguish between two types of measurement when Bernoulli trials are used as the normative model: how often a pattern is to occur (*frequency*, measured by *mean time*) and when a pattern is to first occur (*delay*, measured by *waiting time*). It turns out that despite the equal frequency and mean time among patterns of the same length, streak patterns have the longest delay and waiting time.

Apparently, frequency and delay are two different concepts (for example, think of the experience of waiting for a bus; Gardner 1988, p. 63). Depending on the specific task environment, one measure may be more psychologically relevant than the other. Notably, the statistics of delay may provide normative measures to the descriptive accounts of randomness perception. For example, people judge the frequency of an event on the basis of how it is representative of the underlying population (representativeness), and how easily an example can be brought to mind (availability) (Tversky & Kahneman 1974). Thus, when people think of a random process, a streak pattern might be perceived as the most non-representative – it has the most uneven or clustered distribution over time – and the most unavailable – it is the most delayed pattern and most likely to be preceded by

non-streak patterns. Indeed, it has been reported that people's responses to streak patterns are largely consistent with the statistics of waiting time rather than the mean time (Oppenheimer & Monin 2009).

Since both mean time (frequency) and waiting time (delay) are derived from the same normative model, they provide objective measurements that are invariant across individual human subjects. This leads to another advantage of adhering to the normative models as we can embrace their fast and deep advancement in many different fields. Besides probabilities, the statistics of pattern times literally deal with properties of time and space. Recent neurological studies have drawn attention to some unified theories on representations of space, time, and numbers (e.g., Dehaene & Brannon 2010). It has been reported that people tend to assign a lower probability to an event that is more psychologically distant (Liberman & Trope 2008; Trope & Liberman 2010) and are sensitive to the tradeoffs between the delay in time and the probability (Luhmann et al. 2008; McClure et al. 2007). In financial theories, the delayed waiting time also means a greater variance of pattern times, and variance has been perceived as an essential component of risk (Markowitz 1991; Weber et al. 2004).

Nevertheless, we agree with E&E that normatively “ought” inferences should be taken with caution, and comparing human behaviors with selected products of a normative model does not immediately lead to conclusions on human rationality. In the case of pattern times, the distinction between the frequency and delay may provide us with a more precise tool to tease apart the task environment from which people's perception of randomness might originate, but it does not necessarily exonerate any human fallacies at the macro level (e.g., the gambler's fallacy, cf., Hahn & Warren 2009; 2010; Y. Sun et al. 2010a; 2010b). In addition, people's perception of randomness may not be reduced to a certain set of statistics. To gain a complete picture, we also need descriptive approaches to address perceptual and cognitive mechanisms that come into play, for example, the working memory capacity (Kareev 1992), perception of proportion and symmetry (Rapoport & Budescu 1997), and, subjective complexity (Falk & Konold 1997; Falk et al. 2009).

Normativism versus mechanism

doi:10.1017/S0140525X11000574

Valerie A. Thompson

Department of Psychology, University of Saskatchewan, Saskatoon, SK, S7N 5A5, Canada.

<http://artsandscience.usask.ca/profile/VThompson>

Abstract: Using normative correctness as a diagnostic tool reduces the outcome of complex cognitive functions to a binary classification (normative or non-normative). It also focuses attention on outcomes, rather than processes, impeding the development of good cognitive theories. Given that both normative and non-normative responses may be produced by the same process, normativity is a poor indicator of underlying processes.

Elqayam & Evans (E&E) argue that normative theories are problematic and have impeded theoretical development in the psychology of thinking and reasoning. In this commentary, I articulate additional arguments for the assertion that using normative correctness as a diagnostic tool has impeded our ability to formulate good cognitive theories of the processes that generate inferences. For this, I have assumed traditional models of probability and logic, although the comments are germane to any normative standard used to evaluate performance.

One of the primary goals of a normative theory is to serve as a standard that specifies how people should reason and to then

evaluate how well they achieve this standard. Reliance on normativism in thinking and reasoning research was evident from the first papers published on this subject (e.g., Wilkins 1928), and evolved into a long-standing effort to explain why judgements and inferences often defied normative prescriptions (e.g., Henle 1962; Kahneman & Tversky 1973) that remains to this day.

In fact, as a theoretical and empirical standard for cognitive science, the proscriptions of normativism fail us badly. As cognitive psychologists, our goal is to describe how situations, problems, and goals are represented; how knowledge is recruited to the problem space; the mechanisms by which inferences are computed; and how attentional processes and capacity limitations constrain these operations. Reliance on normative standards diminishes our ability to achieve these goals in numerous ways, not least of which is that it perpetuates reliance on an impoverished form of data, namely, a binary outcome variable. Thus, complex inferential processes that access a wide array of complex cognitive functions are reduced to a binary classification of either normative or non-normative as the (often sole) dependent measure. Alternative measures, such as reporting whether an inference has been accepted, whether a card on the Wason task has been selected, or whether a probability estimate is “close to a base rate,” at least preserve some element of the actual decision: its task-relevant outcome. However, like the normative measure, this single piece of information tells us little about the processes that produced it. As a discipline, the field of thinking and reasoning has been slow to incorporate other processing measures, such as eye tracking, reaction times, confidence measures, and so on, which would provide a much richer view of complex processes than can be gleaned from a single binary measure.

So impoverished is a normative index as an outcome measure that it can be flatly misleading. Indeed, it is easily demonstrable that the same processes may underlie both normative and non-normative responses. For example, on conditional reasoning tasks, both logically valid and invalid responses can be achieved by the same process, namely, the recruitment of counterexamples to the conditional (Thompson 1994). Conversely, a normative (or non-normative) response may be achieved by a variety of processes. A normatively incorrect response may arise because a salient heuristic process produced a prepotent answer; because the reasoner did not understand the instructions; because the reasoner tried but failed to compute a normatively correct answer; and so on. Similarly, a normatively correct response may be achieved because the reasoner knows and applies the relevant normative standard; because the reasoner guessed; because he or she used a strategy, such as the recognition heuristic (Goldstein & Gigerenzer 1999) that often produces correct answers; and so forth.

As an argument in favour of the utility of normative theories, one might point to the evidence demonstrating that persons of higher cognitive ability are able to suppress a tendency to give a “heuristic” response and produce normatively correct answers on reasoning tasks (Stanovich 1999). This research implies that there is value in the normative label, in that it embodies a quality that differentiates between high and low capacity reasoners. Again, however, having such a limited measure of reasoning performance means that these correlations between capacity and normative responses, in and of themselves, tell us little about how such suppression is achieved and may, in fact, tell us more about the reasoner's capacity to understand and encode instructions than about fundamental processes of reasoning. In support of this argument, consider that many concepts such as logical necessity are likely new to most reasoners, such that applying them consistently to a set of novel reasoning problems will certainly require capacity. In cases where the task instructions are less onerous, correlations with cognitive capacity may arise for any number of reasons, ranging from having access to appropriate normative rules, to recognizing their relevance in the current context, and having the motivation to do so (Stanovich 2009a).

Further evidence that normative versus non-normative responses tell us little about the mechanisms that produced them comes from a series of studies recently conducted in my lab. These studies employ a paradigm developed to study processes that monitor inferences and determine when additional analytic thinking is warranted. In this paradigm, participants are asked to give an initial, intuitive answer to a problem, make a metacognitive judgement about how right that answer feels, and are then allowed to take as much time as needed to give a final answer. The metacognitive experience that accompanies an initial decision predicts a lot about the effort that people put into solving a problem. Specifically, a strong feeling of rightness about a decision determines the amount of time spent thinking about a problem and whether or not the initial answer is changed in favour of another. Importantly, a strong initial feeling of rightness does not reliably predict whether the final answer is normatively accurate, despite the fact that it exerts substantial control over subsequent analytic reasoning. Moreover, in several studies, we have observed that when people do change their answers, they are often just as likely to change from a normatively correct answer to a wrong one, as vice versa (Shynkaruk & Thompson 2006; Thompson et al., under review). Understanding what motivates an answer change, what constraints an answer must satisfy to be retained, the information that is recruited to rework the answer, and so on, will tell us a lot about human reasoning. Knowing that the final outcome is normative tells us virtually nothing about the underlying mechanisms.

Neurath’s ship: The constitutive relation between normative and descriptive theories of rationality

doi:10.1017/S0140525X11000586

Michael R. Waldmann

Department of Psychology, University of Göttingen, 37073 Göttingen, Germany.

michael.waldmann@bio.uni-goettingen.de

<http://www.psych.uni-goettingen.de/waldmann>

Abstract: I defend the claim that in psychological theories concerned with theoretical or practical rationality there is a constitutive relation between normative and descriptive theories: Normative theories provide idealized descriptive accounts of rational agents. However, we need to resist the temptation to collapse descriptive theories with any specific normative theory. I show how a partial separation is possible.

The target article by Elqayam & Evans (E&E) reminds us of a plausible methodological norm: Do not confuse normative with descriptive theories! In many areas of psychology this distinction seems trivial. Models of memory or perception try to elucidate how these systems actually function, not how they should function. But what about theories of reasoning and acting, the classical domains of theoretical and practical rationality? Is it prudent to cleanly separate normative and descriptive theories here as well? My claim is that in psychological theories of rationality there is an intrinsically constitutive relation between normative and descriptive theories (Spohn 2002). I exemplify this claim in two different psychological research areas, one from the domain of theoretical, the other from practical rationality. I also argue that we need to resist the temptation to collapse descriptive theories with any specific normative theory.

Theories of *theoretical rationality* are concerned with how we achieve the goal of acquiring correct beliefs about the world. The close ties between normative and descriptive theories can easily be seen when we consider a scientist conducting an experiment. To explain the observed behavior, for example, randomization of subjects, it seems natural to invoke normative theories of how

experiments *should* be conducted. The reason why normative theories work as descriptive explanations here is that we understand the scientist as a rational, goal-oriented person. Everyday reasoning is, in contrast, rarely guided by explicit methodological knowledge. Nevertheless, people’s beliefs also have a normative force. People view themselves as motivated by reasons, and strive for optimality to achieve their goals. They distinguish between true and false beliefs; knowing they may err, they revise their beliefs and accept corrections. We would not say that a person believes that *A causes B* without simultaneously assuming that this person considers this proposition to be true. In sum, people conceive of themselves as rational agents, which make normative theories the natural candidate to explain their behavior.

Recent research about *causal reasoning* has obtained a wealth of evidence showing that we try to go beyond observations to obtain knowledge about causal relations in the world. People are rarely aware of how they acquire causal knowledge, but they understand what it means to respond to a causal query. Thus, a natural place to look for candidate explanations is normative theories of causal inference. In fact, virtually all currently competing theories of causal reasoning can be structured according to the preferred normative account motivating the theory (Waldmann & Hagmayer, in press). Causal reasoning cannot be modeled without some normative theory that tells us what causal judgments are.

Moral reasoning, which belongs to the domain of *practical rationality*, is another example of the constitutive relation between normative and descriptive theories. When we request moral judgments from subjects we do not want to learn about their preferences, wants, or inhibitions; rather, we want them to provide a normative evaluation of whether an act is *right* or *wrong*. Again, we conceive of subjects as rational agents who offer us responses to normative requests. Therefore, it is not surprising that an overview of research on moral judgment reveals that psychological theories use concepts from normative theories of morality (Waldmann et al., in press). For example, when interpreting responses to the famous trolley problem about whether it is permissible to sacrifice one person to save five, many theories focus on acts, outcomes, or values, concepts that have been highlighted as morally relevant in normative theories. And even if other non-moral factors are included in the explanations, we still interpret subjects’ responses as driven by the motive to provide a normative assessment. Otherwise, we could not say anymore that we are studying moral judgments.

Although I claim that in theories of rationality, normative theories are constitutive for the theoretical description of the target phenomena, I agree with E&E that there is a danger of oversteering the empirical adequacy of specific theories. Normative theories come in many variants, which compete. Their development is governed by factors such as coherence and consistency, which are less influential in everyday reasoning. Moreover, normative theories may restrict themselves to aspects of the target domain that are less relevant in everyday reasoning. Hence, it is unlikely that any specific normative position can be directly used as a descriptive theory. Responses in causal reasoning tasks have been interpreted as evidence for Bayesian causal network models, although it may often be possible to provide a more parsimonious account for individual phenomena by stripping away unnecessary untested assumptions implicit in these models. Similarly, the interpretation of responses as consequentialist or deontological in the trolley problem can often be more parsimoniously explained without attributing these global philosophical positions to subjects.

A sensible research strategy, therefore, seems to be to use one of the competing normative theories as a starting point, but to then ask whether all the assumptions inherent in these theories are empirically validated and necessary to explain the target behavior. Later, we can even go outside the realm of rationally relevant explanatory concepts. Rationality is certainly an idealization of thinking; many of our thoughts and actions are

influenced by factors that, on reflection, nobody would consider as legitimate. For example, Eskine et al. (2011) found that moral judgments about transgressions (e.g., stealing; taking a bribe) tend to be harsher when subjects were drinking a shot of a bitter beverage than when they were given water or a sweet beverage. The taste of a beverage certainly does not constitute an acceptable argument for a moral judgment.

In short, it may be necessary to rebuild the selected normative framework, delete components, sacrifice coherence and consistency, and even add non-normative factors. Some core components will stay invariant, to guarantee that we still model the target competency, such as causal or moral reasoning. The end result of the revision process may be a caricature of any accepted normative theory, but its normative foundation will still be discernible. Or as the philosopher Otto Neurath famously claimed: “We are like sailors who have to rebuild their ship on the open sea, without ever being able to dismount it in dry-dock and reconstruct it from the best components” (see Quine 1960, p. vii).

What is evaluative normativity, that we (maybe) should avoid it?

doi:10.1017/S0140525X11000604

Jonathan M. Weinberg

Department of Philosophy, University of Arizona, Tucson, AZ 85721-0027.
jmweinberg@email.arizona.edu

Abstract: Elqayam & Evans (E&E) argue that we should avoid evaluative normativity in our psychological theorizing. But there are two crucial issues lacking clarity in their presentation of evaluative normativity. One of them can be resolved through disambiguation, but the other points to a deeper problem: Evaluative normativity is too tightly-woven in our theorizing to be easily disentangled and discarded.

Elqayam & Evans (E&E) advise us to jettison the normative element in theoretical psychological investigations. To follow their advice, though, we need to know just what it is we are to jettison. There are two crucial areas of unclarity in their presentation of normativism. The first can simply be disambiguated, but the second indicates a deeper problem with their project.

First, some of their formulations of normativism involve a formal system (as in their Abstract: “human thinking reflects a normative system against which it should be measured and judged”). Elsewhere in the target article, normativism involves an appeal to something that is “a priori,” “unconditional,” or “universal.” But at other points, systems drop out, and all that is involved is a particular kind of “oughtness”; and indeed, at other points in the text, merely an appeal to rightness and wrongness, or “error” is deemed sufficient to count as normativist. It is dangerous to elide the difference between the notion of formal or a priori systems, and evaluative normativity itself, for there are evaluative normative theories that traffic in oughtness, but not on the basis of any such formal or a priori principles; epistemic reliabilism, for example (Goldman 1979; Kornblith 2003), and, as I will argue in a moment, the ecological rationality of Gigerenzer, Todd, and others. Some disambiguations are therefore in order.

E&E’s arguments mostly aim at universality (e.g., the “arbitration problem”) and the is-ought/ought-is inferences. So let’s tease a apart several different sub-claims of their anti-normativism that might be at issue:

(*Humeanism*) We should not make any hasty inferences from “is” to “ought,” or vice versa.

(*Anti-Universalism*) There is no unique correct evaluative normative framework applicable to all persons and situations.

(*Descriptivism*) We should not make any substantive (as opposed to merely “inspirational”) inclusion of evaluative normativity as a part of our psychological theorizing.

The authors’ arguments for the first two anti-normativisms are generally well-taken (see also, Weinberg 2007). But in some places they are perhaps over-zealous. For example, there are surely other explanations available for the focus on naive as opposed to trained subjects, on the very plausible presupposition that only a small portion of the population receives any such training; and so what we learn from studying such subjects may not generalize well. And, for that matter, expert reasoning populations have long been studied as well, as with Herbert Simon’s highly influential research. Also, the authors’ discussion of Humeanism could perhaps be more sensitive to a minor quandary they put themselves in: Having rejected a priorist approaches to the question of evaluative normativity, they either have to endorse some “is” facts as evidence for “ought” facts, or end up as full-blown skeptics about such normativity.

Disentangling Descriptivism from Humeanism and Anti-Universalism, however, does not yet render it clear, for “evaluative normativity” itself needs further clarification. One worry is that E&E have not actually succeeded in identifying a distinct form of normativity that can be cleanly set aside. They distinguish it initially from instrumental, bounded, ecological, and evolutionary forms of normativity, which they find unproblematically descriptive in nature. But they are too hasty in assimilating ecological normativity, for psychologists in that school seem to traffic in forms of correctness that cannot be boiled down to “Darwinian and Skinnerian algorithms.” Ecologists’ writings are rife with discussions of a “fit” between environment and mind, and of “success” or simply “good reasoning” that is neither constituted by (even if perhaps highly correlated with) the satisfaction of specific desires of individuals nor grounded in any appeal to what has historically promoted, or would today promote, reproductive fitness. One way of seeing how ecological normativity cannot be explicated in terms of instrumental normativity, is that the former is very often analyzed in terms of the relationship between a heuristic and an environment, without reference to any (even hypothetical) desires of some agent (see Over 2000). So, although E&E place ecologists in the “low normativism” zone, this is due to the ecologists’ conflating evaluative normativity itself with the use of normative systems; ecological rationality eschews the latter, but is robustly committed to the former. This also explains the difficulty E&E have with those researchers apparently deploying “oughtness” more than their location in the “low normativism” zone would expect.

E&E also go on later to invoke yet another form of allegedly unproblematic normativity, “epistemic rationality, in the sense of holding well-calibrated beliefs” (sect. 5.2, para. 4). But such calibration is not an unevaluative notion. Indeed, the logicians or Bayesians could easily claim that such normativity is exactly what they themselves are theorizing in terms of. For example, the reason, according to logicians, why you should not affirm the consequent is that it will generally lead you to believe falsehoods. The reason, according to the likes of Kahneman and Tversky, why heuristics tend to lead to biases is that they leave us susceptible to various forms of false beliefs.

There is a fundamental problem here: The sort of normativity that E&E want to set aside is simply too tightly wound around our psychological theorizing to be disentangled and discarded. That the authors find themselves writing in terms of instrumental rationality where agents and goals drop out of consideration altogether, and of the same epistemic rationality that their opponents would embrace, suggests that they are just as deeply enmeshed in evaluative normativity as everyone else, if in an unintentionally cryptic way. The authors are concerned that normative thinking has a too powerful, “biasing” influence on scientists’ minds. That very fact, however, suggests that trying to get scientists to swear off of normative thinking altogether may be simply impossible. Trying to do without it may only result in

such thinking being performed in subterranean ways, leading to violations of both Humeanism and Anti-Universalism. Wiser instead to *regiment* normative thinking. A methodological norm wherein researchers identify and articulate their evaluative commitments would put normativity out in the open, to be debated where fruitful, ignored where not, and allow for compensation for any biases. Where abstinence is unrealistic and perhaps even undesirable, we are better off adopting instead a norm of full disclosure.

Authors' Response

Towards a descriptivist psychology of reasoning and decision making

doi:10.1017/S0140525X11001440

Jonathan St. B. T. Evans^a and Shira Elqayam^b

^aSchool of Psychology, Faculty of Science, University of Plymouth, Drake Circus, Plymouth PL4 8AA, United Kingdom; ^bDivision of Psychology, School of Applied Social Sciences, Faculty of Health and Life Sciences, De Montfort University, The Gateway, Leicester LE1 9BH, United Kingdom.

jevans@plymouth.ac.uk <http://www.plymouth.ac.uk/staff/jevans>
selqayam@dmu.ac.uk <http://www.psy.dmu.ac.uk/elqayam>

Abstract: Our target article identified *normativism* as the view that rationality should be evaluated against unconditional normative standards. We believe this to be entrenched in the psychological study of reasoning and decision making and argued that it is damaging to this empirical area of study, calling instead for a descriptivist psychology of reasoning and decision making. The views of 29 commentators (from philosophy and cognitive science as well as psychology) were mixed, including some staunch defences of normativism, but also a number that were broadly supportive of our position, although critical of various details. In particular, many defended a position that we call “soft normativism,” which sees a role for normative evaluation within boundaries alongside more descriptive research goals. In this response, we clarify our use of the term “instrumental rationality” and add discussion of “epistemic rationality,” defining both as descriptive and non-normative concepts. We consider the debate with reference to dual-process theory, the “new paradigm” psychology of reasoning, and empirical research strategy in these fields. We also discuss cognitive variation by age, intelligence, and culture, and the issue of relative versus absolute definitions of norms. In conclusion, we hope at least to have raised consciousness about the important boundaries between norm and description in the psychology of thinking.

R1. Introduction

Our purpose has been to argue that the study of higher cognitive processes in psychology, especially in the large fields dedicated to the study of deductive reasoning and of judgement and decision making (JDM), has been unduly influenced by normative theory. We have suggested that most of the psychological paradigms involved have multiple normative theories that could be applied, and that psychologists in consequence ended up making dubious “is-ought” inferences to decide the right normative system on the basis of the observed behaviour. We have also argued that framing and interpreting

research within normativist frameworks has led to systematic biases in the way that research is conducted and has confused the task of constructing and testing descriptive accounts of the process involved.

Given the prevalence of normative thinking in these research fields, we were not expecting to win any popularity contest with these arguments. However, while some commentators are as resistant to our views as expected, we are pleasantly surprised by the amount of support that is also to be found in these commentaries. We also note that our target article has attracted as much comment from philosophers as psychologists. This should not be surprising since our article, as **Foss** remarks, is essentially philosophical in nature. Several commentators observe that our own paper is ironically normativist in that we discuss how psychologists should and should not practise their science. But there is no contradiction here – we are not constructing a theory of human thinking in the present discussion but rather debating the philosophy and practice of the science that underlies attempts to do so. Our unabashed contention is that we *ought* not to focus on “ought” questions while constructing such a theory. While our focus here is on a specific area of science – the psychology of higher cognitive processes – some commentators have also raised wider issues about epistemology and the philosophy of science. We address these as well as the more empirical concerns of the psychological commentators.

R2. Between normativism and descriptivism: Definitions and boundaries

In our target article we defended a notion of *instrumental rationality* defined as “Behaving in such a way as to achieve one’s personal goals” (sect. 2, para. 1) as a non-normative principle in our account of the human mind. However, it is apparent that the term has a normative meaning for others, exemplified by **Stanovich**’s claim that we attempt to drive a false wedge between Bayesianism and instrumental rationality, as violations of Bayesian principles can be shown to be sub-optimal by Dutch book arguments and the like. To clarify, we were not using the term “instrumental rationality” in a philosophical or absolute sense of *maximising* or *optimising* achievement of goals. Neither a *a priori* argument nor actual observation would support the view that people are perfect computational Bayesians. We meant only to convey that the mechanisms of the mind *strive* for the achievement of goals and thus achieve an (imperfect) degree of instrumental rationality. Hence, our use of the term is descriptivist and not normativist. We discuss Bayesianism in detail in section R5.2.

Both **Schurz** and **Hrotic** point to a fuzzy boundary between the evaluative and instrumental *ought*, arguing that what we considered evaluative oughts may themselves be instrumental to other values (Schurz), and questioning whether the researcher should be aware of such instrumental relations (Hrotic). If evaluative and instrumental *oughts* cannot be clearly differentiated, we lose the ability to tell normativism from descriptivism, respectively. A related complaint is that, by focusing on unconditional deontic norms, our definition of normativism is overly narrow (**Fuller & Samuels**; Schurz).

At the formal level, we concede the boundary may be difficult to define. But our concern is at the practical level of how researchers go about designing their studies, analysing and interpreting their findings. It is not usually hard to tell whether they are pursuing normativist or descriptivist goals (what they think they are doing is another matter). If they are focused on what it is that people do, and providing an account of how they do it, then this is descriptivist. If, on the other hand, they want to evaluate, score, and interpret the behaviour they study as “good” or “bad,” “right” or “wrong,” against some *a priori* norm, then it is normativist. A good example is the study of conditional reasoning which combines inferences seen as valid and invalid by standard logic (see target article). Researchers who report and account for the inferences people make are descriptivist. Those (still common in the developmental literature) who report the number of logically correct or erroneous decisions made are being normativist. This will usually be followed by a discussion which refers to some participants reasoning “better” than others, and so on.

R2.1. Soft normativism and the is-ought problem

Supporters of normativist research programmes often argue that the normativist and descriptivist questions cannot be disentangled. We believe this arises mostly from equating norm with competence, which we discuss in section R2.2. However, we observe that a number of commentators, including some who focus on important psychological questions in their own research, support a mixed model in which descriptivism and normativism go side by side, each contributing usefully to the other (Achourioti, Fugard, & Stenning [Achourioti et al.]; Douven; Fuller & Samuels; McNair & Feeney; Nickerson; Quintelier & Fessler; Schurz; Stuppel & Ball; Weinberg). We will dub this position “soft normativism.” According to this view, (a) normativism inspires research questions or whole paradigms that provide important descriptive findings, while (b) such findings feedback into the design and construction of normative theories. We deal with (a) later in section R7, but (b) still seems to leave us with the problem of is-ought inference. How can this be solved?

One suggestion, from Brase & Shanteau, is to adopt Platt’s strong inference strategy (Platt 1964). They claim that this method will “go a very long way towards avoiding the problems associated with prescriptive normativism” but they fail to explain how. Strong inference is, in our view, a very difficult research strategy to adopt, as people naturally focus on one hypothesis at a time in their hypothetical thinking (Evans 2010b). But even if researchers can be persuaded to pursue multiple hypotheses, based on multiple normative systems, how will this help? If we end up adjudicating between normative systems on the basis of behaviour then we are certainly going from “is” to “ought” no matter how many we consider.

The two main types of solution proposed by our commentators are (a) to avoid is-ought inference by attempts to bridge the “Humean bifurcation” (Frankena 1939) between the “is” and the “ought,” or (b) to adhere to the Humean bifurcation, proceeding instead from given first principles or values. Schurz seems to support a *first principles* solution: If we accept a small number of

fundamental norms as given extra-scientifically, we can then draw on empirical evidence to infer *derived* norms for a practically-applied psychology. He proposes truth conduciveness as such a fundamental epistemic norm. The solution is close to the one we offered for applied science in section 8 of the target article, and we welcome Schurz’s emphasis on acknowledging the extra-scientific nature of fundamental norms. However, we have our doubts about truth-conduciveness as a candidate for a fundamental norm, as we discuss in section R3.2.

In contrast, *bridging* solutions to the is-ought problem identify constructs (concepts, principles, postulates) which combine both descriptive and evaluative/deontic elements, and allow evaluative conclusions to be derived from them (and them only). For example, as Searle (1964) famously argued, if Jones *promised* to pay Smith \$5, then Jones *ought* to pay Smith \$5. The promise, a deontic speech act, combines the descriptive linguistic fact with the deontic, value-laden social obligation. We cannot infer “ought” from pure “is,” but we can infer “ought” from “is-with-ought,” and bridge constructs supply us with the latter. Such solutions are as prevalent among our commentators (Nickerson; Pfeifer; Quintelier & Fessler; Sun & Wang; and, more implicitly, Stuppel & Ball and perhaps Douven) as they are in the literature (e.g., Frankena 1939; Searle 1964; Williams 1985; but cf. Schurz 1997). The commentaries draw on a wide variety of bridge concepts, from “cannot implies not-is” (Quintelier & Fessler) to the relations between constitutive and regulative norms (Achourioti et al.).

Buckwalter & Stich discuss the influence of Goodman’s (1965) reflective equilibrium, brought by Cohen (1981) into the rationality debate. The idea is that norms and psychological evidence justify each other by being brought into coherence, creating a “virtuous” circle of justification. We note that reflective equilibrium is itself a bridge solution – it presupposes that such coherence between norms and behaviours is possible. Anderson’s six-step procedure of rational analysis (Hahn; Oaksford & Chater) is a derivative of reflective equilibrium: It seeks to establish coherence between behavioural observation and “optimal” (note the evaluative term) behavioural function. The Kantian “ought implies can” is a classic bridge construct, and so is its contrapositive “cannot implies not-ought” (Quintelier & Fessler). Similarly, Simon’s (1982) *bounded rationality* (Gold, Colman, & Pulford [Gold et al.]; Hahn; Nickerson), the idea that rationality should be judged relative to biological and cognitive limitations, derives from “cannot implies not-ought” (and see sect. R2.4). Another bridge construct is Davidson’s (1973) “principle of charity” (Gold et al.; Hahn), the idea that behaviour should be interpreted in a way that makes the agent appear rational (i.e., can implies ought).

The viability of bridge solutions has been debated for the better part of a century, and we will not solve it here one way or the other. As we show in section R6.2, however, bridge solutions preclude universal or absolute normative evaluation, leaving the more radical type of normativism vulnerable to internal inconsistencies.

R2.2. Normative theory ≠ competence theory

Both Buckwalter & Stich and Fuller & Samuels agree that maintaining the distinction between competence

theory and normative theory is important, but Buckwalter & Stich suggest that we have underestimated how difficult it could be to undermine the view that equates them. Alas, they are right. Despite our best efforts to the contrary (see sections 4 and 7 of the target article), many commentators persist either explicitly (**Foss; Pothos & Busemeyer**) or implicitly (**Achourioti et al.; Pfeifer; Sun & Wang**) in drawing the parallel. A special case of this is commentaries that propose or discuss Bayesianism as a normative system (**Stanovich; McNair & Feeney**), a position we discuss in section R5.2.

Buckwalter & Stich trace the origins of the normative–competence equation to reflective equilibrium (see sect. R2.1), and offer a persuasive counter-argument: Since psychological inference is subserved by two distinct systems (“minds” in our preferred terminology – see R4), there is no single coherent solution (and see R6). Another potential source of confusion seems to be that computational theories in reasoning and decision making often double as normative theories. One commenting strategy (**McNair & Feeney; Pfeifer; Pothos & Busemeyer; Sun & Wang**) is therefore to showcase a particular such theory or system, argue for its descriptive adequacy, and from that infer that normativism is indispensable. But this is not enough: such a strategy only demonstrates that the theory in question is a good computational-level (i.e., competence) theory.

The separation of norm and competence is, of course, crucial for our argument. To “de-normativise” competence theories, **Buckwalter & Stich’s** argument from dual processing provides us a strong first step. We clarify again that both levels, competence-computational (*what*) and algorithmic-processing (*how*), are in our opinion indispensable to a *descriptivist* psychology of reasoning. A descriptivist psychology is not just about processing; and conversely, normativism does not hold a monopoly on computational analysis. On the contrary, we argue that it is theorists who insist on equating norm with competence who drive an unjustified wedge between computational and processing analysis. For us, the *what* stops short of the (evaluative) *ought* – that the system computes a specific function does not mean that this function is normatively justified. To claim otherwise is to invite an is-ought inference.

R2.3. The role of formal theories

Both **Pfeifer** and **Nickerson** express concerns about our depiction of the relation between formal and psychological theories, believing that we had shown an asymmetric relationship, in which the main flow appears to be top-down. Such was not our intention. On the contrary, we argued, following Marr (1982), that the main type of relationship, constraints on theory construction, goes both ways. We are more than happy to concede that the same goes for inspiration. However, we point out that inspiration equally holds when formal theories are construed as purely computational. By reclaiming computational analysis for descriptivism, all we have to lose by giving up normativism is the evaluative function (and, of course, several dubious inferences and a handful of research biases).

We also referred to “formal systems” or “formal theories” as the level relating to computational analysis. Whereas commentators tended to agree with our

emphasis on the non-instrumental, fundamental, and universal evaluative nature of normativism, some challenged our assumptions about normative theories as formal systems (**Schurz; Weinberg**). Historically, psychological theories of higher mental processes have taken their normative standards from formal systems such as classical logic or Bayesianism, but we agree with Schurz and Weinberg that neither formality nor systematicity is crucial. [As an aside, this would obviate the need to define what a formal system means (**Spurrett**), which we have only done by extensional definition (i.e., by providing examples).] Instead, we propose that the crucial feature is the *a priori* nature of a principle (**Pothos & Busemeyer; Schurz**). The same argument we applied to formal systems is still relevant here: *a priori* principles can have a role in a computational level analysis, but cannot assume a role of fundamental value and yet be supported by empirical evidence. In practice, of course, computational level analyses will generally have to draw upon formal systems for their precision and detail.

R2.4. Distinguishing types of rationality

Lastly for this section, we reopen our classification of other types of rationality. Gigerenzer’s ecological rationality, the idea that thinking must be adaptive, proved rather challenging to evaluate. In the target article we classified it as descriptive, albeit with some misgivings, noting its “is-ought” inference, and the weakness of avoiding computational-level theories. (Incidentally, the latter was why – contrary to **Hrotic’s** evaluation – we never presented ecological rationality as the best way to conduct descriptivist research.) Thus, we can agree with **Schurz** that the “ought” of ecological adaptation seems to lie in-between an instrumental and evaluative ought, somewhere in the fuzzy chain connecting instrumental values to increasingly fundamental values (sect. R2); as Hrotic noted, Gigerenzer may simply be replacing one norm with another.

Commentators have agreed with us that, unlike other evaluative oughts in the literature, the ecological ought is informal rather than formal (**Pothos & Busemeyer; Schurz; Weinberg**); however, Schurz and Weinberg point out that this need not preclude its classification as normativist, and we have already conceded (in sect. R2.3) that formality is unnecessary to define normativism. Moreover, prototypical normativism tends to be absolute and universal, and, again, ecological rationality proves difficult to evaluate on that score. While Pothos & Busemeyer and Schurz emphasise the non-universal nature of individual heuristics, Weinberg argues, convincingly to our mind, that true instrumental rationality must be sensitive to the relative context of the agent and her goals, whereas ecological rationality authors often discuss “fit” or “success” in general and evaluative terms. All in all, the weight of evidence seems to point in favour of reclassifying ecological rationality as normativist. At least, as Schurz argues, it is no less normativist than rational analysis.

Brase & Shanteau argue that massive modularity is a better descriptivist system than hypothetical thinking theory (HTT), due to its sensitivity to evolutionary principles. This is beside the point of our article, which was not to promote any particular descriptive account: we referred to HTT only as an example of a descriptive theory which nevertheless still draws on formal theories

for its computational level analysis. It is, however, debatable whether massive modularity, grounded in evolutionary rationality, is a truly descriptive account. As **Weinberg** discusses, ecological rationality cannot be seen as instrumental in the absence of sensitivity to the desires of individual agents, and genetic goals can dramatically clash with the goals of the individual (see Stanovich 2004, for a detailed discussion). So the assumption on *a priori* principles that behaviour ought to be adapted to the environment in which an organism operates, is arguably normative, as noted above.

Last, we refer to Simon’s (1982) bounded rationality, in itself an uncontroversial and widely accepted notion, which several commentators refer to as a potential basis for normative analysis (**Gold et al.**; **Hahn**; **Nickerson**). We have judged bounded rationality to be unproblematic in terms of normativism, and we see no reason to retract this. We see bounded rationality as open to a range of interpretations. With its emphasis on costs and benefits of cognitive effort, bounded rationality comfortably sits with instrumental rationality. Bounded rationality can also serve as a basis for soft versions of normativism discussed above. However, since bounded rationality means that relative considerations of cognitive and biological limits need to be taken into account, it can never fit with the more radical types of normativism, which rely on universal and absolute norms.

R3. Epistemic rationality and self-knowledge

A number of commentators (**Dutilh Novaes**; **Proust**; **Spurrett**; **Stanovich**; **Waldmann**; **Weinberg**) have implicitly or explicitly raised the issue of *epistemic rationality* in contrast with the (descriptivist) notion of instrumental rationality that we focused on in our target article. We agree that the topic of epistemic rationality is an important omission in our account that needs to be rectified. First, we start with some comments about the nature of epistemic norms.

R3.1. Epistemic norms

Our critique focused specifically on the psychological study of higher mental processes (although some of it might be applicable to empirical philosophy as well). Some of the commentators, however, have addressed the broader issue of epistemic norms and their justification (**Dutilh Novaes**; **Proust**; **Quintelier & Fessler**; **Schurz**). To clarify, we never argued that epistemic norms are irrelevant for reasoning – only that it is fruitless for psychological science to focus on norms, especially when it actively interferes with the business of conducting research. As long as norms are kept apart from empirical evidence (e.g., by acknowledging, as Schurz proposes, their extra-scientific status), for us there is no issue.

Several commentators have also discussed the role of epistemic norms as constitutive rules justifying is-ought inference within the system they define, linking the constitutive function to the regulative one either explicitly (**Achourioti et al.**) or implicitly (**Fuller & Samuels**; **Proust**; **Waldmann**). Constitutive rules are definitional: they are what the system is about; regulative rules are deontic: they regiment behaviour. The classic example is

the constitutive rules of chess, which define what chess is as much as they regulate it. If chess rules are violated, one is no longer playing chess. Table manners, in contrast, are purely regulative, as eating exists independently. Searle (1964) argued that inference from is to ought can be justified for constitutive rules only and Achourioti et al., at least, readily accept the moderate relativism it imposes (see also Stenning & van Lambalgen [2008], and our discussion of bridge solutions in sect. R2.1). However, we cannot accept Achourioti et al.’s proposal that “without an ‘ought’ there can’t be an ‘is’ in the first place” (the “ought” in question being constitutive ought). Constitutive rules define what an institution *is*, not what it *ought* to be. Again, we take our example from language. Word order is a constitutive rule; “kicked Jack ball the” is not a poor sentence, it is not a sentence at all. This is a descriptive fact, not an evaluative judgement. In contrast, double negation is a purely regulative rule: “I don’t know nothing” is a sentence, albeit a non-normative one and condemned by linguistic purists. As this example clarifies, constitutive rules are competence rules, and therefore descriptive, not normative. Their “ought” is an epistemic ought rather than a deontic one.

R3.2. Epistemic rationality and normativism

A traditional view of epistemic rationality is that it is *truth-seeking* and several of our commentators believe that we need a normative account of it for this reason (**Douven**; **Fuller & Samuels**; **Proust**; **Spurrett**; **Stanovich**; **Waldmann**). As **Dutilh Novaes** observes, this reflects a Kantian tradition of logicist philosophy that has been very influential in psychology and still is, judging by some of the commentaries. Waldmann, for example, tells us that people distinguish between true and false beliefs, rejecting the false ones when evidence requires it – in the spirit of Popper, it seems. This truth-seeking notion of epistemic rationality needs to be complemented by a *truth-preserving* form of deductive inference (**Schurz**) in a traditional logicist theory of reasoning (and science). However, the latter has already been largely dispensed with by psychologists pursuing the “new paradigm” psychology of reasoning (sect. R5), so why not the former? For example, if we think of people as Bayesian, then we might expect them to hold all propositions with an infinitely variable degree of belief. Beliefs may not rationally be held with certainty, however: a prior probability of 0 or 1 remains fixed and cannot be revised in the light of new evidence. Hence, Bayesian theorists require both a different normative and computational account of belief acquisition than do logicists – another example of the multiple-norm problem.

However, we reject the idea that we need a normative account of belief at all. Our view is that cognitive representations are viewed not as veridical, but as being fit for purpose. They reflect the evolutionary history and instrumental needs of the organism. Take the case of the visual system. It is fallacious to think that there is a “true” picture of the world which our eyes and brains deliver faithfully to us. There is a mass of information in light, which could be interpreted and constructed in many ways. In addition, our visual systems have clear limitations. Some snakes have evolved the ability to see infra-red light because they hunt at night and need to sense heat. We have no

such facility. To take another example, our episodic memory system is designed not to preserve memories but to forget most of them rapidly after a relatively short interval unless they have great personal significance: This is adaptive, as generally we need access to the most recent memories which should not be confused with earlier ones. Our point is that cognitive representations are only veridical to the extent and in the manner required to serve our goals. Using the terms in a descriptive psychological sense, “epistemic rationality” is subservient to “instrumental rationality,” and an account of belief-generating systems is part of the *computational* account of human behaviour that cognitive science seeks. As Anderson (1978) pointed out some years ago, all computational accounts of cognitive functions constitute *representation-process pairs*. Hence, the problem of cognitive representations is mopped up in the computational-level accounts that cognitive scientists provide for instrumentally rational behaviours.

R3.3. Self-knowledge and self-regulation

Proust argues that epistemic norms are used by reasoners to regulate their own performance; **Waldmann** similarly claims that people view themselves as rational agents and act accordingly. **Achourioti et al.** make a similar claim but on a relativist basis: There can be no single norm since participants switch between systems and goals. Achourioti et al. conclude that a descriptivist account without norms is thus doomed to failure. Similarly, some commentators (**Gold et al.**; **Quintelier & Fessler**; Waldmann) propose moral judgement as a test case in which the psychology of higher mental processing cannot proceed without normative judgement. But we believe there is a fundamental problem with these arguments. We can have a descriptivist account of how people pursue their own epistemic goals and norms. A psychological theory of moral reasoning, for example, may include a good computational description of the functions involved in rendering moral judgement, but this cannot (and should not!) justify the moral theory in question. Put simply, psychology may show how people *think* on “right” and “wrong”, but it cannot be used to decide what *is* right or wrong.

It is certainly part of our own cognitive theory that people pursue epistemic goals in the “new mind” (see sect. R4) and this may indeed involve moral reasoning, higher order preferences and motivation to be rational (Stanovich 2008). We fail to see why a descriptivist account cannot include self-regulation by rules and systems if that is where the evidence leads us. But such systems need not be normative in the generally accepted sense. For example, **Bonnefon** describes a theory of his own in which he assumes that people follow folk axioms of decisions which are non-normative. There are, however, great difficulties in measuring the actual goals and rules that guide people’s reasoning, even at the explicit level. **Stanovich** points us to evidence that when confronted with their behaviour and a normative explanation, people will often judge their own performances to be erroneous in hindsight. But does this really provide evidence for the psychological reality of norms?

Unfortunately, people are highly prone to confabulation and self-deception when it comes to the accounts they

provide of their own thinking (Bargh 2006; Carruthers 2009; Evans 2010b; Wegner 2002; Wilson 2002). As an example, in one study the experimenters unkindly suggested to participants that one of several common incorrect (by the standard norm) solutions to the Wason selection task was the right answer (Evans & Wason 1976). In all cases, participants confidently agreed that the proposed solution was correct; not one individual objected that it was wrong. So the explanation of the “right” answer offered only has to be plausible and not actually normative to get the approval of participants that **Stanovich** describes.

Thompson also refers us to some of her own recent work on meta-cognition showing that the feelings of rightness that people have in their own judgements and inferences are unrelated to their normative accuracy. The more fluently and quickly an intuitive response comes to mind, the more likely people are to be confident in it. But this fluency may be a function of a non-normative factor such as the believability of a syllogistic conclusion (Shynkarkuk & Thompson 2006). A number of reasoning studies have allowed participants to give reasons for their choices, but to our knowledge no participant has *ever* declared a known cognitive bias, such as matching or belief bias, to be the cause of their decision. Instead, the evidence indicates that people give a reason which makes their choices appear rational in the context of the experimental instructions (Lucas & Ball 2005; Wason & Evans 1975). So, even if people talk as though they were rationally motivated to follow their own norms and rules, it will be very difficult to establish that this is what they are actually doing. For example, a person may justify their preference to vote for a political party on the basis of belief in a principle, such as free market economics. Psychological study may, however, show that the true cause is something else, for example, identification with a social group the individual would like to join, and the stated reason no more than a *theory* for the behaviour supplied by the media. People are not lying in such cases, but they are confabulating.

R4. Normativism and descriptivism in dual-process research

We support a strong form of dual-process theory which is a *two minds* theory (Evans 2010b), a development from earlier dual-system accounts. In this account, an old (intuitive) and new (reflective) mind co-exist within the human cognitive architecture, cooperating and sometimes competing in the control of our behaviour. While emotions and meta-cognitive feelings arise in the intuitive mind, we reject any suggestion that emotion should be contrasted with rationality, as stated by some dual-process theorists (Epstein 1994) and perhaps implied by **Gold et al.**’s discussion of moral reasoning. Instead, both minds are, in our sense, instrumentally rational, while often pursuing different goals by different mechanisms (see also Stanovich 2004). The old mind evolved early and combines evolutionary programming with experiential learning, thus providing a form of instrumental rationality which we share with many other animals – essentially we repeat behaviours that have been successful in the past. New mind rationality can also be seen as instrumental

but is driven by anticipation of the future, with cognitive structures that support mental simulations and consequential decision making. Its representations are only “epistemically rational” in the special sense that we describe in section R3.2: that is, they must be fit for purpose, so that their processing helps the new mind to achieve its goals.

The treatment of rationality within a dual-process framework is a contentious issue. We argued that it is fallacious to equate abstract reasoning and normative responding with System 2 or Type 2 thinking, and contextualised reasoning and cognitive biases with System 1 or Type 1 thinking (see also Evans 2008; in press a). However, some authors continue to do precisely this, even going so far in some cases as to equate System 2 with the operation of a mental logic (Ricco & Overton, in press). Although **Stanovich** focuses on his points of disagreement with us in his commentary, we know from his other writing that he agrees strongly with us on this point (e.g., Stanovich, in press). Although favouring a normative assessment of good reasoning, he nevertheless agrees that Type 1 processing can sometimes produce this while a number of conditions are required for Type 2 processing to do so, including perceiving the need for reasoning, and possessing the necessary cognitive capacity and “mindware” to achieve it (Stanovich 2010b).

Stuppel & Ball provide one cheer for normativism in their comments about dual process theory. They agree with us that it is dangerous to infer the presence of Type 1 and Type 2 processing from the normative correctness of responding, but suggest that norms can play a useful if limited role in the psychology of reasoning. For example, under certain conditions normative responding might be an unambiguous indicator of Type 2 processing effort. We actually agree with this point, but do not feel that it endorses normativism in the sense with which we are generally concerned. While the “new paradigm” psychology of reasoning (sect. R5) employs a number of methods, the traditional deductive paradigm still has its uses in a dual-process context. For example, Evans et al. (2010) compared deductive with pragmatic reasoning instructions, when people were asked to draw inferences from causal conditional statements. What was predicted – and found – was that participants of higher cognitive ability were more easily able to ignore beliefs and focus on the logical structure of the problems, but only if given deductive reasoning instructions. Our interpretation is that the instruction to draw logically necessary conclusions provides a need for Type 2 reasoning in order to address a novel and difficult problem. It has nothing to do with mental logic – one could get the same effect by giving people lengthy anagrams, for example. So with due care, as Stuppel & Ball suggest, correct answers by a relevant norm *can* sometimes be indicative of Type 2 processing effort. **Thompson**, who is highly supportive of the descriptivist approach, also notes in passing a similarly limited but useful application of norms in reasoning research. None of this amounts to the kind of evaluative normativism critiqued in our target article.

Buckwalter & Stich’s comments imply a relativist position which they relate to dual-process theory when they comment that, “If each person’s reasoning is subserved by two quite different systems, and the second system varies significantly among individuals and cultures, the

claim that an individual’s reasoning competence must be normatively impeccable is very implausible indeed.” Of course we agree, and this makes us all the more puzzled about **Stanovich**’s defence of normativism, given that he, more than anyone, has stressed the importance of acquired mindware (explicit rules and procedures) for Type 2 reasoning (e.g., Stanovich 2010b). Indeed, in his own comment, Stanovich asserts that “normative models are tools of rationality [that] undergo cultural change and revision.” But if we view cognitive psychology as a branch of biology that is concerned with uncovering the (culturally independent) mechanisms of the human mind, then the nature of such shifting rules and norms is only relevant in the limited sense discussed above – as an indicator of what a motivated individual is trying to achieve with high effort reasoning.

R5. The new paradigm psychology of reasoning

R5.1. From deduction to probabilistic accounts of reasoning: The paradigm shift

Sternberg complains about the domination of the deduction paradigm in the psychology of reasoning, and its irrelevance to everyday reasoning, commenting that “few problems of consequence in our lives have deductive or even any meaningful kind of ‘correct’ solution.” In effect, he makes a case for one of our proposed research biases: researchers focus on tasks for which clear norms are available. He is right that the psychology of deductive reasoning has been dominant, although fortunately other kinds are now receiving more attention, including the study of inductive inference (Feeney & Heit 2007), causal inference (Sloman 2005) and counterfactual reasoning (Byrne 2005). What may be less apparent to those outside of the field, however, is that the deduction tradition itself has undergone a recent Kuhnian revolution and paradigm shift. One of the drivers for this was the desire of researchers to integrate theoretical ideas about reasoning and decision making. While **Bonnefon** discusses this in the recent context, theorists have been attempting to do this since the early 1990s (e.g., Evans & Over 1996; Manktelow & Over 1991; Oaksford & Chater 1994).

The deduction paradigm has also collapsed under the weight of its own normative system – standard bivalent logic. As evidence accumulated that much reasoning was (a) logically erroneous and (b) highly sensitive to content and context amassed by use of the paradigm, the paradigm came under attack from within its own ranks. Researchers criticised logicism for its failure to deliver either an appropriate norm or a relevant description of human reasoning (Evans 2002; Oaksford & Chater 1998a). This eventually led a number of authors to declare recently that we have a “new” paradigm psychology of reasoning, freed from the constraints of logic and the deduction paradigm (see especially the two recent edited collections of Manktelow et al. 2011; Oaksford & Chater 2010). The emphasis is no longer on truth and deduction but on probabilistic and Bayesian approaches, degrees of belief, and pragmatic factors. Participants are often asked to draw plausible or pragmatic inferences or to make direct judgements about the believability of conditional statements, the plausibility of conditional advice, and so on. And, as **Bonnefon** points out, researchers are free to focus upon how

reasoning supports decision making, thus shifting the focus from truth and epistemic rationality towards action and instrumental rationality. Researchers are also becoming more interested in informal reasoning (e.g., Stanovich & West 2007), thus hopefully taking us somewhat closer towards the everyday reasoning that concerns **Sternberg**.

The new paradigm is, however, divided on precisely the issue with which our target article is concerned. While some contemporary reasoning researchers (e.g., **Thompson**) support our critique of norms and evaluative approaches or advocate great caution in the use of norms (**Stuppel & Ball**), others (e.g., **Stanovich**; and now less clearly, **Oaksford & Chater**) wish the new paradigm to adopt a normative framework, albeit no longer restricted to bivalent logic. By far the most popular alternative is Bayesianism, to which we now turn.

R5.2. Bayesianism – description or prescription?

Bayesian inference has two big advantages over logicism for the psychology of reasoning: (1) It permits people to hold and draw inferences about beliefs with variable degrees of uncertainty; and (2) it provides direct linkage between reasoning and decision making. People viewed as Bayesians would be able to draw upon all relevant sources of belief and knowledge in order to support consequential decision making. As **McNair & Feeney** illustrate, however, there seems to be some confusion as to whether the current focus on Bayesian theory is intended as a prescriptive normative account or merely a descriptive one.

Oaksford and Chater (2007) appeared explicitly to advocate Bayesianism as a normative system for human reasoning, as illustrated in detail in our target article. We, on the other hand, see Bayesianism as a broadly correct descriptive account of human reasoning and decision making, without any expectation of precise conformity to its computational rules (Evans 2007; Evans et al. 2008). Broadly speaking, the evidence indicates that people do represent degrees of belief, revise these beliefs in the light of evidence, and apply them in their decision making. Perhaps surprisingly, **Oaksford & Chater** (in contrast with **Hahn**) launch no strong defense of normativism in their commentary and take the stance that they are in fact descriptive rather than prescriptive Bayesians, thus avoiding is-ought inference. We are glad to hear it and note that we had made it quite clear in the target article that we have no objection in principle either to rational analysis or to the use of formal, computational theories in the manner they describe. However, we find it hard to reconcile the stance taken in their commentary with the numerous apparently evaluative uses of normative theory in their earlier writings.

In our view, the new paradigm will gain little from an approach that simply substitutes Bayesian norms for those of logic. But we also believe that researchers must go beyond demonstrations that Bayesianism is a good descriptive-computational account of aspects of human inference (Jones & Love 2011) for two reasons. First, an account of the cognitive processes that can produce such behaviour is needed at the algorithmic level, as there are many cognitive mechanisms and constraints to be studied. A major constraint, for example, is that people apparently focus on single rather than multiple hypotheses

when engaged in hypothetical thinking and mental simulation (Evans 2007). It should also be noted that there are areas of psychology where the Bayesian model fails spectacularly as a descriptive account. Descriptive Bayesianism can be hard to refute, since both beliefs and utilities can be assigned subjectively on an individual basis. However, if we examine a number of pieces of evidence in sequence, revising our beliefs after each one, it should make no difference at all to the order in which we see them (as an example, consider the evidence given to a jury during a courtroom trial). In contrast, research on belief updating shows marked primacy and recency effects on final belief judgements, as much research by social psychologists has shown. This has led researchers in this field to dismiss the Bayesian model out of hand and focus on alternative descriptive accounts (Hogarth & Einhorn 1992).

Despite the fact that many philosophers find Bayesianism a satisfying or even compelling normative account of instrumental rationality (see **Stanovich's** commentary and our section R2), we still see little value in prescriptive Bayesianism in the psychology of decision making. One might argue (*ought to is*) that decision makers are shaped by evolution and learning to be instrumentally rational, therefore they must be Bayesian and understood accordingly. But this is an epistemic, “stands to reason”-type *ought*, not a deontic one; and moreover, there is no empirical force to it. Bayesian models fit approximately well to behaviour in some tasks studied by psychologists but very poorly in others, as we have mentioned. The only role that prescriptive Bayesianism can serve in such circumstances is to enable its advocates to declare behaviour *irrational* and in need of remediation. But this does not seem to advance theoretical understanding; primacy and recency effects in belief updating, for example, are descriptive facts that must be understood with reference to the relevant cognitive mechanisms. Instrumental rationality in its *descriptive* sense needs to account for the individual agent in her specific epistemic context (as **Weinberg** argues in his commentary and we discuss in R2.4) and her set of (often conflicting) goals. We fail to see how normative judgement assists this task in any way.

R6. Cognitive variability

No one disputes that human cognition varies broadly and systematically. The question is what it means for normative analysis. We find it remarkable that empirical research on cognitive variability in higher mental processing has tended to come in two distinct camps. Individual differences, including developmental differences, tend to be associated with normativist approaches. Perhaps, as **Dutilh Novaes** has commented, this is due to the Piagetian-Kantian heritage (see our sect. R6.1). In contrast, cultural differences were generally taken up by the camp advocating moderate relativism. Taking full stock of cognitive variability inevitably leads to at least moderate normative relativism.

One of the initial reviewers of our target article complained that relativism “stalked the pages,” and **Weinberg** identifies an “anti-universalism” argument in the article. Although we have not explicitly defended relativism (our target was normativism), we readily concede that

moderate relativism works well with our critique. It is a potential consequence, especially if one wishes to keep a modicum of normative evaluation (sect. R6.2). Relativism has an undeservedly bad name in cognitive science, perhaps because many authors focus on the wilder types of radical relativism – the postmodern, Feyerabend-style “anything goes.” However, a moderate form of relativism, one which acknowledges both cognitive universals and cognitive variation, is not only easier to defend (see Stich 1990; Weinberg 2007), but may well be as indispensable for a descriptive psychology of human thinking (sect. R6.1).

R6.1. Individual and cultural differences in reasoning

From a normative point of view, some people reason better than others, or make better judgements and decisions. There are a number of factors which underlie such individual differences. One is age: children’s reasoning and decision making changes as they grow older. Performance on the raw score tests that underlie IQ projections, for example, increases up until about the age of 16 (actual IQs are age-corrected). Adults vary considerably in IQ, of course (and in working-memory capacity which is closely correlated); and this is one of several factors that have been implicated in the ability to produce normatively correct answers. Another important source of variation is culture, as an increasing number of research studies shows important differences in thinking styles; for example, between individuals in Western and Eastern societies (Buchtel & Norenzayan 2009).

Stanovich argues that, since norms of rationality are culturally transmitted tools, the analogy we draw to language is “inapt.” But the cultural transmission thesis makes our analogy to language more apt, not less. Cognitive variability is a well-established fact in all areas of cognition, and in language more than anywhere else – otherwise, how can there be thousands of different languages in the world? Even language universals, these prototypical, hard core cognitive universals, vary in how they are instantiated in specific languages; and it has been suggested (Evans & Levinson 2009) that language diversity far exceeds any universal core. (For normativist discussions of the analogy to language, see, e.g., Cohen 1981; Stein 1996; for relativist discussions, see Elqayam 2011; Stich 1990.)

For a while, it seemed that normative assessment might provide a useful tool for dual-process research. In his earlier individual differences studies, Stanovich (1999; Stanovich & West 2000a) found many examples of reasoning and decision tasks where participants of higher cognitive ability were more able to find correct answers by the standard norm applied to each task. This was attributed to their superior ability for System 2 reasoning. But as **Thompson** points out, there are multiple possible causes of correlations between cognitive capacity and normative performance. These include a disposition for rational thinking, and the presence of relevant education and training (Stanovich 2009b; 2010b). In recent research, identification of exceptions to the correlation of ability with normative performance has increased and we now know that many decision biases are fully independent of cognitive ability (Stanovich & West 2008). Stanovich’s most recent theory of the mind (Stanovich 2010b) reflects

these multiple requirements for normative responding. Ironically, in the developmental context, Stanovich (in press) has recently shown how simplistic and erroneous are some recent attempts to test dual-process theories of cognitive development, precisely because they assume that normativity reflects System 2 thinking. So why has **Stanovich** presented a staunch defence of normativism here? Because, it seems, his goal remains that of accounting for good and bad reasoning: he just thinks it is a complicated business. We ourselves have no mixed motives in critiquing the highly normative traditions of developmental psychology from Piaget onwards. This tradition is endorsed by **Giroto** who remarks that: “If one does not use normative standards, how could one compare the answers produced by children of different ages?” The answer is: by observing what the children actually do, as we discuss below.

A key feature of dual-process (system, mind) theories is that humans have a uniquely developed system for general purpose thinking and reasoning (objections of **Brase & Shanteau** noted). This is enabled by the acquisition of many beliefs, heuristics, and procedures in the new mind (or System 2) which are dependent upon education and cultural environment. So, while the object of cognitive psychology is to uncover the biological mechanism of the mind, we must always be aware of the culturally determined nature of much of its contents and sometimes even its processes. Hence, the importance of cross-cultural research in determining cognitive universals and the danger of approaching research with *a priori* normative framework which inevitably reflects the culture in which the researchers themselves were raised. Thinking about rationality, for example, has very different traditions in Eastern and Western cultures (Nisbett et al. 2001). We absolutely agree with **Kim & Park** that it is most important to study individual and cultural differences in a descriptive psychology of thinking, and to avoid application of universal norms in the process. However, we cannot accept at face value their proposal that cultural variations preclude any useful generalisation of cognitive architecture. On the contrary, we argue that cognitive variability (cultural, individual, historical) is essential to understanding universal cognitive architecture: without it, the picture portrayed by theorists may be too limited (Henrich et al. 2010; Norenzayan & Heine 2005). Consider, for example, Buchtel and Norenzayan’s (2009) proposal that contextualising, holistic thinking might be acquired in Eastern cultures as the preferred form of *Type 2* thinking. If this is right, it contributes a powerful evidence against the received view that *Type 2* processes are essentially abstract and decontextualised (Evans, in press).

From a descriptivist viewpoint, influence of age, intelligence, and culture are observable simply because people *do different things*. If this were not the case, the normative assessments advocated by **Giroto** and many others could not differ. The observable change is seen in *what* people do: the *correctness* of the responding is, as always, an externally imposed interpretation. Such differences in responding may be quantitative or qualitative. For example, as children grow older they are able to recall longer lists of words or digits, which is a quantitative change. They may also quite suddenly start to take account of another person’s perspective in responding to

a theory of mind task: a qualitative change. The same applies to individual differences in adults and between cultures. In both cases, we can observe what it is that people do differently without normatively assessing it. It is no coincidence that the descriptivist approach of cross-cultural research described by **Kim & Park** leads often to discovery of qualitative differences.

We agree with **Thompson** that a focus on normativity may lead us into reducing complex processes to binary measurements, and this danger is particularly strong when investigating individual and cultural differences. As a result, qualitative individual differences may be overlooked. Suppose we empirically observe four ways of responding, only one of which conforms with the norm we have in mind? Why is it any less important to understand how each of the other three arise? As an example, in a study of Bayesian reasoning, Evans et al. (2000) reported a number of qualitatively different responses; some participants gave the base rate, some the false positive rate, some, say, one minus the false positive rate, and so on. These variations have gone almost totally unnoticed in the large literature on the base rate fallacy, because authors time and again focus on normative responding or its absence. The right/wrong framing has dominated perceptions of findings in this literature, as in so many others.

R6.2. Normative relativism and why the bridge solution requires it

Normative relativism means that normative evaluation should be relative to specific agents with their specific goals and within a specific cultural and historical context (e.g., Elqayam 2011; Stich 1990). Soft normativism (see sects. R2.1, R7) often leads to such normative relativism, as some commentators have readily acknowledged (**Achourioti et al.** [and see Stenning & van Lambalgen 2008]; **Proust; Weinberg**). In particular, normative relativism is inevitable where soft normativism attempts to bridge *is* and *ought* (sect. R2.2), drawing on constructs such as reflective equilibrium and constitutive norms.

The evaluative and deontic functions reflected by bridge constructs draw heavily on cultural and historical context (Elqayam 2011). Here are a few examples. First, as Stich (1990) famously argued (and **Buckwalter & Stich** remind us), reflective equilibrium can result in a radical, anything-goes relativism. Moreover, Searle’s (1964) constitutive norms (sect. R3.1) are social institutions, and depend on cultural context. But the classic example is the Kantian “ought implies can,” perhaps the prototypical bridge construct, with its corollaries, “cannot implies not-ought” (bounded rationality), and “can implies ought” (principle of charity). Not for nothing did Nickerson (2008) classify bounded rationality under “relativity.” “Can” and “cannot” are fuzzy concepts. Ultimately, they rely on biology and cognitive architecture, but how cognitive architecture is instantiated in individual agents depends a great deal on cultural and historical context (Henrich et al. 2010; Norenzayan & Heine 2005). Where “cannot” implies “not-ought,” the descriptive facts of cognitive variability should force the normativist’s hand to adopt a position of normative relativism. Where theorists wish to bridge the “is” with the “ought,” yet avoid the relativist implications, they allow an inconsistency into their system.

R7. Descriptivism versus normativism in conduct of empirical research on thinking

An important issue for us, as well as many of our commentators, is whether normativism helps or hinders the empirical investigation of human thinking. We argued strongly for the descriptivist approach in our target article, also claiming that normativism has biased what we study, how we study it, and the interpretations of our findings. Broad or strong support for the descriptive approach can be found in several commentaries (e.g., **Bonnefon; Sternberg; Thompson**), while a small number of others present robust defences of the use of normativism in psychological research (**Giroto; Hahn; Pfeifer; Stanovich**). However, a number of other commentators take the in-between stance that we have described as “soft normativism.” This is the view that both normativism and descriptivism have their place in empirical research, and that while the former can lead to biases it can also be valuable (**Douven; Hroic; McNair & Feeney; Nickerson; Quintelier & Fessler; Schurz; Stuppel & Ball; Sun & Wang; Waldmann**).

Closer examination of these middle ground positions suggests that they are based on two distinct rationales. Some authors fully endorse our psychological and descriptive objectives, but think that normativism has nevertheless had a useful heuristic value in the conduct of research; others see it as an essential tool for theory construction. Evidence for heuristic value is that the paradigms inspired by normativism have led to a number of important psychological findings (e.g., **Stuppel & Ball; McNair & Feeney**, perhaps **Waldmann**). It is hard to dispute this: Jonathan Evans has spent much of his career working with the deduction paradigm and would prefer not to think he was wasting his time. In fact, Evans (2010a) shows how research on deductive reasoning, intended to investigate reflective Type 2 thinking, actually yielded a wealth of information about Type 1, intuitive thought. This happened because enough psychologists were open to what the data were telling them (although many also were not). As soon as the extent of logical reasoning biases, content and context effects, et cetera, became apparent, psychologists started theorising about them and designing new experiments to test explanations of what were initially seen as cognitive biases. So much systematic study did lead to important psychological advances in the end, as well as the eventual abandonment by many of the original paradigm. So in a sense, it is true in this case that normativism inspired important empirical studies. But we very much doubt that it was the most efficient way of doing so. (We also agree with **Sternberg** that a number of important kinds of reasoning – of everyday relevance – were neglected in the process.) In the same way, we agree that much psychological value came out of the heuristics and biases research programme inspired by Tversky and Kahneman (**McNair & Feeney**) despite the normativist framing of much of this work.

We assumed, in our target article, the research objectives of cognitive psychology. Of course, psychology does not have a monopoly on the study of human behaviour which features in many other disciplines, including economics, sociology, geography, history, and even the study of literature. So we are not surprised, and nor do we object to the fact, that some philosophers wish to study

human reasoning in order to assess the relevance and empirical grounding of normative theories (**Douven; Quintelier & Fessler; Schurz**; also **Hrotic**, an anthropologist). We can assure Quintelier & Fessler that we have no desire to outlaw experimental philosophy. What does surprise us is the enthusiastic support for empirical assessment of normative theories from some of the psychologists, such as **Nickerson**, who insists that we need to know not only how we reason but how we *should* reason. For our part, while we can see that empirical studies of thinking are of interest to those who earn their living debating the value of normative theories, such as logicians and economists, we do not see their objectives as being directly *psychological*. For their part, they still have to solve the problem of is-ought inference.

A second argument for the value of normativism in psychological research is that a close comparison of the discrepancy between normative theory and actual behaviour is of direct value in constructing psychological theories (**Hahn; Nickerson; Oaksford & Chater; Pfeifer; Waldmann**). A general theme is that of iterative refinement of normative theories in the light of psychological evidence, with an increased convergence between normative and descriptive accounts (see also **Buckwalter & Stich** on reflective equilibrium). Thus, Pfeifer talks about how the psychology of reasoning has moved on from bivalent to probability logics, which are now seen as a much closer approximation of actual human reasoning. Hahn describes the iteration involved in ideal observer analysis and Oaksford & Chater lay out in detail the research strategy for their rational analysis programme. We do wonder, however, whether some of these authors are really talking of competence or computational rather than normative accounts (see sect. R2). It seems to us that the force of a normative theory lies in its *a priori*, evaluative oughtness. We cannot imagine, for example, a strict religious sect with a ban on premarital sex, revising its dogma in the light of observation of regular deviations in the behaviour of its young members.

If the process of iteration described by **Oaksford & Chater** really means comparing formal computational accounts to observed behaviour, and iteratively refining them, then we have no problem at all with rational analysis. It is a method strongly founded in the directive ought, but need not be evaluative. We just wonder why their earlier writing on this topic has needed to emphasise good and bad reasoning in the way it has. The same probably applies to some other critical commentators, such as **Pfeifer** and **Sun & Wang**, who talk of the value of developing alternative normative theories on the basis of their favoured empirical research paradigm. If that is really what they are doing, then they are committing is-ought inferences which are hard to defend. But if, instead, they are developing improved computational accounts of the processes, drawing on available formal theories as they do so, then we do not have a problem with this, as we made quite clear in our target article.

R8. Conclusions

We think it important that researchers in any field of science raise their heads above the parapets of their paradigms from time to time, and reflect more broadly on what

they are doing. Our purpose was to stimulate such head-raising for those engaged in the psychological study of reasoning and decision making: fields which, while purportedly signing up to the methods and objectives of the much wider field of cognitive psychology, also have a history of application of normative theory that sets them apart. We are grateful to the number of colleagues who have taken the time to comment and note the considerable diversity in the views expressed. Whether or not we succeed in moving the field towards a more descriptivist approach, we hope at least that we have raised consciousness of the important issue of normativism in these fields, and that researchers will think a little more carefully and clearly about what they are doing, following this lively and informative debate.

ACKNOWLEDGMENT

We thank David Over for a critical reading of an earlier draft of this response paper.

References

[The letters “a” and “r” before author’s initials stand for target article and response references, respectively.]

- Adams, E. W. (1998) *A primer of probability logic*. CSLI Press. [GS]
- Anderson, J. R. (1978) Arguments concerning representations for mental imagery. *Psychological Review* 85:249–77. [r]StBTE]
- Anderson, J. R. (1990) *The adaptive character of thought*. Erlbaum. [aSE, MO]
- Anderson, J. R. (1991) Is human cognition adaptive? *Behavioral and Brain Sciences* 14:471–517. [aSE, MO]
- Anderson, N. H. & Shanteau, J. (1977) Weak inference with linear models. *Psychological Bulletin* 84:1147–70. [GLB]
- Austin, J. L. (1961) Other minds. In: *Philosophical papers*, ed. J. O. Urmson & G. J. Warnock, pp. 44–84. Oxford University Press. [aSE]
- Ayton, P. & Fischer, I. (2004) The hot hand fallacy and the gambler’s fallacy: Two faces of subjective randomness? *Memory and Cognition* 32(8):1369–78. [YS]
- Baker, C. L., Saxe, R. & Tenenbaum, J. B. (2009) Action understanding as inverse planning. *Cognition* 113:329–49. [J-FB]
- Ball, L. J., Phillips, P., Wade, C. N. & Quayle, J. D. (2006) Effects of belief and logic on syllogistic reasoning: Eye-movement evidence for selective processing models. *Experimental Psychology* 53:77–86. [EJNS]
- Barbey, A. K. & Sloman, S. A. (2007) Base rate respect: From ecological rationality to dual processes. *Behavioral and Brain Sciences* 30:241–97. [aSE, SMcN]
- Bar-Eli, M., Avugos, S. & Raab, M. (2006) Twenty years of “hot hand” research: Review and critique. *Psychology of Sport and Exercise* 7(6):525–53. [YS]
- Bargh, J. A., ed. (2006) *Social psychology and the unconscious*. Psychology Press. [r]StBTE]
- Baron, J. (2008) *Thinking and deciding*, 4th edition. Cambridge University Press. [aSE]
- Barrett, H. C. & Kurzban, R. (2006) Modularity in cognition: Framing the debate. *Psychological Review* 113:628–47. [GLB]
- Barrouillet, P., Markovits, H. & Quinn, S. (2001) Developmental and content effects in reasoning with causal conditionals. *Journal of Experimental Child Psychology* 81:235–48. [aSE]
- Benferhat, S., Bonnefon, J. F. & Da Silva Neves, R. (2005) An overview of probabilistic handling of default reasoning, with experimental studies. *Synthese* 146:53–70. [J-FB]
- Bierce, A. (1971) *The Enlarged Devil’s dictionary*. Penguin. [DS]
- Binmore, K. (2009) *Rational decisions*. Princeton University Press. [NG]
- Bishop, M. A. & Trout, J. D. (2005) *Epistemology and the psychology of human judgment*. Oxford University Press. [GS]
- Blake, A., Bulthoff, H. H. & Sheinberg, D. (1996) Shape from texture: Ideal observers and human psychophysics. In: *Perception as Bayesian inference*, ed. D. Knill & W. Richards, pp. 287–321. Cambridge University Press. [MO]
- Bombardieri, M. (2005) Summers’ remarks on women draw fire. Retrieved from: http://www.boston.com/news/education/higher/articles/2005/01/17/summers_remarks_on_women_draw_fire/ [SH]

- Bonnefon, J. F. (2009) A theory of utility conditionals: Paralogical reasoning from decision-theoretic leakage. *Psychological Review* 116:888–907. [J-FB]
- Bonnefon, J. F., Dubois, D., Fargier, H. & Leblois, S. (2008) Qualitative heuristics for balancing the pros and cons. *Theory and Decision* 65:71–95. [J-FB]
- Bornstein, B. H. & Emler, A. C. (2001) Rationality in medical decision making: A review of the literature on doctors' decision making biases. *Journal of Evaluation in Clinical Practice* 7:97–107. [SMcN]
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S. & Cohen, J. D. (2001) Conflict monitoring and cognitive control. *Psychological Review* 108(3):624–52. [JP]
- Braine, M. D. S. & O'Brien, D. P. (1998) The theory of mental-propositional logic: Description and illustration. In: *Mental logic*, ed. M. D. S. Braine & D. P. O'Brien, pp. 79–89. Erlbaum. [aSE]
- Broome, J. (1999) Normative requirements. *Ratio* 12:398–419. [JP]
- Buchtel, C. & Norenzayan, A. (2009) Thinking across cultures: Implications for dual processes. In: *In two minds: Dual processes and beyond*, ed. J. St. B. T. Evans & K. Frankish, pp. 217–38. Oxford University Press. [r]StBTE]
- Burns, B. D. (2004) Heuristics as beliefs and as behaviors: The adaptiveness of the “hot hand”. *Cognitive Psychology* 48:295–331. [YS]
- Busemeyer, J. R., Pothos, E. M., Franco, R. & Trueblood, J. (2011) A quantum theoretical explanation for probability judgment errors. *Psychological Review* 118:193–218. [EMP]
- Busemeyer, J. R., Wang, Z. & Townsend, J. T. (2006) Quantum dynamics of human decision-making. *Journal of Mathematical Psychology* 50:220–41. [EMP]
- Byrne, R. M. J. (2005) *The rational imagination*. MIT Press. [r]StBTE]
- Byrne, R. M. J. & Handley, S. J. (1997) Reasoning strategies for suppositional deductions. *Cognition* 62:49. [aSE]
- Byrne, R. M. J., Handley, S. J. & Johnson-Laird, P. N. (1995) Reasoning from suppositions. *The Quarterly Journal of Experimental Psychology* 48A:915–44. [aSE]
- Carruthers, P. (2006) The case for massively modular models of mind. In: *Contemporary debates in cognitive science*, ed. R. J. Stainton, pp. 3–21. Blackwell. [GLB]
- Carruthers, P. (2009) How we know our own minds: The relationship between mindreading and metacognition. *Behavioral and Brain Sciences* 32(3):121–38. [r]StBTE]
- Chater, N. & Oaksford, M. (1999) The probability heuristics model of syllogistic reasoning. *Cognitive Psychology* 38:191–258. [MO]
- Chater, N. & Oaksford, M. (in press) Normative systems: Logic, probability, and rational choice. In: *The Oxford handbook of thinking and reasoning*, ed. K. Holyoak & R. Morrison. Oxford University Press. [MO]
- Chater, N., Tenenbaum, J. & Yuille, A., eds. (2006) *Probabilistic models of cognition*. Special Issue: *Trends in Cognitive Sciences* 10. [MO]
- Chomsky, N. (1995) *The minimalist program*. MIT Press. [NG]
- Chua, H. F., Boland, J. E. & Nisbett, R. E. (2005a) Cultural variation in eye movements during scene perception. *Proceedings of the National Academy of Sciences USA* 102(35):12629–33. [KK]
- Chua, H. F., Leu, J. & Nisbett, R. E. (2005b) Culture and diverging views of social events. *Personality and Social Psychology Bulletin* 31(7):925–34. [KK]
- Cohen, L. J. (1981) Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences* 4:317–70. [WB, aSE, r]StBTE, KES]
- 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* 12:251–74. [aSE]
- Coletti, G. & Scozzafava, R. (2002) *Probabilistic logic in a coherent setting*. Kluwer Academic. [NP]
- Cosmides, L. & Tooby, J. (1996) Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition* 58:1–73. [aSE]
- Davidson, D. (1973) Radical interpretation. *Dialectica* 27:313–28. [r]StBTE, NG]
- Dehaene, S. & Brannon, E. M. (2010) Space, time, and number: A Kantian research program. *Trends in Cognitive Sciences* 14(12):517–19. [YS]
- De Neys, W. (2006) Dual processing in reasoning: Two systems but one reasoner. *Psychological Science* 17(5):428–33. [EJNS]
- De Neys, W. & Glumicic, T. (2008) Conflict monitoring in dual process theories of thinking. *Cognition* 106:1248–99. [aSE, SMcN]
- De Neys, W., Vartanian, O. & Goel, V. (2008) Smarter than we think. *Psychological Science* 19(5):48389. [EJNS]
- Dennett, D. C. (1987) *The intentional stance*. MIT Press. [aSE]
- Dennett, D. C. (2008) Fun and games in fantasyland. *Mind and Language* 23:25–31. [aSE]
- De Saussure, F. (1916/1966) *Course in general linguistics*. (Original publication, 1916). McGraw-Hill. [aSE]
- Descartes, R. (1985) *The philosophical writings of Descartes, vol. I*, trans. J. Cottingham, R. Stoothoff, & D. Murdoch. Cambridge University Press. [CDN]
- DeScioli, P., Christner, J. & Kurzban, R. (in press) The omission strategy. *Psychological Science*. [NG]
- de Soto, D. (1539–1540) *Aeditio Secunda Summularum*. Salamanca. [CDN]
- Dijksterhuis, A., Bos, M. W., Nordgren, L. F., & van Baaren, R. B. (2006) On making the right choice: The deliberation-without-attention effect. *Science*, 311:1005–1007. [aSE]
- Douven, I. & Dietz, R. (2011) A puzzle about Stalnaker's hypothesis. *Topoi* 30:31–37. [ID]
- Douven, I. & Romeijn, J. W. (in press) A new resolution of the Judy Benjamin problem. *Mind*. [ID]
- Dretske, F. (2000) Norms, history, and the constitution of the mental. In: *Perception, knowledge and belief*. pp. 242–58. Cambridge University Press. [JP]
- Edgington, D. (1995) On conditionals. *Mind* 104:235–329. [aSE, MO]
- Edgington, D. (2003) What if? Questions about conditionals. *Mind and Language* 18:380–401. [aSE]
- Edgington, D. (2008) Conditionals. In: *Stanford encyclopedia of philosophy*, ed. E. N. Zalta. Stanford University. (Online publication. Available at: <http://plato.stanford.edu/archives/win2008/entries/conditionals/>). [aSE]
- Elqayam, S. (2003) Norm, error and the structure of rationality: The case study of the knight-knave paradigm. *Semiotica* 147:265–89. [aSE]
- Elqayam, S. (2006) The collapse illusion effect: A pragmatic–semantic illusion of truth and paradox. *Thinking and Reasoning* 12:180. [aSE]
- Elqayam, S. (2011) Grounded rationality: A relativist framework for normative rationality. In: *The science of reason: A Festschrift in honour of Jonathan St. B. T. Evans*, ed. K. I. Manktelow, D. E. Over & S. Elqayam, pp. 397–420. Psychology Press. [aSE, r]StBTE]
- Elster, J. (1989) *Nuts and bolts for the social sciences*. Cambridge University Press. [NG]
- Epstein, S. (1994) Integration of the cognitive and psychodynamic unconscious. *American Psychologist* 49:709–24. [r]StBTE]
- Eskine, K. J., Kacirik, N. A. & Prinz, J. J. (2011) A bad taste in the mouth: Gustatory disgust influences moral judgment. *Psychological Science* 22:295–99. [MRW]
- Evans, J. St. B. T. (1972) Interpretation and matching bias in a reasoning task. *Quarterly Journal of Experimental Psychology* 24:193–99. [aSE]
- Evans, J. St. B. T. (1982) *The psychology of deductive reasoning*. Routledge. [aSE, NP]
- Evans, J. St. B. T. (1989) *Bias in human reasoning: Causes and consequences*. Erlbaum. [aSE]
- Evans, J. St. B. T. (1993) Bias and rationality. In: *Rationality: Psychological and philosophical perspectives*, ed. K. I. Manktelow & D. E. Over, pp. 6–30. Routledge. [aSE]
- Evans, J. St. B. T. (2000) Thinking and believing. In: *Mental models in reasoning*, ed. J. Garcia-Madruga, N. Carriedo & M. J. González-Labra, pp. 41–56. UNED (Universidad de Nacional de Education a Distancia). [EJNS]
- Evans, J. St. B. T. (2002) Logic and human reasoning: An assessment of the deduction paradigm. *Psychological Bulletin* 128:978–96. [aSE, r]StBTE]
- Evans, J. St. B. T. (2003) In two minds: Dual process accounts of reasoning. *Trends in Cognitive Sciences* 7:454–59. [aSE]
- Evans, J. St. B. T. (2006) The heuristic–analytic theory of reasoning: Extension and evaluation. *Psychonomic Bulletin and Review* 13:378–95. [aSE]
- Evans, J. St. B. T. (2007) *Hypothetical thinking: Dual processes in reasoning and judgement*. Psychology Press. [aSE, r]StBTE]
- Evans, J. St. B. T. (2008) Dual-processing accounts of reasoning, judgment and social cognition. *Annual Review of Psychology* 59:255–78. [aSE, r]StBTE]
- Evans, J. St. B. T. (2010a) Intuition and reasoning: A dual-process perspective. *Psychological Inquiry* 21:313–26. [r]StBTE]
- Evans, J. St. B. T. (2010b) *Thinking twice: Two minds in one brain*. Oxford University Press. [aSE, r]StBTE]
- Evans, J. St. B. T. (in press a) Dual-process theories of reasoning: Facts and fallacies. In: *The Oxford handbook of thinking and reasoning*, ed. K. Holyoak & R. G. Morrison. Oxford University Press. [r]StBTE]
- Evans, J. St. B. T. (in press b) Reasoning. In: *The Oxford handbook of cognitive psychology*, ed. D. Reisberg. Oxford University Press. [aSE]
- 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. [aSE, JP]
- Evans, J. St. B. T. & Coventry, K. (2006) A dual-process approach to behavioral addiction: The case of gambling. In: *Handbook of implicit cognition and addiction*, ed. R. W. Wiers, A. W. Stacy, R. pp. 29–43. Sage. [aSE]
- Evans, J. St. B. T. & Frankish, K., eds. (2009) *In two minds: Dual processes and beyond*. Oxford University Press. [aSE]
- Evans, J. St. B. T. & Handley, S. J. (1999) The role of negation in conditional inference. *Quarterly Journal of Experimental Psychology* 52A:739–69. [aSE]
- Evans, J. St. B. T., Handley, S., Neilens, H., Bacon, A. M. & Over, D. E. (2010) The influence of cognitive ability and instructional set on causal conditional inference. *Quarterly Journal of Experimental Psychology* 63:892–909. [r]StBTE]

- Evans, J. St. B. T., Handley, S. J. & Over, D. E. (2003) Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory and Cognition* 29(2):321–55. [aSE, GS]
- Evans, J. St. B. T., Handley, S. J., Perham, N., Over, D. E. & Thompson, V. A. (2000) Frequency versus probability formats in statistical word problems. *Cognition* 77:197–213. [aSE, rJStBTE, SMcN]
- Evans, J. St. B. T., Neilens, H., Handley, S. & Over, D. E. (2008) When can we say “if”? *Cognition* 108:100–16. [rJStBTE]
- Evans, J. St. B. T., Newstead, S. E. & Byrne, R. M. J. (1993) *Human reasoning: The psychology of deduction*. Erlbaum. [aSE, EJNS]
- Evans, J. St. B. T. & Over, D. E. (1996) *Rationality and reasoning*. Psychology Press. [aSE, rJStBTE]
- Evans, J. St. B. T. & Over, D. E. (2004) *If*. Oxford University Press. [aSE]
- 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* 112:1040–52. [aSE]
- Evans, J. St. B. T. & Wason, P. C. (1976) Rationalisation in a reasoning task. *British Journal of Psychology* 63:205–12. [rJStBTE]
- Evans, N. & Levinson, S. C. (2009) The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences* 32(5):429–48. [rJStBTE]
- Falk, R., Falk, R. & Ayton, P. (2009) Subjective patterns of randomness and choice: Some consequences of collective responses. *Journal of Experimental Psychology: Human Perception and Performance* 35(1):203–24. [YS]
- Falk, R. & Konold, C. (1997) Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review* 104(2):301–18. [YS]
- Feeney, A. (2007) Individual differences, dual processes, and induction. In: *Inductive reasoning*, ed. A. Feeney & E. Heit, pp. 302–27. Cambridge University Press. [aSE]
- Feeney, A. & Heit, E., eds. (2007) *Inductive reasoning: Experimental, developmental, and computational approaches*. Cambridge University Press. [rJStBTE, SMcN]
- Fernbach, P. M., Darlow, A. & Sloman, S. A. (2011) Asymmetries in predictive and diagnostic reasoning. *Journal of Experimental Psychology: General* 140:168–85. [SMcN]
- 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. [aSE]
- Fitzgerald, F. S. (1936) The Crack-Up. *Esquire Magazine*, February 1936. Available at: <http://www.esquire.com/features/the-crack-up>. [GLB]
- Fodor, J. A. (1983) *The modularity of mind: An essay in faculty psychology*. MIT Press. [GLB]
- Fodor, J. A. (2008) Against Darwinism. *Mind and Language* 23:1–24. [aSE]
- Foot, P. (1967) The problem of abortion and the doctrine of double effect. *Oxford Review* 5:5–15. [NG]
- Ford, M. (2004) System LS: A three-tiered nonmonotonic reasoning system. *Computational Intelligence* 20(1):89–108. [NP]
- Ford, M. (2005) On using human nonmonotonic reasoning to inform artificial systems. *Psychologica Belgica* 45:57–70. [NP]
- Fox, J. (2003) Probability, logic and the cognitive foundations of rational belief. *Journal of Applied Logic* 1:197–224. [RSN]
- Fox, J., Beveridge, D. & Glasspool, D. (2003) Understanding intelligent agents: Analysis and synthesis. *AI Communications* 16:139–52. [RSN]
- Fox, J. & Parsons, S. (1998) Arguing about beliefs and actions. In: *Applications of uncertain formalisms*, ed. A. Hunter & S. Parsons, pp. 266–302. Springer. [RSN]
- Frankena, W. (1939) The naturalistic fallacy. *Mind* 48:464–77. [aSE, rJStBTE]
- Frankena, W. (1963) *Ethics*. Prentice-Hall. [GS]
- Fugard, A. J. B., Pfeifer, N. & Mayerhofer, B. (2011a) Probabilistic theories of reasoning need pragmatics too: Modulating relevance in uncertain conditionals. *Journal of Pragmatics* 43:2034–42. [TA]
- Fugard, A. J. B., Pfeifer, N., Mayerhofer, B. & Kleiter, G. D. (2011b) How people interpret conditionals: Shifts towards the conditional event. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 37(3):635–48. [TA, NP]
- Gal, Y. & Pfeffer, A. (2008) Networks of influence diagrams: A formalism for representing agents’ beliefs and decision-making processes. *Journal of Artificial Intelligence Research* 33:109–47. [J-FB]
- Gardner, M. (1988) *Time travel and other mathematical bewilderments*. Freeman. [YS]
- Gardner, W. L., Gabriel, S. & Lee, A. Y. (1999) “I” value freedom, but “we” value relationships: Self-construal priming mirrors cultural differences in judgment. *Psychological Science* 10(4):321–26. [KK]
- Geisler, W. S. (1987) Ideal-observer analysis of visual discrimination. In: *Frontiers of visual science*. National Academy Press. [UH]
- Gigerenzer, G. (1991) How to make cognitive illusions disappear: Beyond “heuristics and biases.” In: *European review of social psychology*, ed. W. Stroebe & M. Hewstone, pp. 83–115. Wiley. [aSE]
- Gigerenzer, G. (2007) *Cut feelings: The intelligence of the unconscious*. Viking Press. [aSE, SH]
- Gigerenzer, G., Hell, W. & Blank, H. (1988) Presentation and content: The use of base rates as a continuous variable. *Journal of Experimental Psychology: Human Perception and Performance* 14:513–25. [SMcN]
- Gigerenzer, G. & Hoffrage, U. (1995) How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review* 102:684–704. [aSE]
- Gigerenzer, G. & Selten, R. (2001) *Bounded rationality: The adaptive toolbox*. MIT Press. [aSE]
- Gigerenzer, G. & Todd, P. M. (1999) Fast and frugal heuristics: The adaptive toolbox. In: *Simple heuristics that make us smart*, ed. G. Gigerenzer, P. M. Todd & the ABC Research Group, pp. 3–34. Oxford University Press. [aSE]
- Gigerenzer, G., Todd, P. M. & the ABC Research Group, eds. (1999) *Simple heuristics that make us smart*. Oxford University Press. [aSE, SH, EMP, GS]
- Gilovich, T., Griffin, D. & Kahneman, D. (2002) *Heuristics and biases: The psychology of intuitive judgement*. Cambridge University Press. [aSE]
- Gilovich, T., Vallone, R. & Tversky, A. (1985) The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology* 17:295–314. [YS]
- Giroto, V. & Gonzalez, M. (2001) Solving probabilistic and statistical problems: A matter of information structure and question form. *Cognition* 78(3):247–76. [VG, SMcN]
- Giroto, V. & Gonzalez, M. (2008) Children’s understanding of posterior probability. *Cognition* 106:325–44. [VG]
- Goel, V. & Dolan, R. J. (2003) Explaining modulation of reasoning by belief. *Cognition* 87:B11–B22. [JP, EJNS]
- Goldman, A. (1979) What is justified belief? In: *Justification and knowledge*, ed. G. Pappas, pp. 1–23. Reidel. [JMW]
- Goldman, A. (1986) *Epistemology and cognition*. Harvard University Press. [TF, GS]
- Goldman, A. (1999) *Knowledge in a social world*. Oxford University Press. [GS]
- Goldstein, D. G. & Gigerenzer, G. (1999) The recognition heuristic: How ignorance makes us smart. In: *Simple heuristics that make us smart*, by G. Gigerenzer, P. M. Todd, & The ABC Research Group (1999), pp. 37–58. Oxford University Press. [VAT]
- Goldszmidt, M. & Pearl, J. (1996) Qualitative probabilities for default reasoning, belief revision and causal modeling. *Artificial Intelligence* 84:57–112. [GS]
- Goodman, N. (1965) *Fact, fiction, and forecast*. Bobbs-Merrill. [WB, rJStBTE]
- Gottwald, S. (2001) *A treatise on many-valued logics. (Studies in Logic and Computation, vol. 9)*. Research Studies Press. [aSE]
- 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 Society of London, Series B* 205:581–98. [aSE]
- Govier, T. (1987) Problems in argument analysis and evaluation. Foris. [UH]
- Greene, J. D. (2007) Why are VMPFC patients more utilitarian? A dual-process theory of moral judgment explains. *Trends in Cognitive Sciences* 11:322–23. [NG]
- Hagen, E. H. (2005) Controversial issues in evolutionary psychology. In: *The handbook of evolutionary psychology*, ed. D. M. Buss, pp. 145–73. John Wiley. [GLB]
- Hagmayer, Y. & Sloman, S. A. (2009) Decision makers conceive of their choices as interventions. *Journal of Experimental Psychology: General* 138:22–38. [J-FB]
- Hahn, U. (2009) Explaining more by drawing on less. Commentary on Oaksford & Chater. *Behavioral and Brain Sciences* 32:90–91. [UH]
- Hahn, U. & Oaksford, M. (2007) The rationality of informal argumentation: A Bayesian approach to reasoning fallacies. *Psychological Review* 114:704–32. [MO]
- Hahn, U. & Warren, P. A. (2009) Perceptions of randomness: Why three heads are better than four. *Psychological Review* 116(2):454–61. [aSE, YS]
- Hahn, U. & Warren, P. A. (2010) Why three heads are a better bet than four: A reply to Sun, Tweney, and Wang (2010) *Psychological Review* 117(2):706–11. [YS]
- Hanson, N. R. (1958) *Patterns of discovery*. Cambridge University Press. [TA]
- Hare, R. M. (1969) Descriptivism. In *The is-ought question: A collection of papers on the central problem in moral philosophy*, ed. W. D. Hudson, pp. 240–58. Macmillan. [aSE]
- Hare, R. M. (1993) *Essays in ethical theory*. Oxford University Press. [aSE]
- Harman, G. (1975) Moral relativism defended. *The Philosophical Review* 84(1):3–22. [KJPQ]
- Harrington, S. P. M. (1993) Bones and bureaucrats: New York’s great cemetery imbroglio. *Archaeology* 46(2):28–38. [SH]
- Harris, R. (1980) *The language-makers*. Duckworth. [aSE]
- Harris, R. (1981) *The language myth*. Duckworth. [aSE]
- Hauser, M. D. (2007) *Moral minds: How nature designed our universal sense of right and wrong*. Little, Brown. [NG]
- Henle, M. (1962) On the relation between logic and thinking. *Psychological Review* 69(4):366–78. [VAT]

- Henrich, J., Heine, S. J. & Norenzayan, A. (2010) The weirdest people in the world? *Behavioral and Brain Sciences* 33(2):61–83. [r]StBTE]
- Hergenhahn, B. R. (2009) *An introduction to the history of psychology*, 6th edition. Wadsworth. [CDN]
- Hilton, D. J. (1995) The social context of reasoning: Conversational inference and rational judgment. *Psychological Bulletin* 118(2):248–71. [UH]
- Hogarth, R. M. & Einhorn, H. J. (1992) Order effects in belief updating: The belief adjustment model. *Cognitive Psychology* 24:1–55. [r]StBTE]
- Houdé, O. & Tzourio-Mazoyer, N. (2003) Neural foundations of logical and mathematical cognition. *Nature Reviews Neuroscience* 4:507–14. [JP]
- Howes, A., Lewis, R. L. & Vera, A. (2009) Rational adaptation under task and processing constraints: Implications for testing theories of cognition and action. *Psychological Review* 116(4):717–51. [UH]
- Howson, C. & Urbach, P. (1993) *Scientific reasoning*, 2nd edition. Open Court. [aSE]
- Hsee, C. K. & Weber, E. U. (1999) Cross-national differences in risk preference and lay predictions. *Journal of Behavioral Decision Making* 12(2):165–79. [KK]
- Hudson, W. D., ed. (1969) *The is-ought question: A collection of papers on the central problem in moral philosophy*. Macmillan. [aSE]
- Hume, D. (1739–1740) *A treatise of human nature. Being an attempt to introduce the experimental method of reasoning into moral subjects*. Available at: <http://www.gutenberg.org/ebooks/4705>. [KJPQ]
- Hume, D. (1739-1740/2000) *A treatise on human nature*. (Original work published in 1739–1740). Clarendon Press. [aSE]
- Inhelder, B. & Piaget, J. (1958) *The growth of logical thinking*. Basic Books. [aSE]
- Jackendoff, R. (2002) *Foundations of language: Brain, meaning, grammar, evolution*. Oxford University Press. [aSE]
- Jepson, C., Krantz, D. & Nisbett, R. (1983) Inductive reasoning: Competence or skill? *Behavioral and Brain Sciences* 6:494–501. [KES]
- Johnson-Laird, P. N. & Byrne, R. M. J. (1991) *Deduction*. Erlbaum. [aSE]
- Johnson-Laird, P. N. & Byrne, R. M. J. (2002) Conditionals: A theory of meaning, pragmatics and inference. *Psychological Review* 109:646–78. [aSE]
- Jones, M. & Love, B. C. (2011) Bayesian Fundamentalism or Enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences* 34(4):169–231. [r]StBTE, SMcN]
- Joyce, J. (1998) A nonpragmatic vindication of probabilism. *Philosophy of Science* 65(4):575–603. [ID]
- Joyce, J. M. (2004) Bayesianism. In: *The Oxford handbook of rationality*, ed. A. R. Mele & P. Rawling, pp. 132–55. Oxford University Press. [KES]
- Kahneman, D. & Frederick, S. (2002) Representativeness revisited: Attribute substitution in intuitive judgement. In: *Heuristics and biases: The psychology of intuitive judgement*, ed. T. Gilovich, D. Griffin & D. Kahneman, pp. 49–81. Cambridge University Press. [aSE]
- Kahneman, D. & Tversky, A. (1972) Subjective probability: A judgment of representativeness. *Cognitive Psychology* 3:430–54. [aSE]
- Kahneman, D. & Tversky, A. (1973) On the psychology of prediction. *Psychological Review* 80:237–51. [SMcN, VAT]
- Kahneman, D. & Tversky, A. (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47:263–91. [aSE]
- Kant, I. (1781/1787/1929) *Kritik der Reinen Vernunft*. Johann Friedrich Hartknoch. (First edition = A, 1781. Second edition = B, 1787.) English edition, 1929: *Critique of pure reason*, trans. N. Kemp Smith. Macmillan. [CDN]
- Kant, I. (1781/1998) *Critique of pure reason*, trans. P. Guyer & A. W. Wood. The Cambridge edition of the works of Immanuel Kant. Cambridge University Press. (Original German publication 1781). [TA]
- Kareev, Y. (1992) Not that bad after all: generation of random sequences. *Journal of Experimental Psychology: Human Perception and Performance* 18:1189–94. [YS]
- Kemp, C. & Tenenbaum, J. B. (2009) Structured statistical models of inductive reasoning. *Psychological Review* 116:20–58. [SMcN]
- Kim, K. & Markman, A. B. (2006) Differences in fear of isolation as an explanation of cultural differences: Evidence from memory and reasoning. *Journal of Experimental Social Psychology* 42(3):350–64. [KK]
- Kim, K., Grimm, L. R. & Markman, A. B. (2007) Self-construal and the processing of covariation information in causal reasoning. *Memory and Cognition* 35:1337–43. [KK]
- Klauer, K. C. (1999) On the normative justification for information gain in Wason's selection task. *Psychological Review* 106:215–22 [MO]
- Klauer, K. C., Musch, J. & Naumer, B. (2000). On belief bias in syllogistic reasoning. *Psychological Review* 107:852–84. [aSE]
- Klein, G. A. (1998) *Sources of power: How people make decisions*. MIT Press. [aSE, RSN]
- Koehler, D. J. & James, G. (2009) Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition* 113:123–27. [KES]
- Koehler, J. J. (1996) The base rate fallacy reconsidered: Descriptive, normative and methodological challenges. *Behavioral and Brain Sciences* 19:1–53. [aSE]
- Koriat, A. & Goldsmith, M. (1996) Monitoring and control processes in the strategic regulation of memory accuracy. *Psychological Review* 103(3):490–517. [JP]
- Komblith, H. (2003) *Knowledge and its place in nature*. Oxford University Press. [JMW]
- Korsgaard, C. (1998) *The sources of normativity*. Cambridge University Press. [CDN]
- Kraus, S., Lehmann, D. & Magidor, M. (1990) Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44:167–207. [NP]
- Krebs, J. R. & Davies, N., eds. (1996) *Behavioural ecology: An evolutionary approach*, 4th edition. Blackwell. [MO]
- Kreps, D. M. (1992) *A course in microeconomic theory*. Harvester Wheatsheaf. [MO]
- Kripke, S. (1975) Outline of a theory of truth. *Journal of Philosophy* 72:690–716. [aSE]
- Krueger, J. (2001) Null hypothesis significance testing: On the survival of a flawed method. *American Psychologist* 56:16–26. [GLB]
- Krynski, T. R. & Tenenbaum, J. B. (2007) The role of causality in judgement under uncertainty. *Journal of Experimental Psychology: General* 136:430–50. [SMcN]
- Kuhn, D. (1991) *The skills of argument*. Cambridge University Press. [DS]
- Kuhn, T. S. (1962) *The structure of scientific revolutions*. University of Chicago Press. [TA, NG, EMP]
- Kuhn, T. S. (1996) *The structure of scientific revolutions*, 3rd edition. University of Chicago Press. [SH]
- Leitgeb, H. (2004) *Inference on the low level*. Kluwer. [GS]
- Leplin, J. (2009) *A theory of epistemic justification*. Springer. [GS]
- Levis, G. L. (1999) *The Turkish language reform: A catastrophic success*. Oxford University Press. [aSE]
- Liberman, N. & Trope, Y. (2008) The psychology of transcending the here and now. *Science* 322(5905):1201–205. [YS]
- Lieberman, M. D. (2007) Social cognitive neuroscience: A review of core processes. *Annual Review of Psychology* 58:259–89. [aSE]
- Loftus, G. R. (1996) Psychology will be a much better science when we change the way we analyze data. *Current Directions in Psychological Science* 5:161–71. [GLB]
- Longuenesse, B. (1998) *Kant and the capacity to judge*. Princeton University Press. [CDN]
- Lopes, L. L. (1991) The rhetoric of irrationality. *Theory and Psychology* 1:65–82. [aSE]
- Lopes, L. L. & Oden, G. C. (1987) Distinguishing between random and nonrandom events. *Journal of Experimental Psychology: Learning Memory and Cognition* 13(3):392–400. [YS]
- Lucas, E. J. & Ball, L. J. (2005) Think-aloud protocols and the selection task: Evidence for relevance effects and rationalisation processes. *Thinking and Reasoning* 11(1):35–66. [r]StBTE, EJNS]
- Luhmann, C. C., Chum, M. M., Yi, D.-J., Lee, D. & Wang, X.-J. (2008) Neural dissociation of delay and uncertainty in intertemporal choice. *Journal of Neuroscience* 28(53):14459–66. [YS]
- MacFarlane, J. (2000) *What does it mean to say that logic is formal?* Doctoral Dissertation, University of Pittsburgh. Available at: <http://johnmacfarlane.net/diss.html> [CDN]
- MacIntyre, A. (1988) *Whose justice? Whose rationality?* Notre Dame University Press. [RSN]
- MacIntyre, A. C. (1959) Hume on “is” and “ought.” *Philosophical Review* 68:451–68. [aSE]
- Manktelow, K. I. & Over, D. E. (1991) Social roles and utilities in reasoning with deontic conditionals. *Cognition* 39:85–105. [aSE, r]StBTE]
- Manktelow, K. I., Over, D. E. & Elqayam, S., eds. (2011) *The science of reason: A Festschrift for Jonathan St B. T. Evans*. Psychology Press. [r]StBTE]
- Markman, A. B., Grimm, L. R. & Kim, K. (2009) Culture as a vehicle for studying individual differences. In: *Understanding culture: Theory, research and application*, ed. R. S. Wyer, C. Y. Chiu & Y. Y. Hong, pp. 93–106. Taylor and Francis. [KK]
- Markowitz, H. M. (1991) Foundations of portfolio theory. *Journal of Finance* 46(2):469–77. [YS]
- Marr, D. (1982) *Vision: A computational investigation into the human representation and processing of visual information*. W. H. Freeman. [aSE, r]StBTE]
- Martin, R. L., ed. (1984) *Recent essays on truth and the liar paradox*. Oxford University Press. [aSE]
- McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G. & Cohen, J. D. (2007) Time discounting for primary rewards. *Journal of Neuroscience* 27(21):5796–804. [YS]
- McGee, V. (1985) A counterexample to modus ponens. *Journal of Philosophy* 82:462–71. [aSE]
- Merton, R. K. (1968) Matthew Effect in science. *Science* 159(3810):56–63. [SH]
- Mikhail, J. (2007) Universal moral grammar: Theory, evidence, and the future. *Trends in Cognitive Sciences* 11:143–52. [NG]

- Millikan, R. G. (1984) *Language, thought, and other biological categories: New foundations for realism*. MIT Press. [aSE]
- Millikan, R. G. (1995) *White Queen psychology and other essays for Alice*. MIT Press. [aSE]
- Millikan, R. G. (1996) Pushmi-pullyu representations. In: *Minds and morals*, ed. L. May, L. Friedman & A. Clark, pp. 145–62. MIT Press. [aSE]
- Moore, G. E. (1903) *Principia Ethica*. Available at: <http://fair-use.org/g-e-moore/principia-ethica> [KJPQ]
- Moore, G. E. (1903) *Principia ethica*. Cambridge University Press. [aSE]
- Murphy, G. L. (2002) *The big book of concepts*. MIT Press. [SMcN]
- Nelson, J. (2005) Finding useful questions: On Bayesian diagnosticity, probability, impact, and information gain. *Psychological Review* 112:979–99. [MO]
- Nickerson, R. S. (2000) Null hypothesis significance testing: A review of an old and continuing controversy. *Psychological Methods* 5:241–301. [GLB]
- Nickerson, R. S. (2002) The production and perception of randomness. *Psychological Review* 109(2):330–57. [YS]
- Nickerson, R. S. (2008) *Aspects of rationality: Reflections on what it means to be rational and whether we are*. Psychology Press. [aSE, r]StBTE]
- Nickerson, R. S. & Butler, S. F. (2009) On producing random binary sequences. *American Journal of Psychology* 122(2):141–51. [YS]
- Nisbett, R. E., Krantz, D. H., Jepson, D. H. & Kunda, Z. (1983) The use of statistical heuristics in everyday inductive reasoning. *Psychological Review* 90:339–63. [aSE]
- Nisbett, R., Peng, K., Choi, I. & Norenzayan, A. (2001) Culture and systems of thought: Holistic vs. analytic cognition. *Psychological Review* 108:291–310. [r]StBTE]
- Norenzayan, A. & Heine, S. J. (2005) Psychological universals: What are they and how can we know? *Psychological Bulletin* 131:763–84. [r]StBTE]
- Noveck, I. A. (2001) When children are more logical than adults: Experimental investigations of scalar implicature. *Cognition* 78:165–88. [VG]
- Nozick, R. (1993) *The nature of rationality*. Princeton University Press. [aSE, TF]
- Oaksford, M. & Chater, N. (1991) Against logicist cognitive science. *Mind and Language* 6:1–38. [aSE]
- Oaksford, M. & Chater, N. (1994) A rational analysis of the selection task as optimal data selection. *Psychological Review* 101:608–31. [aSE, r]StBTE, UH, MO, EMP]
- Oaksford, M. & Chater, N. (1996) Rational explanation of the selection task. *Psychological Review* 103:381–91. [aSE]
- Oaksford, M. & Chater, N. (1998a) *Rationality in an uncertain world*. Psychology Press. [aSE, r]StBTE, MO]
- Oaksford, M. & Chater, N., eds. (1998b) *Rational models of cognition*. Oxford University Press. [MO]
- Oaksford, M. & Chater, N. (2003) Optimal data selection: Revision, review and re-evaluation. *Psychonomic Bulletin and Review* 10:289–318. [MO]
- Oaksford, M. & Chater, N. (2007) *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press. [GLB, aSE, r]StBTE, SMcN, MO, GS]
- Oaksford, M. & Chater, N., eds. (2010) *Cognition and conditionals*. Oxford University Press. [r]StBTE]
- Oaksford, M., Chater, N. & Larkin, J. (2000) Probabilities and polarity biases in conditional inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26:883–99. [MO]
- Oppenheimer, D. M. & Monin, B. (2009) The retrospective gambler’s fallacy: Unlikely events, constructing the past, and multiple universes. *Judgment and Decision Making* 4(5):326–34. [YS]
- Osherson, D., Perani, D., Cappa, S., Schnur, T., Grassi, F. & Fazio, F. (1998) Distinct brain loci in deductive versus probabilistic reasoning. *Neuropsychologia* 36(4):369–76. [JP]
- Oskarsson, A. T., Van Boven, L., McClelland, G. H. & Hastie, R. (2009) What’s next? Judging sequences of binary events. *Psychological Bulletin* 135(2):262–85. [YS]
- Over, D. E. (2000) Ecological issues: A reply to Todd, Fiddick, and Krauss. *Thinking and Reasoning* 6(4):385–88. [JMW]
- Over, D. E. (2007) Content-independent conditional inference. In: *Integrating the mind: Domain general versus domain specific processes in higher cognition*, ed. M. J. Roberts, pp. 83–104. Psychology Press. [aSE]
- Over, D. E. & Evans, J. St. B. T. (1997) Two cheers for deductive competence. *Current Psychology of Cognition* 16:255–78. [aSE]
- Over, D. E., Evans, J. St. B. T. & Elqayam, S. (2010) Conditionals and non-constructive reasoning. In: *Cognition and conditionals: Probability and logic in human thinking*, ed. M. Oaksford & N. Chater, pp. 135–51. Oxford University Press. [aSE]
- Papineau, D. (1999) Normativity and judgment. *Proceedings of the Aristotelian Society, Supplementary Volumes* 73:17–43. [JP]
- Peng, K. P. & Nisbett, R. E. (1999) Culture, dialectics, and reasoning about contradiction. *American Psychologist* 54(9):741–54. [KK]
- Peterson, C. R. & Beach, L. R. (1967) Man as an intuitive statistician. *Psychological Bulletin* 68:29–46. [SMcN]
- Pfeifer, N. & Kleiter, G. D. (2002) Experiments on nonmonotonic reasoning: The coherence of human probability judgments. In: *Proceedings of the 1st Salzburg Workshop on Paradigms of Cognition*, ed. H. Leitgeb & G. Schurz. Universität Salzburg, Institut für Philosophie. [Shortened version published in 2005 in *Synthese* 146(1–2):93–109.] [NP]
- Pfeifer, N. & Kleiter, G. D. (2005) Towards a mental probability logic. *Psychologica Belgica* 45(1):71–99. [NP]
- Pfeifer, N. & Kleiter, G. D. (2006) Inference in conditional probability logic. *Kybernetika* 42:391–404. [NP]
- Pfeifer, N. & Kleiter, G. D. (2009) Framing human inference by coherence based probability logic. *Journal of Applied Logic* 7(2):206–17. [NP]
- Pfeifer, N. & Kleiter, G. D. (2010) The conditional in mental probability logic. In: *Cognition and conditionals: Probability and logic in human thought*, ed. M. Oaksford & N. Chater, pp. 153–73. Oxford University Press. [NP]
- Pfeifer, N. & Kleiter, G. D. (2011) Uncertain deductive reasoning. In: *The science of reason: A Festschrift for Jonathan St. B.T. Evans*, ed. K. I. Manktelow, D. E. Over & S. Elqayam, pp. 145–66. Psychology Press. [aSE, NP]
- Pinker, S. (1997) *How the mind works*. W. W. Norton. [GLB]
- Pinker, S. (2002) *The blank slate: The modern denial of human nature*. Viking. [GLB]
- Pirolli, P. (2007) *Information foraging: A theory of adaptive interaction with information*. Oxford University Press. [MO]
- Platt, J. R. (1964) Strong inference. *Science* 164:347–53. [GLB, r]StBTE]
- Pliske, R. & Klein, G. (2003) The naturalistic decision-making perspective. In: *Emerging perspectives on judgment and decision research*, ed. S. L. Schneider & J. Shanteau, pp. 559–85. Cambridge University Press. [RSN]
- Pothos, E. M. & Busemeyer, J. R. (2009) A quantum probability explanation for violations of “rational” decision theory. *Proceedings of the Royal Society of London B* 276:2171–78. [EMP]
- Prinz, J. J. (2007) *The emotional construction of morals*. Oxford University Press. [KJPQ]
- Proust, J. (2009) Adaptive control loops as an intermediate mind-brain reduction basis. In: *Reduction and elimination in philosophy of mind and philosophy of neuroscience*, ed. H. Leitgeb & A. Hieke, pp. 191–219. Ontos. [JP]
- Proust, J. (in press) Mental acts as natural kinds. In: *Decomposing the Will*, ed. T. Vierkant, A. Clark & J. Kiverstein. Oxford University Press. [JP]
- Prowse Turner, J. A. & Thompson, V. A. (2009) The role of training, alternative models, and logical necessity in determining confidence in syllogistic reasoning. *Thinking and Reasoning* 15(1):69. [E]NS]
- Quine, W. V. O. (1960) *Word and object*. MIT Press. [MRW]
- Quinn, W. S. (1989) Actions, intentions, and consequences: The doctrine of doing and allowing. *Philosophical Review* 98:287–312. [NG]
- Ramsey, F. P. (1931) *The foundations of mathematics and other logical essays*. Routledge and Kegan Paul. [MO]
- Ramsey, F. P. (1931/1990) General propositions and causality. In: *Philosophical papers*, ed. D. H. Mellor, pp. 145–63. Cambridge University Press. (Original publication, 1931). [aSE]
- Rapoport, A. & Budson, D. V. (1997) Randomization in individual choice behavior. *Psychological Review* 104(3):603–17. [YS]
- Rawls, J. (1955) Two concepts of rules. *The Philosophical Review* 64(1):3–32. [TA]
- Rawls, J. (1971) *A theory of justice*. Belknap Press of Harvard University Press. [WB]
- Raylu, N. & Oei, T. P. S. (2002) Pathological gambling: A comprehensive review. *Clinical Psychology Review* 22:1009–61. [aSE]
- Rescher, N. (1969) *Many-valued logics*. McGraw-Hill. [aSE]
- Reverberi, C., Shallice, T., D’Agostini, S., Skrape, M. & Bonatti, L. L. (2009) Cortical bases of elementary deductive reasoning: Inference, memory, and metaduction. *Neuropsychologia* 47:1107–16. [JP]
- Reyna, V. F. (2004) How people make decisions that involve risk: A dual-processes approach. *Current Directions in Psychological Science* 13:60–66. [aSE]
- Ricco, R. B. & Overton, W. F. (in press) Reasoning development and dual systems processing: Competence-procedural developmental systems theory. *Developmental Review*. [r]StBTE]
- Rips, L. J. (1989) The psychology of knights and knaves. *Cognition* 31:85–116. [aSE]
- Rips, L. J. (1994) *The psychology of proof*. MIT Press. [aSE]
- Ritov, I. & Baron, J. (1992) Status-quo and omission bias. *Journal of Risk and Uncertainty* 5:49–61. [NG]
- Rosenkrantz, R. (1992) The justification of induction. *Philosophy of Science* 59(4):527–39. [ID]
- Rossiter, M. W. (1993) The Matilda Effect of science. *Social Studies of Science* 23(2):325–41. [SH]
- Samuels, R. (2000) Massively modular minds: Evolutionary psychology and cognitive architecture. In: *Evolution and the human mind: Modularity, language and meta-cognition*, ed. P. Carruthers & A. Chamberlain, pp. 13–46. Cambridge University Press. [GLB]
- Savage, L. J. (1954) *The foundations of statistics*. Wiley. [aSE]
- Savage, L. J. (1972) *The foundations of statistics*, 2nd edition. Dover. [NG]

- Schmidt, J. & Thompson, V. (2008) “At least one” problem with “some” formal reasoning paradigms. *Memory and Cognition* 36(1):217–29. [EJNS]
- 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. [aSE]
- Schroyens, W. (2010) Logic and/in psychology: The paradoxes of material implication and psychologism in the cognitive science of human reasoning. In: *Cognition and conditionals: Probability and logic in human thinking*, ed. M. Oaksford & N. Chater, pp. 69–84. Oxford University Press. [aSE]
- 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. [aSE]
- Schurz, G. (1997) *The is-ought problem: An investigation in philosophical logic*. Kluwer. [aSE, r]StBTE, GS]
- Schurz, G. (1998) Probabilistic semantics for Delgrande’s conditional logic and a counterexample to his default logic. *Artificial Intelligence* 102:81–95. [GS]
- Schurz, G. (2005) Non-monotonic reasoning from an evolutionary viewpoint. *Synthese* 146(1–2):37–51. [GS]
- Schurz, G. (2007) Human conditional reasoning explained by non-monotonicity and probability: An evolutionary account. In: *Proceedings of EuroCogSci07. The European Cognitive Science Conference 2007*, ed. S. Vosniadou, D. Kayser & A. Protopapas, pp. 628–33. Erlbaum. [GS]
- Schurz, G. (2008) The meta-inductivist’s winning strategy in the prediction game: A new approach to Hume’s problem. *Philosophy of Science* 7:278–305. [GS]
- Schurz, G. (2009) Meliorative reliabilist epistemology: Where externalism and internalism meet. *Grazer Philosophische Studien* 79:41–62. [GS]
- Schurz, G. (2010) Non-trivial versions of Hume’s is-ought thesis and their pre-suppositions. In: *Hume on “Is” and “Ought”*, ed. C. Pigden, pp. 198–216. Palgrave Macmillan. [GS]
- Searle, J. R. (1964) How to derive “ought” from “is”. *Philosophical Review* 73:43–58. [aSE, r]StBTE]
- Searle, J. R. (1970) *Speech acts*. Cambridge University Press. [TA]
- Shafir, E. (1993) Intuitions about rationality and cognition. In: *Rationality: Psychological and philosophical perspectives*, ed. K. Manktelow & D. Over, pp. 260–83. Routledge. [KES]
- Shafir, E., ed. (2003) *Preference, belief, and similarity: Selected writings of Amos Tversky*. MIT Press. [KES]
- Shafir, E. & Tversky, A. (1995) Decision making. In: *Thinking, vol. 3*, ed. E. E. Smith & D. N. Osherson, pp. 77–100. The MIT Press. [KES]
- Shepard, R. N. (1992) The perceptual organization of colors: An adaptation to regularities of the terrestrial world? In *The adapted mind*, ed. J. H. Barkow, L. Cosmides & J. Tooby, pp. 495–532. Oxford University Press. [EMP]
- Shynkarkuk, J. M. & Thompson, V. A. (2006) Confidence and accuracy in deductive reasoning. *Memory and Cognition* 34:619–32. [r]StBTE, VAT]
- Simon, H. A. (1982) *Models of bounded rationality*. MIT Press. [aSE, r]StBTE]
- Skyrms, B. (1980) *Causal necessity*. Yale University Press. [ID]
- Sloman, S. A. (1996) The empirical case for two systems of reasoning. *Psychological Bulletin* 119:3–22. [aSE]
- Sloman, S. A. (2005) *Causal models*. Oxford University Press. [r]StBTE]
- Sloman, S. A. (2007) Taxonomising induction. In: *Inductive reasoning*, ed. A. Feeney & E. Heit, pp. 328–43. Cambridge University Press. [SMcN]
- Slovic, P. & Tversky, A. (1974) Who accepts Savage’s axiom? *Behavioral Science* 19:364–71. [aSE]
- 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* 4:108–31. [aSE]
- Smith, J. D., Shields, W. E. & Washburn, D. A. (2003) The comparative psychology of uncertainty monitoring and metacognition. *Behavioral and Brain Sciences* 26(3):317–73. [JP]
- Sober, E. (1993) *Philosophy of biology*. Oxford University Press. [MO]
- Spohn, W. (2002) The many facets of the theory of rationality. *Croatian Journal of Philosophy* 2:247–62. [MRW]
- Stanovich, K. E. (1999) *Who is rational? Studies of individual differences in reasoning*. Erlbaum. [WB, aSE, r]StBTE, VAT]
- Stanovich, K. E. (2004) *The robot’s rebellion: Finding meaning in the age of Darwin*. Chicago University Press. [aSE, r]StBTE]
- Stanovich, K. E. (2008) Higher-order preferences and the Master Rationality Motive. *Thinking and Reasoning* 14:111–27. [r]StBTE]
- Stanovich, K. E. (2009a) Distinguishing the reflective, algorithmic, and autonomous minds: Is it time for a tri-process theory? In: *In two minds: Dual processes and beyond*, ed. J. St. B. T. Evans & K. Frankish, pp. 55–88. Oxford University Press. [aSE, VAT]
- Stanovich, K. E. (2009b) *What intelligence tests miss: The psychology of rational thought*. Yale University Press. [aSE, r]StBTE]
- Stanovich, K. E. (2010a) *Decision making and rationality in the modern world*. Oxford University Press. [aSE]
- Stanovich, K. E. (2010b) *Rationality and the reflective mind*. Oxford University Press. [r]StBTE]
- Stanovich, K. E. (in press) The complexity of developmental predictions from dual process models. *Developmental Review*. [r]StBTE]
- Stanovich, K. E. & West, R. F. (1998) Cognitive ability and variation in selection task performance. *Thinking and Reasoning* 4:193–230. [aSE]
- Stanovich, K. E. & West, R. F. (2000a) Advancing the rationality debate. *Behavioral and Brain Sciences* 23:701–26. [r]StBTE]
- Stanovich, K. E. & West, R. F. (2000b) Individual differences in reasoning: Implications for the rationality debate. *Behavioral and Brain Sciences* 23:645–726. [aSE, EJNS]
- Stanovich, K. E. & West, R. F. (2003) The rationality debate as a progressive research program. *Behavioral and Brain Sciences* 26:531–34. [aSE]
- Stanovich, K. E. & West, R. F. (2007) Natural myside bias is independent of cognitive ability. *Thinking & Reasoning* 13:225. [r]StBTE]
- Stanovich, K. E. & West, R. F. (2008) On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology* 94:672–95. [aSE, r]StBTE]
- Stein, E. (1996) *Without good reason: The rationality debate in philosophy and cognitive science*. Oxford University Press/Clarendon Press. [aSE, r]StBTE, TF]
- Stenning, K. & van Lambalgen, M. (2008) *Human reasoning and cognitive science*. MIT Press. [TA, aSE, r]StBTE]
- Stenning, K. & van Lambalgen, M. (2010) The logical response to a noisy world. In: *Cognition and conditionals: Probability and logic in human thought*, ed. M. Oaksford & N. Chater, pp. 85–102. Oxford University Press. [TA]
- Stenning, K. & Yule, P. (1997) Image and language in human reasoning: A syllogistic illustration. *Cognitive Psychology* 34:109–59. [TA]
- Stich, S. (1981) Inferential competence: Right you are if you think you are. *Behavioral and Brain Sciences* 4(3):353–54. [WB]
- Stich, S. (1990) *The fragmentation of reason: Preface to a pragmatic theory of cognitive evaluation*. MIT Press. [WB, aSE, r]StBTE]
- Stich, S. (2001) Plato’s method meets cognitive science. *Free Inquiry* 21(2):36–38. [WB]
- Stuppel, E. J. N. & Ball, L. J. (2008) Belief-logic conflict resolution in syllogistic reasoning: Inspection-time evidence for a parallel-process model. *Thinking and Reasoning* 14(2):168–81. [EJNS]
- Stuppel, E. J. N. & Waterhouse, E. F. (2009) Negations in syllogistic reasoning: Evidence for a heuristic-analytic conflict. *Quarterly Journal of Experimental Psychology* 62(8):1533–41. [EJNS]
- Sun, S., Pan, W. & Wang, L. L. (2010) A comprehensive review of effect size reporting and interpreting practices in academic journals in education and psychology. *Journal of Educational Psychology* 102:989–1004. [GLB]
- Sun, Y. & Wang, H. (2010a) Gambler’s fallacy, hot hand belief, and time of patterns. *Judgment and Decision Making* 5(2):124–32. [YS]
- Sun, Y. & Wang, H. (2010b) Perception of randomness: On the time of streaks. *Cognitive Psychology* 61(4):333–42. [YS]
- Sun, Y., Tweney, R. D. & Wang, H. (2010a) Occurrence and nonoccurrence of random sequences: Comment on Hahn and Warren (2009) *Psychological Review* 117(2):697–703. [YS]
- Sun, Y., Tweney, R. D. & Wang, H. (2010b) Postscript: Untangling the gambler’s fallacy. *Psychological Review* 117(2):704–705. [YS]
- Thompson, V. A. (1994) Interpretational factors in conditional reasoning. *Memory and Cognition* 22:742–58. [VAT]
- Thompson, V. A., Prowse-Turner, J. & Pennycook, G. (under review) Intuition, metacognition, and reason. [VAT]
- Thomson, J. J. (1976) Killing, letting die, and the Trolley problem. *The Monist* 59:204–17. [NG]
- Todd, P. M., & Gigerenzer, G. (2000) Précis of *Simple heuristics that make us smart*. *Behavioral and Brain Sciences* 23:727–41. [EJNS]
- Tooby, J. & Cosmides, L. (1992) The psychological foundations of culture. In: *The adapted mind: Evolutionary psychology and the generation of culture*, ed. J. H. Barkow, L. Cosmides & J. Tooby, pp. 19–136. Oxford University Press. [GLB]
- Torrens, D., Thompson, V. A. & Cramer, K. M. (1999) Individual differences and the belief bias effect: Mental models, logical necessity, and abstract reasoning. *Thinking and Reasoning* 5(1):1–28. [EJNS]
- Trope, Y. & Liberman, N. (2010) Construal-level theory of psychological distance. *Psychological Review* 117(2):440–63. [YS]
- Trueblood, J. S. & Busemeyer, J. R. (in press) A comparison of the belief-adjustment model and the quantum inference model as explanations of order effects in human inference. *Cognitive Science*. [EMP]
- Tversky, A. & Kahneman, D. (1973) Availability: A heuristic for judging frequency and probability. *Cognitive Psychology* 5:207–32. [VG]
- Tversky, A. & Kahneman, D. (1974) Judgment under uncertainty: Heuristics and biases. *Science* 185:1124–31. [YS]
- Tversky, A. & Kahneman, D. (1981) The framing of decisions and the psychology of choice. *Science* 211:453–58. [NG]
- Tversky, A. & Kahneman, D. (1983) Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review* 90:293–315. [VG]

- Verschueren, N., Schaeken, W., & d'Ydewalle, G. (2005) A dual-process specification of causal conditional reasoning. *Thinking and Reasoning* 11: 239–78. [aSE]
- von Neumann, J. & Morgenstern, O. (1947) *Theory of games and economic behavior*, 2nd edition. Princeton University Press. [aSE, NG]
- Wagenaar, W. A. (1988) *Paradoxes of gambling behaviour*. Erlbaum. [aSE]
- Waldmann, M. R. & Hagmayer, Y. (in press) Causal reasoning. In: *Oxford handbook of cognitive psychology*, ed. D. Reisberg. Oxford University Press. [MRW]
- Waldmann, M. R., Nagel, J. & Wiegmann, A. (in press) Moral judgment. In: *The Oxford handbook of thinking and reasoning*, ed. K. J. Holyoak & R. G. Morrison. Oxford University Press. [MRW]
- Wason, P. C. (1960) On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology* 12:12–40. [aSE]
- Wason, P. C. (1966) Reasoning. In: *New horizons in psychology, vol. 1*, ed. B. M. Foss, pp. 106–37. Penguin. [aSE]
- Wason, P. C. (1968) Reasoning about a rule. *Quarterly Journal of Experimental Psychology* 20:273–81. [MO]
- Wason, P. C. & Brooks, P. G. (1979) THOG: The anatomy of a problem. *Psychological Research* 41:79–90. [aSE]
- Wason, P.C. & Evans, J.St.B.T. (1975) Dual processes in reasoning? *Cognition* 3:141–54. [r]StBTE]
- Weber, E. U. & Hsee, C. K. (2000) Culture and individual judgment and decision making. *Applied Psychology – an International Review [Psychologie Appliquee – Revue Internationale]* 49(1):32–61. [KK]
- Weber, E. U., Shafir, S. & Blais, A.-R. (2004) Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review* 111(2):430–45. [YS]
- Weber, M. (1922/1978) *Economy and Society: An outline of interpretive sociology*, trans. E. Fischoff; ed. G. Rothe & C. Wittich. University of California Press. (Original work published in German in 1922). [NG]
- Wedgwood, R. (2007) *The nature of normativity*. Oxford University Press. [CDN]
- Wegner, D. M. (2002) *The illusion of conscious will*. MIT Books. [r]StBTE]
- Weinberg, J. M. (2007) Moderate epistemic relativism and our epistemic goals. *Episteme: A Journal of Social Epistemology* 4(1):66–92. [r]StBTE, [JMW]
- Whitehead, A. N. & Russell, B. (1910/1962) *Principia Mathematica*. Cambridge University Press. (Original publication, 1910). [aSE]
- Wilkins, M. C. (1928) The effect of changed material on the ability to do formal syllogistic reasoning. *Archives of Psychology*, No. 102. [VAT]
- Williams, B. (1985) *Ethics and the limits of philosophy*. Harvard University Press. [aSE, r]StBTE]
- Wilson, T. D. (2002) *Strangers to ourselves*. Belknap Press. [r]StBTE]
- Wilson, T. D. & Schooler, J. W. (1991) Thinking too much: Introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology*, 60:181–92. [aSE]
- Zsombok, C. E. & Klein, G., ed. (1997) *Naturalistic decision making*. Erlbaum. [RSN]