A peace treaty for the rationality wars? External validity and its relation to normative and descriptive theories of rationality

Annika Wallin

Theory Psychology 2013 23: 458 originally published online 12 June 2013
DOI: 10.1177/0959354313489369

Theory & Psychology

Article

A peace treaty for the rationality wars? External validity and its relation to normative and descriptive theories of rationality

Theory & Psychology
23(4) 458–478

DOI: 10.1177/0959354313489369
tap.sagepub.com
SAGE

Annika Wallin
Lund University Cognitive Science and Swedish Collegium for Advanced Study

Abstract

If we know that certain ways of making decisions are associated with real-life success, is this then how we should decide? In this paper the relationship between normative and descriptive theories of decision-making is examined. First, it is shown that the history of the decision sciences ensures that it is impossible to separate descriptive theories from normative ones. Second, recent psychological research implies new ways of arguing from the descriptive to the normative. The paper ends with an evaluation of how this might affect normative theories of decision-making.

Keywords

philosophy, philosophy of science, scientific progress, theory

If we know that certain ways of making decisions—such as following particular guidelines—are associated with success in real life, is this how we should decide? Are we warranted to move from what appears to be descriptively (and contingently) the case to what *should* be done by a decision-maker?

Often when descriptive and normative issues are discussed it is assumed that they can be treated quite independently. I want to show that at the descriptive level, normative considerations often play a significant role: The descriptive theories we end up with depend on the normative views endorsed. In addition I will present a case in which descriptive considerations can affect which normative models are accepted. The demonstration is based on the rationality wars. This debate on how to interpret experimental

Corresponding author:
Annika Wallin, Lund University Cognitive Science, Department of Philosophy, Kungshuset, Lundagård, 222 22 Lund, Sweden.
Email: annika.wallin@lucs.lu.se

results concerning human decision-making has raged for half a century, and its news value may have deflated. In addition to being a clear example of how descriptive and normative issues co-depend, the debate has, however, undergone some new developments. In particular, new ways of arguing from the descriptive to the normative have opened up in recent years. I will end this paper by evaluating how this might affect the debate.

First, I will introduce the rationality wars. Then, I will discuss the different levels at which decision-making can be understood, from ideal models of rationality to real-life decisions. The rationality wars can be understood at many different levels, and in order to evaluate the debate, we have to be clear about at which level we are operating. Understanding this will, for instance, help us see that the rationality wars are not merely a matter of how "rationality" is defined. Secondly, the possibility of a peace treaty will be considered. This will make it clear that the debate's normative and descriptive levels strongly influence each other. How experimental results are interpreted depends on the preferred normative model. Thirdly, however, it is also possible to resolve the conflict by arguing from empirical data to normative model. If certain ways of approaching decision problems are related to real-life success, then, perhaps, it is those approaches that we should prefer?

The rationality wars: Beginnings

In 1954 Ward Edwards suggested that psychologists systematically test the tenets of economic theory—in particular economic man:

> It is easy for a psychologist to point out that an economic man who has the properties discussed above [complete information, infinite sensitivity, rationality] is very unlike a real man. In fact, it is so easy to point this out that psychologists have tended to reject out of hand the theories that result from these assumptions. This isn't fair. … The most useful thing to do with a theory is not to criticize its assumptions but rather to test its theorems. (p. 382)

The research introduced by this paper culminated with Amos Tversky and Daniel Kahneman's heuristics and biases research. Perhaps the real starting point was a *Science* paper (Tversky & Kahneman, 1974), in which violations of some very fundamental "statistical rules" were summarized and described (p. 1130). Among other things, participants did not seem sensitive to the prior probability of an outcome, to sample size, or to predictability. Through such deviations the authors argued for the existence of heuristics: simple decision rules described as being "highly economical and usually effective," but leading to "systematic and predictable errors" (p. 1131). These, and similar results, have been famously described as having "bleak implications for human rationality" (Nisbett & Borgida, 1975, p. 935). How the results should be interpreted, whether participants violate fundamental statistical rules, and whether the implications for human rationality are truly bleak has been discussed since. The resulting debate has been named the rationality wars by Richard Samuels and colleagues (Samuels, Stich, & Bishop, 2002).

But what is the debate *really* about? Let us start by carefully considering the available data. In the following I will use one of Kahneman and Tversky's most well-known

experiments to illustrate the types of situations that the rationality wars have centred on. The experiment concerns the conjunction effect: that is, how participants make probability judgements involving conjuncts (Kahneman & Tversky, 1982; Tversky & Kahneman, 1982, 1983).

In this experiment participants were first presented with a personality sketch, such as: "Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations" (Tversky & Kahneman, 1983, p. 297). Participants were then asked to rank a number of statements about Linda "by their probability" (p. 297). Among them were declarations such as "Linda works in a bookstore and takes Yoga classes" or "Linda is a psychiatric social worker," but also, and more importantly, "Linda is a bank teller" and "Linda is a bank teller and an active feminist" (p. 297). When the number of statements is about eight, a large number of participants (typically over 80%; Kahneman & Tversky 1982; Tversky & Kahneman, 1982, 1983) rank "Linda is a bank teller and an active feminist" as being more probable than "Linda is a bank teller." This is described as a fallacy since the conjunction of two events (being a bank teller and being an active feminist) cannot be more likely than one of the conjuncts (such as being a bank teller). Even when participants are asked to bet on one of the statements ("If you could win $10 by betting on an event, which of the following would you choose to bet on?"), more than half choose "Linda is a bank teller and an active feminist" (Tversky & Kahneman, 1983, p. 300). This is particularly interesting since an actual bet (remember that this one is hypothetical) would imply that participants are willing to act on their statistical estimate.

The rationality wars concern how we should judge participants' behaviour when they, for instance, claim that it is more probable that Linda is a feminist bank teller. Does this demonstrate human irrationality? There are, at least, two aspects to this question. First, is the conjunction effect truly a fallacy, or are participants' conclusions warranted? Second, is it correct to describe participants' behaviour as a violation of the conjunction rule, or are the experimental results in some way misleading? The experiments' validity has been challenged in many ways. Most challenges have focused on the interpretation of the "and" in "Linda is a bank teller and an active feminist" and of the "probability" in the instruction asking participants to order statements "by their probability." If participants understand the task differently than the experimenters intended, participants' behaviour can be explained without accepting that participants violate the conjunction rule. This debate most recently flared up between Ralph Hertwig and Katya Tentori with colleagues (see Hertwig, Benz, & Krauss, 2008; Tentori & Crupi, 2012). But there are also challenges to the conjunction effect that rather question the fallaciousness of sometimes breaking the conjunction rule. For instance, Isaac Levi (1985, 2004) claims that participants presented with the Linda task see themselves as being asked not to estimate probabilities, but rather to judge which of the propositions is best supported by the data. This is a claim regarding the validity of the experiments—whether they indeed demonstrate the phenomenon they are intended to demonstrate. But it is also a claim regarding the normative status of the response: Levi maintains that there are several accounts, including Bayesian ones, which regard "Linda is a bank teller and an active feminist" to be better supported. This, normative, part of Levi's critique becomes very clear in the

following passage: "Kahneman and Tversky seem to have gone around lecturing at medical schools as to the dangers of the conjunction fallacy. In my judgment, if they were persuasive, they would be hazardous to our health" (Levi, 2004, p. 40).

Levi's attack is thus on two fronts: at the descriptive and at the normative level (or, as I will soon call it, the guidelines level). In order to understand the rationality wars, it is important that we consider these levels separately. Therefore I will take some time to introduce four levels at which decision-making can be understood before I continue with my review of the rationality wars. In later sections I will argue that the rationality wars primarily are due to conflicting normative standards for the decision-making of real people (guidelines).

The four levels of decision-making involved in the rationality wars

The highest level at which the rationality wars can be understood concerns *ideal rational behaviour*. How should an agent ideally choose? A number of different models have been constructed to answer this question, each of them based on different assumptions. For instance, not all consider the cancellation criterion or Leonard Savage's "sure thing principle" necessary (for an introduction to the debate, see Gardenfors & Sahlin, 1988). The agent assumed in these models is also ideal in the sense that it is assumed to have, among other things, full information and infinite sensitivity. The intention of models of ideal rational behaviour is not to be realistic, but rather to explore the nature of rational choice. As Edwards pointed out above, an economic man is very unlike a real man. These models are purely normative and few would claim that they have any descriptive component.

The rationality wars are, however, about what *real* people (in contrast to ideal agents) should do. Levi's concern with the advice handed out in medical schools appears to belong on this level. In order to apply an ideal model of rationality to actual decision-making, a number of assumptions have to be made regarding the information available to the decision-maker and her or his preferences. In a manner of speaking, the models of ideal rational behaviour have to be interpreted as, or translated into, recommendations that are applicable to real life. When Kahneman and Tversky's research is described as having bleak implications for human rationality, the assumption is that if people were to follow the statistical rules tested by these authors, their decision-making would be (more) rational. The rules are assumed to ensure appropriate decisions also in the *absence* of full information or infinite sensitivity. I will refer to this as a *guidelines* approach to human decision-making in order to emphasize that the guidelines are intended for actual decision-makers. When the appropriateness of real people's behaviour is judged, it is further assumed that sufficiently clever experimental set-ups will allow us to estimate, among other things, participants' preferences, so that we can determine whether their decision-making follows the appropriate guidelines or not (see Keys & Schwartz, 2007). Guidelines are normative, of course, since they recommend particular ways of making decisions, but they are also, in a manner of speaking, descriptive: A researcher trying to find out whether decision-makers conform to the ways of deciding s/he advocates is testing a potentially descriptive theory. As will become clear later on, at least two competing

sets of guidelines are advocated in the rationality wars, and it is around these that its major battles are fought.

There is, however, also a purely *descriptive* approach to human decision-making. In principle, descriptive models of decision-making do not have to relate to guidelines or ideal rational models at all. The descriptive model could, for instance, be based completely on inductive research. In practice, however, most descriptive models have been influenced by different sets of guidelines. To remain with the previous example, Kahneman and Tversky explain the conjunction effect by proposing that participants judge the probability of a statement after how representative it is of, in this case, Linda's personality sketch. Since "Linda is a bank teller and an active feminist" bears a closer resemblance than "Linda is a bank teller" to the sketch of an outspoken person who has been deeply concerned with issues of discrimination, participants judge the first statement as being more probable. Kahneman and Tversky claim that participants' probability judgements are formed through their use of the representativeness heuristic. Evidence for the heuristic thus comes from deviations from the guidelines advocated by the authors—in this particular case the conjunction rule. Although the heuristic is purely descriptive, and not in any sense normative, it thus still does, to some extent, depend on a guideline. It is less easy to categorize judgements that do not deviate from a guideline as being produced by a particular heuristic. The judgement that it is more probable that Linda is working in a bookstore and taking Yoga classes than that she is a bank teller can be explained in any number of ways (see Sahlin, Wallin, & Persson, 2010).

Ideal rational models, guidelines for appropriate behaviour, and descriptive accounts are not enough, however. They also have to relate to the real-life decisions people make every day. Guidelines are only interesting if actual people are able to follow them. Descriptive theories have to generalize to settings that differ from those in which they have been established. Even if we accept that participants solving a paper-and-pencil task, such as the one about Linda, are using the representativeness heuristic, it remains to be seen whether they do so also outside of the psychologists' laboratory. Ultimately we want to apply a descriptive theory of human decision-making to all sorts of decisions and situations. Unconstrained and spontaneous decision-making thus constitutes the fourth level at which decision-making can be understood. I will refer to it as *real-life decision-making*, although decisions in laboratories are just as real. Real life decision-making may, of course, differ from decisions made in controlled settings, not only with respect to how spontaneous it is, and what resources it can draw upon, but also with respect to how conscious the decision-maker is of what s/he is doing. We can only hope that careful empirical methods will help us understand the mechanisms underlying our decisions well enough.

In sum, decision-making can be understood at, at least, four different levels. In a way they give rise to one another, as in Figure 1. Interpretations of ideal models of rationality give rise to guidelines, guidelines are used as hypotheses and produce various descriptive theories of decision-making, and these theories hopefully generalize to real-life decision-making. In the following I will demonstrate that the levels relate in complicated ways. They influence each other, and although they have to be treated separately, we cannot assume that they are in any way independent. It is through this co-dependence of levels that the normative and descriptive begin to merge in the rationality wars.

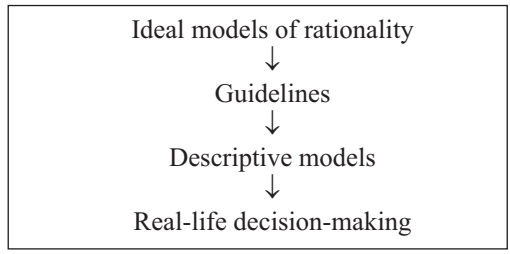

Figure 1. Different levels at which decision-making can be understood.

Dependence of descriptive theories on normative theories in the rationality wars

I have promised to show that the descriptive and normative levels are heavily intertwined in the rationality wars. A first illustration is the fact that a large proportion of empirical research in judgement and decision-making uses recommendations from various guidelines as a base level, or a null hypothesis, such as the representativeness heuristic discussed above. There is, however, more to the co-dependence than this. Owing to the fact that, in the rationality wars, experiments are influenced by guidelines, different views on what the appropriate guidelines are will lead advocates of the various normative camps to interpret empirical results differently. In the following, I will introduce two prominent normative camps and show how the entanglement between their preferred guidelines and empirical work at the descriptive level give rise to the rationality wars. I will come back to the possibility of descriptive research influencing guidelines and ideal rational models at the end of the paper.

In the rationality wars there exist (at least) two different views on what the appropriate guidelines for real life decision-making are. One focuses on the formal properties of the decision-making process, whereas the other is more concerned with an economic use of time and cognitive resources in decision-making. Tversky and Kahneman (1986) are clear proponents of the former.[1] They explicitly describe their research programme as testing the normative rules of "the modern theory of decision making under risk" (p. S252). Among these they count the *transitivity* of preferences (if you prefer *a* to *b* and *b* to *c*, you should also prefer *a* to *c*), *dominance* (if one option is better than others in at least one respect and at least as good as the others in other aspects, chose that option), *invariance* (how an option is represented should not change its value, if representations are isomorphic), and *cancellation* (you do not have to worry about invariant states of the world when you make choices). Edwards (1954) also adopts this sort of guideline. He suggests that we treat economic man as a hypothetical descriptive model of human decision-making: "For instance, if economic man is a model for real men, then real men should always exhibit transitivity of real choices" (p. 382).

In contrast, other researchers focus on the success of real-life decision-making and the constraints under which it operates. The advocates of this camp often refer to Simon as one of their founding fathers, given his clear emphasis on realistic models of decision-making:

> Broadly stated the task is to replace the global rationality of Economic Man with a kind of rational behavior that is compatible with the access to information and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist. (Simon, 1955, p. 99)

That decision-making has to be adapted to its real-life circumstances is one of the core tenets of this line of thought. The "ecological rationality" proposed by Peter Todd and Gerd Gigerenzer is a case in point:

> Traditional definitions of rationality are concerned with maintaining internal order of beliefs and inferences. … But real organisms spend most of their time dealing with the external disorder of their environment, trying to make the decisions that will allow them to survive and reproduce. … To behave adaptively in the face of environmental challenges, organisms must be able to make inferences that are fast, frugal, and accurate. These real-world requirements lead to a new conception of what proper reasoning is: *ecological rationality*. (Todd & Gigerenzer, 2000, p. 736)

On this view, accuracy and cognitive economy are more important than rule-following. Gigerenzer has even advocated decision rules that deviate from such important criteria as transitivity since they have proven to be successful and require little information (Gigerenzer & Goldstein, 1996).

The combatants in the rationality wars thus advocate different guidelines after which human decision-making should be judged. Where one emphasizes the formal properties of the decision-making process, the other completely disregards this, as long as the end result is acceptable. Kenneth Hammond (1996) has named these two camps the coherence (for the rule-followers) and the correspondence approach (for the success-oriented party): "Correspondence researchers are interested in the empirical accuracy of judgments; coherence researchers are interested in the intentional rationality of judgments" (p. 106). I will stick with this terminology, but want to remind the reader that there is no one-to-one mapping between Hammond's use of these concepts and how they are employed in, for instance, epistemology.

Which guidelines you prefer quite obviously affects how experimental studies, such as the Linda example given above, are interpreted. Where Kahneman and Tversky (1982) state that their research raises "doubts about the descriptive adequacy of rational models of judgment and decision making" (p. 124), Gigerenzer (1996) claims that the "norms for evaluating reasoning have been too narrowly drawn, with the consequence that judgments deviating from these norms have been mistakenly interpreted as 'cognitive illusions'" (p. 592). Participants' behaviour is related to different guidelines, and is thus judged differently. If we take the heuristics proposed by Tversky and Kahneman as an example, one and the same outcome can lead to quite different judgements. Tversky and Kahneman (1974) claim their heuristics to be "highly economical and usually effective" at the same time as they lead to "systematic and predictable errors" (p. 1131). Such a distribution of outcomes is entirely acceptable to the correspondence camp, as long as the costs associated with the systematic errors do not outweigh the gains obtained through economy and effectiveness. For the coherence camp this is not the case. The systematic

and predictable errors violate the guidelines advocated by this school of thought. The frequency of such a violation is not particularly important. What matters, however, is whether participants commit an error of application—that is, accept the guidelines they violate—or rather an error of comprehension. In the latter case, we cannot really talk about a violation (see Kahneman & Tversky, 1982).

One might then superficially think that the rationality wars are mostly an issue of semantics. Researchers agree on the evidence, and (basically) on which types of cognitive processes operate in the experimental settings, but disagree as to whether the processes should be labelled rational or not. If the opposing camps would see that this is all there is to the dispute, the rationality wars would end (see Samuels et al., 2002). Such a conclusion is, however, mistaken. First, the disagreement concerning the meaning of "rationality" is more than a matter of semantics. It mirrors a genuine conflict concerning what the appropriate guidelines for human decision-making are. When Levi worries about the advice Kahneman allegedly hands out in medical schools, what he worries about is whether physicians will, in practice, follow it, not what they in the future will label as rational. The conflict concerning appropriate guidelines for real-life decision-making has severe implications and will affect how actual instances of decision-making will be carried out and how they will be judged. The rationality wars cannot end until the guidelines issue is firmly settled. This is made more difficult by the fact that the substantive guidelines issue has not been much discussed in the literature (most exceptions are discussed in the rest of this paper). Second, the disagreement concerning appropriate guidelines affects not only how experimental results are interpreted (is the conjunction effect a fallacy or not?), but also whether they are accepted as genuine experimental results or dismissed as mere experimental artefacts (is there a conjunction effect?). These two fronts of the rationality wars roughly correspond to what Philip Tetlock and Barbara Mellers (2002) call the empirical and the normative boundary conditions debates. There are thus more dimensions to the rationality wars than what is considered normative. In the following I will show why this is the case.

Competing guidelines and empirical research

In the rationality wars, which guidelines are adopted affects not only how experimental results are interpreted but also how valid they are considered to be. If results such as the conjunction effect are to be accepted by psychologists, they have to believe that they can be replicated, that they generalize to other settings, and that they do not depend on non-vital parts of the experimental set-up. These are all issues that concern the *external* and *internal validity* of the experiments. Donald Campbell (1957) describes the two validities in the following way:

> First, and as a basic minimum, is what can be called internal validity: did in fact the experimental stimulus make some significant difference in this specific instance? The second criterion is that of external validity, representativeness, or generalizability: to what populations, settings, and variables can this effect be generalized? (p. 297)

So when it comes to internal validity, we have to estimate how well the experiment is set up: Can we trust the result? Can we, for example, expect test-retest reliability? Are the results statistically significant? Are they experimental artefacts? External validity instead centres on how results generalize. Can we expect them to hold under different circumstances, with different participants, or with variations in the tasks used? Internal validity is often prioritized in psychological research. If results are not trustworthy to begin with, it is fairly useless to worry about whether they generalize or not. Nevertheless questions related to external validity are extremely important to researchers connected to the correspondence camp. There are two reasons for this.

First, the correspondence camp's focus on success pushes external validity to the front. Experimental situations can be constructed where the most odd ways of deciding are successful, but the success that matters on the correspondence interpretation is the one that makes a difference to participants. Decision-makers are assumed to be adaptive, and to adopt decision strategies that generally lead to success: in real life, on average, or in situations that matter. Thus, any experimental test of decision-making must be relevantly similar to the situations and environments that matter to participants. It is quite clear that the correspondence camp must have a clear idea on what these situations and environments are in order to be able to give a coherent account of their decision guidelines. There are several attempts to provide such definitions, from evolutionary psychology's emphasis on the "environment of evolutionary adaptiveness" (Barkow, Cosmides, & Tooby, 1992) to the sampling methods of rational analysis (Anderson & Schooler, 2000), but since it is besides the topic of the current paper, I will not go into it here (see, however, Wallin, 2007). Regardless of how these environments are defined, one of the purposes of such a definition is precisely to guarantee the external validity of any success measurement. In a manner of speaking, the environments identified specify which types of external validity are relevant: that is, to which settings the experimental study is supposed to generalize (or vice versa: which types of settings the experimental set-up is supposed to represent).

Second, the correspondence camp has a historically strong focus on external validity. This school of thought is often traced back to Egon Brunswik's perception research:

> The modern origin of the correspondence view of competence and judgement can be found in the classical treatment of perception by Egon Brunswik. Brunswik challenged the gestalt psychologists' emphasis on perceptual illusions (which, after all, are perceptual inaccuracies) by presenting evidence for the high degree of accuracy of perception regarding the natural world outside the psychologists' artificial laboratory conditions. (Hammond, 1996, p. 109; see also Goldstein & Hogarth, 1997)

Brunswik's insistence on performance in the natural world presupposes external validity and has given rise to the probabilistic view on judgement and decision-making often advocated by correspondence researchers. He is, for instance, well known for research in which he tried to determine to what extent retinal size could be used to predict the actual size of an object. In principle, retinal size is not a good cue for actual size, since both objects' size and their distance to participants can vary. In practice, however, objects tend to be of certain sizes and be looked at from certain distances. Such contingent relations can, and to some extent do, make retinal size a good cue for actual size in

natural environments. Brunswik is, however, not only historically important to the correspondence camp. His emphasis on representative sampling is also a tool for identifying the situations or environments in which decision-making is supposed to succeed (Brunswik, 1944, 1955). In the retinal size study, Brunswik attempted to randomly sample instances in which participants spontaneously looked at objects in their everyday life, and measure the correlation between retinal size and object size in these particular situations. The environment in which the predictive potential of retinal size is measured is thus determined through *representative sampling*. Again, the accuracy of the cognitive process, be it perceptual or judgemental, is determined through its external validity.

Given this, it is not surprising that when a correspondence proponent is presented with research indicating that participants do not behave sensibly, a gut reaction is to immediately question the experimental set-up and how it relates to the everyday settings in which participants operate. The most common critique directed at experimental results such as the conjunction effect is thus to question their external validity. The results could, for instance, be argued to be an experimental artefact produced by unrepresentative stimuli. To take an example, the conjunction effect has been criticized on the grounds that the original research assumes that the "and" in "Linda is a bank teller and a feminist" can be treated as the logical operator "and." In the everyday life of participants (or the linguistic communities they participate in), "and" can, however, be used to convey many other relations, such as temporal order or causality. If participants assume that the "and" used in the experiment is of this less constrained type, the experimental stimuli may take on a different meaning—in which the conjunction effect may not be a fallacy at all, the argument goes. Now, it should be noted that such challenges should be (and are) supplemented by empirical tests in which the suggested explanations for the phenomenon are tested. This is done both by the correspondence camp (who need to demonstrate that their explanation is correct) and by the coherence camp (who need to make sure that the experimental results are internally valid—see, again, the debate between Hertwig et al., 2008 and Tentori & Crupi, 2012).

There is, however, also another way in which the external validity of the experimental findings such as the conjunction effect can be questioned. Even if careful empirical work demonstrates that the conjunction effect does exist in the laboratory, proponents of the correspondence camp can claim that such a finding is irrelevant. If violations of the conjunction rule produce no or only very few negative consequences for participants outside the laboratory setting, we have little reason to try to avoid it, they could argue. The reason for this is, of course, the correspondence guidelines' emphasis on success. The very fact that a particular behaviour leads to success is enough to consider it appropriate, given that it is likely to be successful under a suitable range of circumstances. This stance is evident in the ecological rationality framework. Many fast and frugal heuristics bet on specific environmental structures (e.g., the recognition heuristic bets on a positive relationship between recognition and size—Goldstein & Gigerenzer, 2002). If these environmental structures are very common, then potential mistakes made when the structure is not present are paid for by the average success under "normal" circumstances. Thus, if the heuristic also sometimes, or even often, leads to blatant violations of coherence guidelines, it may still be successful. As a matter of fact, it is perfectly consistent for proponents of the correspondence guidelines to completely disregard experimental

research and focus merely on the success of real-life decision-makers. The drawback of ignoring experimental work is, of course, that it is difficult to establish a sound descriptive theory without it, and few correspondence proponents would therefore do so (see below, or Gigerenzer, 1996).

Both the correspondence camp's focus on success as a criterion for good decisions and its historical roots have led it to emphasize external validity. This type of validity is not as important for those advocating coherence guidelines. A violation of a coherence rule is a violation of a coherence rule regardless of where it occurs. The different emphasis of the coherence and correspondence camps illustrates that the set of guidelines you prefer affect not only how experimental results are interpreted but also whether they are considered as valid experimental results at all.

More on the correspondence and coherence guidelines

Above I have described how different guidelines lead researchers to different normative recommendations and also how they affect how empirical results are interpreted. There are, however, more facets to the rationality wars. Among these, proponents of the coherence and correspondence camps have disagreed on what a successful theory of decision-making should look like, and how empirical research should be conducted. Gigerenzer (1996) has, for instance, criticized Kahneman and Tversky for underspecifying their heuristics:

> The problem with these heuristics is that they at once explain too little and too much. Too little, because we do not know when these heuristics work and how; too much, because, post hoc, one of them can be fitted to almost any experimental result. (p. 592)

Instead, Gigerenzer asks for a careful specification of the heuristic: how information is sampled and integrated and which decisions it gives rise to, and in addition the environmental features which will lead the heuristic to perform better or worse. I will, however, not regard this as a core issue for the rationality wars as such—this is, rather, a disagreement about methodology. There are many researchers in the correspondence camp who have a strong interest in the mechanisms underlying our decisions. For instance, John Payne and colleagues (Payne, Bettman, & Johnson, 1993) have been a major methodological influence on the correspondence camp (see, e.g., Gigerenzer, Todd, & the ABC Group, 1999), although Hammond places them in the coherence camp based on their consistent use of coherence benchmarks.

Making normative theories dependent on descriptive theories in the rationality wars

I hope you are now convinced that the type of guidelines that are adopted in the rationality wars affect both how experimental results are interpreted and whether they are considered as proper psychological findings at all. In the following I will show you that, at least in the rationality wars, it is also possible to argue from the descriptive level (what decision-making looks like in real life) to the normative one (what it should look like).

In addition I will propose that this link from descriptive to normative positively affects the possibilities for a peace treaty in the rationality wars.

Recall the different historical roots of the coherence and correspondence camps. Coherence advocates focus on guidelines linked to normative "rules" such as dominance and independence, adopted from economic man, or other ideal models of rationality. While there is good evidence for why ideal agents should adhere to such rules, proponents of the correspondence camp claim that it is not as clear that they lead to success in real-life situations. The ideal differs too much from the actual. For instance, although intransitive preferences *may* lead to money pumps, it is not self-evident that they will do so for a genuine agent in an everyday setting (see, e.g., Gigerenzer & Goldstein, 1996). For a member of the correspondence camp, potential money pumps carry little weight until they have been linked to (a lack of) real-life success. Coherence guidelines are inherently uninteresting for correspondence advocates.

But there is nothing that prohibits members of the coherence party from being interested in real-life success. The history of the coherence camp has led to an emphasis on the structure of decision-making, but coherence criteria can, of course, be combined with aspirations for real-life success. For instance, arguments involving money pumps get their force from pointing to the potential consequences of not having transitive preferences. In a way this is an attempt to link ideal rationality to real-life outcomes. If it can be demonstrated that following the guidelines set up by the coherence camp leads to successful behaviour, this is potentially good evidence for the innate sensibility of coherence rules. Furthermore, any such link would provide the coherence camp with a powerful weapon against correspondence advocates. If coherence leads to success, it seems as if the correspondence camp is pre-committed to this set of guidelines. In sum, if coherence leads to correspondence, a genuine peace treaty is possible for the rationality wars. Then we would have both a priori and pragmatic reasons for trying to adhere to guidelines such as transitivity and dominance. Note, however, that although the correspondence camp will have to yield to this argument (if it can be constructed), it is in no way necessary for the coherence camp to adopt the strategy. They might instead insist that real-life success is fickle, and that it does not capture the advantages of following their preferred set of guidelines. But such an insistence will, I think, make these guidelines less attractive for the people they are supposed to guide. It is better suited as an argument for the ideal rational models' agents than as a recommendation for how decision-making should proceed in real life.

Interestingly, over the past decade, but most notably the past few years, we have seen researchers examine the real-life success of individuals following or not following coherence guidelines. In the next section, the evidence that is presently available will be reviewed, and then the paper ends with a discussion of the consequences of adopting a strategy that links coherence criteria to success.

Evidence for a link between coherence and correspondence

The first empirical studies linking real-life success to criteria favoured by the coherence camp appear to have emerged in the 1990s. Richard Larrick and colleagues (Larrick,

Nisbett, & Morgan, 1993) examined how effectively achieving "desirable life outcomes" was related to use of "cost–benefit rules" in a set of decision problems designed to be representative of everyday life. The researchers found a positive correlation between salary and answering decision problems in a way that was consistent with cost–benefit reasoning. Similarly the average size of raise obtained by their sample of professors had a weak correlation with cost–benefit reasoning. In a similar vein, Keith Stanovich and colleagues (Stanovich, Grunewald, & West, 2002) investigated how well high school students with multiple suspensions performed on a cost–benefit task (of a more abstract kind). Students with multiple suspensions had a slight tendency to perform worse than students with one or no suspensions, although the groups did not differ in cognitive ability or age.

The purpose of Larrick and colleagues' (1993) study was, quite explicitly, to test the normativeness of the maximization model of microeconomic theory through three predictions. First: "The consequences of using putatively normative rules ought to be superior" (p. 332), and furthermore: "If it were to turn out *not* to be the case that people who use the rules have better outcomes, this would throw doubt on the claims of the model to normativeness" (p. 333). In addition, "intelligent people would be more likely to use cost–benefit reasoning" (p. 333), since such individuals are more likely to identify the most effective decision-making strategies. Last, "people ought to be trainable by the cost–benefit rules in the sense of coming to use them in everyday life choices once they have been exposed to them" (p. 333). Again, the idea is that once decision-makers have observed the (hypothesized) superiority of cost–benefit reasoning, this should become their preferred mode of deciding.

The real-life outcomes linked to cost–benefit reasoning in these studies are, however, rather specific and few. It is not particularly surprising that individuals who tend to engage in cost–benefit reasoning have higher salaries than those that do not. Such a result gives us little reason to assume that this decision-making strategy leads to general success. Furthermore, university professors may not be the best group for investigating the "normativeness" of a model of ideal rational behaviour. (The high school students whom Stanovich and colleagues studied seem far more suitable.)

Keith Stanovich has continued research along these lines together with Richard West, but with a focus on individual differences in decision-making. Their main concern has been how cognitive capacity, measured in various ways (SAT-scores, etc.), correlates with performance on tasks measuring adherence to coherence guidelines. Their initial hypothesis was similar to Larrick and colleagues' in that if there is a benefit to following coherence guidelines, intelligent people ought to be more inclined to do so. In early studies the authors did find some support that higher cognitive capacity increases the likelihood that participants reason in accordance with these guidelines (see, e.g., Stanovich & West, 1998, 2000). Later experiments, however, indicate that the relation between cognitive capacity and coherence guidelines is more complicated. In some cases more gifted participants even deviate more from coherence guidelines than participants of lower cognitive capacity do (Stanovich & West, 2008). Although this research certainly can tell us something about the value of coherence guidelines, it is not focused on external validity, and is thus slightly beside the point of the present paper. Therefore I will not discuss this research in any more detail, but instead present a group of researchers focused on the

decision outcomes obtained by people who do or do not reason in a way that the coherence camp would find acceptable.

The bulk of evidence regarding "real-life success" and adherence to coherence guidelines comes from a set of studies designed by Baruch Fischhoff and colleagues. Their first contribution to the debate is an attempt to formulate a standardized measurement of *decision-making competency*, which can be read as "tendency to follow coherence guidelines" (see below). As Andy Parker and Fischhoff (2005) put it: "If there is a common factor underlying performance on behavioral decision-making tasks, then one can ask more orderly questions about their external validity" (p. 2). Their second major contribution is to identify and measure a wide range of real-life *decision outcomes*. By doing so the researchers can also clarify how decision-making competency and real-life success relate—they can use criteria similar to those advocated by the correspondence camp to evaluate coherence guidelines. I will first give some more details regarding the way decision-making competency and decision outcomes are measured, and then introduce the actual results and discuss their (potential) impact on the rationality wars. Fischhoff and colleagues' measurement of decision-making competency is based on a number of classical experiments gathered from the coherence camp's research. It includes attempts to gauge participants' tendency to succumb to the framing effect, overconfidence, and sunk costs. In particular, items are chosen to measure the *ability to judge probabilities correctly,* the *ability to form values in a coherent way*, the *ability to combine beliefs and values*, and *metacognition*. Ability to judge probabilities correctly consists of items measuring consistency in risk perception. It evaluates whether participants have reasonably calibrated probability judgements, and whether they are able to correctly estimate to what extent other participants in the experiment will agree with a particular norm. Ability to form values in a coherent way measures whether participants' value judgements are insensitive to irrelevant task factors such as sunk costs and framing. Ability to combine beliefs and values measures whether participants can follow a decision rule coherently and correctly and whether they are insensitive to the sequence of events leading to a particular outcome. Metacognition measures participants' degree of overconfidence. All of these tasks can be argued to measure adherence to different coherence guidelines.

I should point out that there are slight variations in how decision competency is measured in the two studies that so far have been published by this group of researchers. The first decision competency study was aimed at youths at risk and used a youth decision-making competency scale (Parker & Fischhoff, 2005). The second and most recently published study attempted to develop and refine the items used and is thus based on an adult decision-making scale instead (Bruine de Bruin, Parker, & Fischhoff, 2007). Although the measurements differ to some extent, they are similar enough for the purposes of this paper, and I will treat them as comparable. There are also later studies, investigating, for instance, the relationship between decision-making competency and executive functions (Del Missier, Mäntylä, & Bruine de Bruin, 2010, 2012), and age (Bruine de Bruin, Parker, & Fischhoff, 2012), but since they do not focus on real-life outcomes, I will not discuss them further here.

The success of everyday decision-making has to be estimated through participants' decision outcomes. In their studies, Fischhoff and colleagues have investigated the outcomes reached by youths at risk (Parker & Fischhoff, 2005), and outcomes reported by a

large group of adults recruited through social service organizations in the Pittsburgh area (Bruine de Bruin et al., 2007). The outcome measures used in these two studies (i.e., how successful real-life decision-making is) differ to a large extent. The youth at risk study focuses on maladaptive risk behaviours, such as delinquency, drug use, and early sexual behaviour. This outcome data had already been assembled (for other purposes) in a longitudinal study, and came from various sources, such as schools, parents, and the youths themselves. In contrast, the adult study tried to measure the success of real-life decision-making through a *decision outcome inventory* including items such as whether participants have ever returned a rented movie without watching it at all or have been kicked out of an apartment or rental property because the lease ran out. It was thus completely based on self-reports. It is, of course, difficult to measure real-life success, and the outcomes gauged in these two studies are not necessarily the best way to operationalize success. As the authors themselves point out, given a specific set of preferences among youths, risk behaviour such as drug use may, for instance, be perfectly reasonable. But success has to be judged somehow, and one way of doing so is to utilize norms commonly accepted in the society. The basic point is that in order to understand the relation between decision-making competency and successful decision-making we have to start somewhere. The important challenge for a critic is not to object to the measures currently used but to improve them. In addition the success measures appear to be as good as (but less specific than) those used by the correspondence camp. In correspondence research successful decision outcomes are commonly identified through constrained tasks (asking participants to guess which of two cities has more inhabitants) or by comparing outcomes to a benchmark set by a decision strategy that is compatible with coherence guidelines (see Wallin, 2007).

Let us accept Fischhoff and colleagues' outcome measures for the time being. What can we then learn about the relation between decision competence—according to the coherence camp's interpretation—and successful real life decision-making? When decision-making competency is understood as one single factor, it correlates significantly with good decision outcomes. Interestingly the correlation holds also when cognitive ability is controlled for (measured by abbreviated IQ tests). Adhering to coherence guidelines is apparently linked to successful real-life decisions irrespective of how intelligent the individual using the rules is. In addition, in both studies, most of the component tasks also correlated positively with real-life decision outcomes. It should be noted that correlations are relatively low, but also that the participants used in these studies are very heterogeneous compared to most psychological studies.

The two studies could be seen as indicating that being able to adhere to coherence guidelines leads to successful real-life decisions. Furthermore, this appears to be the case irrespective of how intelligent the individual using the rules is (assuming that IQ is within the normal range). Where does this leave us? Obviously we do not have enough facts to draw any conclusions at the time being. Results are too few and correlations too low to allow for any strong arguments in the direction from empirical success to normative stance. But if results continue to indicate that being able to follow coherence guidelines leads to improved decision-making, how should the correspondence camp respond to these results? And how should the coherence camp itself deal with its potential success?

Implications

Let us assume for the sake of the argument that the obtained results were completely unambiguous so that adherence to coherence guidelines is reliably related to real-life successful decision-making. We should all be very aware that this is not the case, but if it were the case, where would this leave us?

First of all, it seems clear that an unequivocal relation between decision competency and success commits the correspondence camp to coherence criteria. For instance, Gigerenzer (1991) laments the present lack of connection between descriptive and normative levels:

> Since its origins in the mid-seventeenth century … when there was a striking discrepancy between the judgment of reasonable men and what probability theory dictated—as with the famous St. Petersburg paradox—then the mathematicians went back to the blackboard and changed the equations. … Those good old days have gone. (p. 109)

Gigerenzer advocates changing guidelines so that they fit with what decision-makers do, as long as the decision-makers are "reasonable" people. If a strong relationship between coherence guidelines and real-life success were found, the link would be even stronger. In that case we would know (so we assume) that decision-makers who follow coherence guidelines obtain better decision outcomes than those who do not. That the behaviour we observe is that of reasonable men would then be not only a conjecture, but also an established fact. After all, for the correspondence camp, there is no better way to define being reasonable than getting the outcomes you desire.

In principle, then, the correspondence camp is bound to accept coherence guidelines *if* these are reliably linked to success. In itself this is an important step in the rationality wars. *If* studies such as those of Fischhoff and colleagues are replicated and strengthened, and the results convincingly indicate that coherence is related to success, the two camps should be able to formulate a peace treaty. If this were the case, the advantages of coherence guidelines would come not only from their internal consistency and their relationship to ideal models of rational behaviour, but also from their success: that is, from a fit with correspondence guidelines. As a reviewer pointed out to me, such a (potential) peace treaty is, perhaps, better described as a truce. The difference in guidelines would remain, but since they in this case lead to the same recommended means of decision-making (as proponents of the correspondence camp will recommend the decision-making means that best leads to success), they will cease to matter. I have of course, ignored the issue of frugality here, and merely focused on success, but I think that if it were the case that following coherence consistently led to success for real people and not merely for rational angels, the issue of frugality would become secondary.

We should note that the positive relationship between coherence guidelines and correspondence-like success that we are assuming here could be because adhering to coherence guidelines leads to successful decision-making or because if we select successful ways of making decisions we end up with a set of decision-making strategies that comply with coherence guidelines. Either way we set up our hypothetical case, the results would indicate that coherence and correspondence guidelines overlap. It is, however, interesting to note that the direction of causation could help us determine whether compliance

with coherence guidelines is necessary or sufficient for success. For instance, if selecting the most successful decision strategies would mean that we end up with ways of making decisions that fit with coherence guidelines, then these guidelines appear to be necessary for producing successful decision-making, but there might very well exist decision procedures that fit with coherence guidelines and do not lead to successful decision-making. In contrast, if adhering to coherence guidelines in some way produces success, so that coherence guidelines lead to success, then adhering to coherence guidelines instead is sufficient for successful decision-making. If we leave the hypothetical case for a moment, there are some results from Fischhoff and colleagues' studies that seem to indicate that there are decision strategies that fit with coherence guidelines but do not relate to successful decision-making: Path independence was not related to successful decision outcomes in either study, although attempts to improve the measure have been made (Bruine de Bruin et al., 2007). These results could be taken to indicate that path independence does not affect the success of real-life decision-making (see Bruine de Bruin et al., 2007, p. 949).

In practice, the issue is, of course, more complicated, and the situation we are facing is not in any way similar to the hypothetical case. Both results and measurements are preliminary and the observed correlations relatively weak. So far I have not seen the correspondence camp respond to the evidence outlined above. I suspect that a first move would be to question the measurements, the strength of the correlations, and, presumably, the link between decision-making competency and outcome. My guess is that a core issue would be the definition of success. It is an important issue since the notion of ecological rationality assumes that rationality comes in patches—a heuristic is not assumed to work well over all, but only in certain environments. Presumably the correspondence camp would require results that match the impressing success of fast and frugal heuristics in such environments (see, e.g., Gigerenzer et al., 1999), in order to take them seriously. These are all perfectly sensible reactions, and present many practical difficulties for a truce, but in the hypothetical case, where such doubts could be answered, the rationality wars would have to end. We would at last have agreement on what constitutes good real-life decision-making.

What if the results change though, so that adherence to coherence norms is related to worse life outcomes? For the correspondence camp this would be to revert to status quo, but for the coherence side, consequences are potentially dire. Recall that Larrick and colleagues (1993) stated: "If it were to turn out *not* to be the case that people who use the rules have better outcomes, this would throw doubt on the claims of the model to normativeness" (p. 333). Although I admire the spirit behind this statement, I have to disagree with the conclusion. A non-existent, or even negative, relationship between coherence guidelines and real-life success does not necessarily throw doubts on the coherence guidelines' claim to normativeness. They are not, as the correspondence camp is, precommitted to success. Coherence guidelines are motivated by their relationship to a specific type of ideal models of rational behaviour, and proponents of the coherence camp have no obligation to take real-life success seriously, in particular as they can claim that the need for cognitive economy always present in real-life decision-making (where infinite cognitive capacities are rare) should not affect which decision procedures we agree to be normatively appropriate. We should, however, note that if proponents of the

coherence camp decide to take results indicating a positive relation between coherence guidelines and success as indicating that coherence guidelines are superior, they are also indirectly committed to take less positive results seriously as well. Unfortunately, such a commitment is more one of etiquette than of necessity.

A separate, but important, issue is whether the coherence camp should worry about real-life decision outcomes at all. I am inclined to say that they should, in so far as they want to guide actual decision-making. A researcher satisfied to stay at the level of ideal models of rational behaviour does not have to worry about what real people do, or how their decision-making proceeds. But as soon as a proponent of the coherence camp wants to advocate coherence guidelines as being appropriate for real-life decision-making, success has to be taken into account. First, refusing to deal with decision outcomes should also lead to an abandonment of arguments referring to potential consequences of not following particular decision rules, such as money pump arguments. Second, coherence guidelines will not be particularly convincing if they are completely dissociated from actual decision outcomes. In addition, the types of descriptive theories that we get by ignoring external validity and decision outcomes are potentially misleading. Fischhoff (1996) has phrased it wonderfully:

> It is difficult to study simultaneously how people perceive the world and what they do with those perceptions. A standard strategy in cognitive psychology is to standardize the stimuli, so as to gain access to the process. In this light, participants become something like battery-raised hens, placed in very similar environments, in hope of achieving very similar outcomes. (p. 245)

It is time to release at least some of the hens and see how well they do on their own.

Conclusions

It seems clear that normative and descriptive levels cannot be completely separated in the rationality wars. First, what is taken to be descriptively true is affected by the normative guidelines that are advocated. Second, the correspondence camp questions the division between normativity and facts. Their stance is that we *ought* to do what is in *fact* successful. Thus any attempt to resolve the rationality wars will have to resort to empirical facts. It is possible for the coherence camp to take this route, and indeed some steps have already been taken in that direction. So far, the news is mostly good for coherence advocates: Coherence guidelines appear to be related to successful decision-making. But doing so opens up for new possibilities, such as having to accept that an unsuccessful, but (on the coherence interpretation) normatively correct standard has to be revised.

Acknowledgements

The author wishes to thank Baruch Fischhoff for interesting discussions of decision competence, and several anonymous reviewers for their insightful comments. In addition versions of this paper have been presented to numerous audiences and friends, and I would like to thank them all for their interest and critique.

Funding

The Bank of Sweden Tercentenary Foundation provided financial support for this research.

Note

1. Edwards, and his engineering psychology, is presumably an even stronger proponent, but for the sake of brevity the discussion will be restricted to Tversky and Kahneman in this paper.

References

Anderson, J. R., & Schooler, L. J. (2000). Reflections of the environment in memory. *Psychological Science*, *2*, 396–408.

Barkow, J. H., Cosmides, L., & Tooby, J. (1992). *The adapted mind: Evolutionary psychology and the generation of culture*. New York, NY: Oxford University Press.

Bruine de Bruin, W., Parker, A. M., & Fischhoff, B. (2007). Individual differences in adult decision-making competence. *Journal of Personality and Social Psychology*, *92*, 938–956.

Bruine de Bruin, W., Parker, A. M., & Fischhoff, B. (2012). Explaining adult age differences in decision-making competence. *Journal of Behavioral Decision Making*, *25*, 352–360.

Brunswik, E. (1944). Distal focussing of perception: Size constancy in a representative sample of situations. *Psychological Monographs*, *56*(1), Whole No.

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, *62*, 193–217.

Campbell, D. T. (1957). Factors relevant to the validity of experiments in social settings. *Psychological Bulletin*, *54*, 297–312.

Del Missier, F., Mäntylä, T., & Bruine de Bruin, W. (2010). Executive functions in decision making: An individual differences approach. *Thinking & Reasoning*, *16*, 69–97.

Del Missier, F., Mäntylä, T., & Bruine de Bruin, W. (2012). Decision-making competence, executive functioning, and general cognitive abilities. *Journal of Behavioral Decision Making*, *25*, 331–351.

Edwards, W. (1954). The theory of decision-making. *Psychological Bulletin*, *51*, 380–417.

Fischhoff, B. (1996). The real world: What good is it? *Organizational Behavior and Human Decision Processes*, *65*, 232–248.

Gardenfors, P., & Sahlin, N.-E. (1988). *Decision, probability, and utility: Selected readings*. Cambridge, UK: Cambridge University Press.

Gigerenzer, G. (1991). How to make cognitive illusions disappear. In W. Stroebe & M. Hewstone (Eds.), *Review of social psychology* (pp. 83–115). Chichester, UK: Wiley.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky (1996). *Psychological Review*, *103*, 592–596.

Gigerenzer, G., & Goldstein, D. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *98*, 506–526.

Gigerenzer, G., Todd, P. M., & the ABC Group (1999). *Simple heuristics that make us smart*. Oxford, UK: Oxford University Press.

Goldstein, G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, *109*, 75–90.

Goldstein, W. M., & Hogarth, R. M. (1997). Judgement and decision research: Some historical context. In W. M. Goldstein & R. M. Hogarth (Eds.), *Research on judgment and decision making: Currents, connections, and controversies* (pp. 3–65). Cambridge, UK: Cambridge University Press.

Hammond, K. R. (1996). *Human judgment and social policy: Irreducible uncertainty, inevitable error, unavoidable injustice*. Oxford, UK: Oxford University Press.

Hertwig, R., Benz, B., & Krauss, S. (2008). The conjunction fallacy and the many meanings of and. *Cognition*, *108*, 740–53.

Kahneman, D., & Tversky, A. (1982). On the study of statistical intuitions. *Cognition*, *11*, 123–141.

Keys, D. J., & Schwartz, B. (2007). "Leaky" rationality: How research on behavioural decision making challenges normative standards of rationality. *Perspectives on Psychological Science*, *2*, 162–180.

Larrick, R. P., Nisbett, R. E., & Morgan, J. J. (1993). Who uses the cost–benefit rules of choice? Implications for the normative status of microeconomic theory. *Organizational Behaviour and Human Decision Processes*, *56*, 331–347.

Levi, I. (1985). Illusions about uncertainty. *The British Journal for the Philosophy of Science*, *36*, 331–340.

Levi, I. (2004). Jaakko Hintikka. *Synthese*, *140*, 25–35.

Nisbett, R. E., & Borgida, E. (1975). Attribution and the psychology of prediction. *Journal of Personality and Social Psychology*, *32*(5), 932–943.

Parker, A. M., & Fischhoff, B. (2005). Decision-making competence: External validation through an individual-differences approach. *Journal of Behavioral Decision Making*, *18*, 1–27.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge, UK: Cambridge University Press.

Sahlin, N.-E., Wallin, A., & Persson, J. (2010). Decision science: From Ramsey to dual process theories. *Synthese*, *172*, 129–143.

Samuels, R., Stich, S., & Bishop, M. (2002). Ending the rationality wars: How to make disputes about human rationality disappear. In R. Elio (Ed.), *Common sense, reasoning and rationality* (pp. 236–268). Oxford, UK: Oxford University Press.

Simon, H. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118.

Stanovich, K. E., Grunewald, M., & West, R. F. (2002). Cost–benefit reasoning in students with multiple secondary school suspensions. *Personality and Individual Differences*, *35*, 1061–1072.

Stanovich, K. E., & West, R. F. (1998). Individual differences in rational thought. *Journal of Experimental Psychology: General*, *127*, 161–188.

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

Stanovich, K. E., & West, R. F. (2008). On the relative independence of thinking biases and cognitive ability. *Personality Processes and Individual Differences*, *94*, 672–695.

Tentori, K., & Crupi, V. (2012). On the conjunction fallacy and the meaning of *and*, yet again: A reply to Hertwig, Benz and Krauss (2008). *Cognition*, *122*, 123–134.

Tetlock, P. E., & Mellers, B. A (2002). The great rationality debate. *Psychological Science*, *13*, 94–99.

Todd, P. M., & Gigerenzer, G. (2000). Précis of *Simple heuristics that make us smart*. *Behavioral and Brain Sciences*, *23*, 727–780.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.

Tversky, A., & Kahneman, D. (1982). Judgments of and by representativeness. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 84–98). New York, NY: Cambridge University Press.

Tversky, A., & Kahneman, D. (1983). Extentional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, *90*, 293–315.

Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, *59*, S251–S278.

Wallin, A. (2007). Explanation and environment: The case of psychology. In J. Persson & P. Ylikoski (Eds.), *Rethinking explanation* (pp. 163–176). Dordrecht, The Netherlands: Boston Studies in the Philosophy of Science.

Author biography

Annika Wallin is associate professor in Cognitive Science at Lund University and a Pro Futuris fellow at the Swedish Collegium for Advanced Study in Uppsala. She is fascinated by everyday decision-making, and how it is best understood, from both a normative and a descriptive point of view. Address: Lund University Cognitive Science, Department of Philosophy, Kungshuset, Lundagård, 222 22 Lund, Sweden. Email: annika.wallin@lucs.lu.se