Relevance and rationality[*]

NICHOLAS ALLOTT

Abstract

Subjects' poor performance relative to normative standards on reasoning tasks has been supposed to have 'bleak implications for rationality' (Nisbett & Borgida, 1975). More recent experimental work suggests that considerations of relevance underlie performance in at least some reasoning paradigms (Sperber et al., 1995; Girotto et al., 2001; Van der Henst et al., 2002). It is argued here that this finding has positive implications for human rationality since the relevance theoretic comprehension procedure is computationally efficient and well-adapted to the ostensive communicative environment: it is a good example of bounded and adaptive rationality in Gigerenzer's terms (Gigerenzer and Todd, 1999), and, uniquely, it is a fast and frugal satisficing heuristic which seeks optimal solutions.

1 Bleak implications for rationality?

> I wonder who it was defined man as a rational animal. It was the most premature definition ever given. (Oscar Wilde *The Picture of Dorian Gray*, quoted in Stein (1996: 277).)

A substantial body of research in experimental psychology over the last four decades has established that subjects in widely differing experiments give responses that systematically deviate from logical and probabilistic norms of rationality. (See Kahneman et al., 1982 for examples.) This has been taken to have 'bleak implications for [human] rationality' in general. (Nisbett & Borgida, 1975)

As an example, consider Wason's selection task (Wason, 1966), in which four cards are presented, for example the following together with a conditional statement ' If a card has a 6 on the front it has an E on the back'.

Figure 1

*Many thanks to my supervisor, Deirdre Wilson, for detailed comments on a draft of this paper as well as discussions and advice on the topics discussed. Opinions and mistakes are my own.

The subject is asked which of the cards should be turned over to check the truth of the conditional statement. The normative or 'correct' response, by application of the logical rule of modus ponens, is 6 and A. This is typically given by only a small percentage of subjects. Similarly poor performance is generally observed in experiments in which subjects are asked to estimate probabilities: the responses given are typically far from normative Bayesian answers. (Kahneman and Tversky, 1973; Tversky and Kahneman, 1983)

The heuristics and biases school of thought stresses the negative implications of experiments like these for human rationality. The results are explained with a claim that subjects use non-logical heuristics, rather than normative rules of logic, to make decisions. (See Kahneman et al., 1982, for the heuristics and biases programme; and Evans, 1984, for a heuristic account of the abstract selection task).

A contrasting programme of research is associated with Gerd Gigerenzer and evolutionary psychologists. (See Samuels et al., 2002, for a summary.) They have stressed that performance on reasoning tasks is often much better when the information is presented in a different format, even if the task is formally equivalent. For example, Gigerenzer found that subjects estimate probabilities well when they are presented as frequencies of events rather then fractions or percentages, giving orthodox Bayesian answers that take into account base rates.

Similarly, subjects typically give the normative response to the following version of the selection task:

You are a police officer on duty. It is your job to ensure that people obey the rule 'If a person is drinking beer then that person must be at least 18 years old.' The cards give information about four people in a bar with the person's age on one side, and what they are drinking on the other. Which cards should you turn over to determine whether or not they are breaking the rule?

Drinking beer	Drinking coke	22 years old	16 years old

Figure 2
(Adapted from Griggs and Cox, 1982)

The answer which conforms with modus ponens and is chosen by most subjects is cards 1 and 4.

A number of theorists have proposed that humans are good at reasoning in specific domains. For example Cosmides (1989) proposes that we have an evolved domain-specific mental mechanism for detecting violations of social contracts - a so-called 'cheater detection' faculty. Similarly Gigerenzer has proposed that our probabilistic reasoning abilities only work with frequency data because that is the

format in which humans were exposed to probabilistic data in the evolutionary environment. (Gigerenzer 1994, p142; see also Cosmides and Tooby 1996, p3)

2 An emerging consensus

In a recent paper Samuels, Stich and Bishop (2002) have pointed out that the pictures of human reasoning given by the heuristics and biases school and by evolutionary psychologists are largely congruent, despite apparent disagreements, which are shown to be mainly a matter of emphasis or rhetoric.

The heuristics and biases school's claim that

> people's intuitive judgements on a large number of problems ... regularly deviate from appropriate norms of rationality (op cit, p240)

is entirely compatible with the claims of the opposing school that

> there are many reasoning problems ... on which people's intuitive judgements *do not* differ from appropriate norms of rationality. (op cit, p244)

Samuels et al. sketch out an emerging consensus picture of reasoning in which evolved domain-specific mechanisms do a good, normative job when confronted with problems in formats that they can take as input while other mechanisms deal with other problems and do less well, giving non-normative answers.

Samuels et al. also note a further point of agreement:

> When evaluating human reasoning both evolutionary psychologists and proponents of the heuristics and biases program typically presuppose what Stein has called the standard theory of rationality. (op cit, p 247)

I want to take issue here with both of these points of agreement, drawing on recent work in relevance theory and on a strand of Gigerenzer's work, the fast and frugal heuristics programme. I will argue that the 'standard theory of rationality' is being superseded by naturalised pictures of rationality including bounded rationality and adaptive rationality. Before this, I will briefly describe recent work in relevance theory which partly challenges and partly supplements the emerging picture of reasoning.

3 Relevance theory

Sperber et al. (1995) argue that what underlies successful performance on the selection task is not a domain-specific faculty such as a cheater detection mechanism (Cosmides, 1989) or a pragmatic permission schema (Cheng and Holyoak, 1985) but pragmatic factors affecting interpretation of the conditional statement. They state that what matters is the way that the statement 'If P then Q' achieves relevance. To make this point clear it is necessary to give a brief description of relevance theory.

Relevance theory is a general theory of cognition which defines relevance as a property of inputs to cognitive processes. The relevance of an input is a positive function of the cognitive effects achieved by processing it and a negative function of the effort required to process it. In the case of ostensive-inferential communication, utterances create a presumption of optimal relevance: the speaker is justified in assuming that an utterance is at least relevant enough to be worth processing, and is moreover the most relevant one compatible with the speaker's abilities and preferences. This means that the hearer is justified in following a path of least effort in deriving the explicit meaning and implications of an utterance, stopping when an interpretation has been reached that satisfies his expectations of relevance. This is called the relevance theoretic comprehension procedure. (See Wilson and Sperber, 2002, for a thorough introduction to relevance theory.)

From the definition of relevance as a positive function of cognitive effects and a negative function of processing effort, it follows that in an experimental situation, different interpretations can be made more likely by manipulating the effects which will be achieved by deriving a particular conclusion or the effort a subject will need to expend to derive it.

Returning to the selection task, a conditional statement 'If P then Q' has a number of derivable consequences including the following: that the consequent Q will be true when the antecedent P is satisfied; that P and Q will be true together; that P and not Q will not be true together. Choosing cards on the basis of these interpretations leads respectively to selection of the P card only (6); the P and Q cards (6 and E); or the P and not Q cards (6 and A). To make it likely that participants make the normative choice of the P and not-Q cards the corresponding interpretation must be more relevant then the others in the context. In most contexts this is not the case, but by manipulating the effort and effects involved Sperber et al. were able to obtain a majority of correct responses. The successful scenario involved a card printing machine which is supposed to comply with the conditional statement 'If a card has a 6 on the front it has an E on the back' but which had malfunctioned, printing As instead of Es. Here the conditional statement becomes relevant by implying that the machine will no longer print cards with a 6 on one side and an A on the other. In this scenario, as predicted, these cards were

preferred. (See Sperber et al. 1995; there is also a useful summary in Sperber and Wilson 2002, section 6)

These and further experiments on the selection task (Girotto et al. 2001) demonstrate that the kind of reasoning that is tested in the selection task uses neither domain-general reasoning abilities, nor abilities that are domain-specific in Cosmides' or Cheng and Holyoak's sense. Instead, the mental apparatus which deals with ostensive stimuli appears to be used. This result seems to generalise from the selection task to other experimental reasoning paradigms: for example, Van der Henst, Politzer and Sperber (2002) have recently shown that relevance theory makes successful predictions of subjects' responses in indeterminate relational problems. Thus it seems that at least some of what has been taken to be accomplished by mechanisms dedicated to human reasoning, either domain-general or domain-specific, will turn out to rely on the pragmatic faculty, where by this is understood the mental machinery which deals with ostensive stimuli by making inferences guided by the search for optimal relevance, as described. If this is so, then the question of whether humans are rational depends partly on the rationality of the pragmatics module.

I will argue that the relevance theoretic model of pragmatics presents a procedure that is rational in the way that it exploits environmental structure to arrive at conclusions efficiently without needing to consider all theoretical possibilities. It is important to note that this view of the procedure's rationality rules out an approach based on dividing the question of the rationality of the pragmatics module into two separate questions: (1) is it rational for (certain kinds of) reasoning to be done by the pragmatics module? (2) is the pragmatics module itself rational? I believe that the second of these questions is ill-formed and unanswerable in principle, since the rationality of a mental ability or procedure cannot be assessed in abstraction from the problems with which it deals (or those which it evolved to deal with).

Returning to the general question, then, it is clear that according to the standard model of rationality, the pragmatics module must be seen as irrational or at best only occasionally rational, if we accept that it underlies subjects' performance in the selection task. The module generates different answers to the selection task depending on the context, sometimes P, sometimes P and Q, sometimes P and not Q, and sometimes other combinations, whereas in the literature it is agreed that the 'correct' answer according to the standard model is P and not Q. Rather than accepting this conclusion, I want, as I have said above, to take issue with the idea that the standard model is the correct yardstick for evaluating human rationality. To do so it will be necessary first to look at the motivations behind the standard model and define it.

Following this I will look at some of the most vigorous recent work on fleshing out alternatives to the standard model, the work of Gigerenzer and the ABC research group. I will propose a way in which the relevance theoretic

comprehension procedure can be seen as fitting into this framework as an example of bounded and adaptive rationality, that is, a reasoning mechanism which respects the finite capabilities of human beings and exploits certain facts about the environment in which human beings are situated.

4 Consistency, the standard model and naturalized rationality

Philosophical models of cognitive rationality are often built on the avoidance of inconsistency since from any two inconsistent statements, P and not P, any arbitrary statement follows logically. Consistency is then sometimes given as a minimum criterion for rationality, as in Elster's (1983) 'thin theory of rationality' in which a system of beliefs is rational only if the beliefs are all consistent with each other. Although intended as a minimum criterion for rationality (hence the term 'thin theory') this criterion is too strong to be useful in describing human reasoning since, as a matter of fact, people often hold inconsistent beliefs. Furthermore, Elster's theory is of limited use for cognitive science for another reason: since it only aims to evaluate a set of beliefs as a whole, it has nothing to say about the ways in which beliefs are arrived at.

The 'standard theory of rationality' (Stein, 1996) is a theory which describes what manipulations and combinations of beliefs are allowed while trying to ensure that only consistent beliefs are arrived at. Stein (1996) describes the standard model of rationality in cognitive science as one in which logical principles (and the axioms of probability theory) are transformed into instructions for combining beliefs. Thus modus ponens: 'If *P* and *if P then Q*, then *Q*' gives rise to the principle of reasoning 'If you believe *P* and you believe *if P then Q*, then you should believe *Q*'. Since only logically valid operations are allowed, no inconsistencies will arise if all beliefs used as input are consistent.

There are a number of drawbacks with the standard theory. First, it has nothing to say about dealing with inconsistent beliefs, although these are bound to arise, for example when new information from the senses contradicts an existing belief, as in the following scenario from Kitamura (1995). Gogol believes (1) *You should trust someone unless you have good reason not to.* A wolf tells him, "We will look after your coats." Having no reason to mistrust the wolf he forms belief (2) *The wolves will look after our coats.* Later, returning from swimming in the sea, he sees that the coats and wolves have gone and forms belief (3) *The wolves stole our coats*. Gogol now has two inconsistent beliefs, (2) and (3). He proceeds by discarding (3), although alternative results are possible. The point is that the principles of reasoning involved in dealing with contradictory beliefs cannot be simply derived from principles of logic, since these take only consistent premises as input or deliver arbitrary conclusions.

A second question about the standard theory is discussed by Sperber and Wilson (1986/95, pp 96-7) who point out the consequences of having a set of principles for reasoning which include analogues for all normative logical principles. Believing P forces derivation of an infinite number of beliefs, as for example by or-introduction: P --> P or Q -> P or Q or R ->... Thus, the standard model cannot include all logical principles thought to be normative and it becomes an empirical question which, if any of these beliefs, form part of human reasoning ability.[1]

Finally, the theory, like Elster's 'thin theory', attempts to do too much in trying to ensure that only consistent beliefs are held. It is far from clear that the consequences of having inconsistent beliefs are damaging in the way that inconsistent premises in a logical derivation are damaging. A realistic theory of human reasoning cannot outlaw inconsistent beliefs; it must instead describe procedures for reconciling them or at least failing to derive arbitrary conclusions from them.[2]

It should be noted that Stein characterises but does not advocate the standard model. Instead he puts forward a 'naturalised model of rationality' according to which 'there are normative principles of reasoning' and they can be investigated by balancing 'our first-order judgements about what counts as good reasoning, our more general intuitions about what the normative principles of reasoning are, and various philosophical and scientific theories' (1996, pp 254-5). While I agree that this is a better model for human rationality it is not clear to me how Stein's proposed method of investigating reasoning improves on a general description of good scientific practice. On the other hand it recommends heavier reliance on pre-theoretical intuitions than one might expect, since scientific investigation starts with intuitions but, if it is successful, tends to leave them behind.

The work of Gigerenzer and the ABC research group is a more specific source of ideas about naturalized rationality. Gigerenzer (2002) has argued that consistency is, at most, only a secondary criterion for good reasoning, coming well behind accuracy, speed and transparency.[3] His work contrasts classical visions of unbounded rationality with bounded rationality, a naturalized vision. Unbounded rationality suggests 'building models that perform as well as possible with little or no regard for how time consuming or informationally greedy such models may be'. (Gigerenzer et al., 2002, p 149). These models will therefore be poor models of

---

[1] Sperber and Wilson make the radical proposal that the deductive device has only elimination rules available for use. (Sperber and Wilson, 1986/95, p 96)

[2] See Sperber and Wilson (1986/95, pp 114-5) for a proposal about how the deductive device resolves contradictions.

[3] It is beyond the scope of this paper to assess these criteria in detail; however it seems to me that transparency is not a good criterion for mental processes which are not accessible to consciousness, such as the relevance theoretic comprehension procedure.

human reasoning. Instead, he advocates bounded rationality, which 'suggests designing models specifically to fit the peculiar properties and limits of the mind and the environment'. (op cit., p 149) From the beginning, the assumptions of relevance theory have made the relevance theoretic comprehension procedure a type of bounded rationality. I want to show that there are strong affinities here between relevance theory and Gigerenzer's research programme. I will try to specify how the relevance theory comprehension procedure is rational in Gigerenzer's terms by arguing that it is a special case of a satisficing fast and frugal heuristic: one (perhaps the only one possible) which exploits environmental structure in order to optimize.

5 Visions of rationality

> Cognition is the art of focusing on the relevant and deliberately ignoring the rest. (Gigerenzer and Todd, 1999, section 5.2)

Gigerenzer and Todd (1999, section 2) give the following figure:

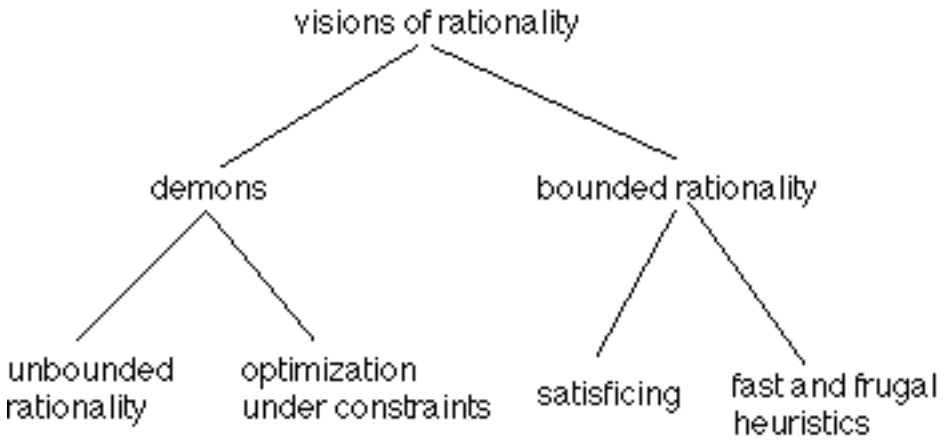

Figure 3

5.1 Unbounded rationality

As noted, unbounded rationality imagines a rational agent as one who weighs up all of the information that could possibly bear on a problem to make a choice or inference. This has been a popular view of rationality in economics and in philosophy. For example, this seems to be Jerry Fodor's vision of central cognition. (Fodor, 2000) He reasons that since we do not know (and cannot,

perhaps, imagine) mechanisms that would allow a finite agent to weigh up all information in a finite time then central cognition must be beyond empirical investigation: 'Classical computations have no way to model [abductive reasoning].' (2000, p 77)

A more optimistic conclusion is that the vision of rationality as unbounded should be abandoned. This conclusion is particularly compelling in light of the fact that there can be no stopping point for information search in this model, since consulting all information known to a subject is not enough - the environment is also full of information, all of which, according to this model, may have to be brought to bear on any decision. (Sperber & Wilson, 1996, p 530. Gigerenzer and Todd, 1999, make a similar point in section 2.2)

5.2 Optimization under constraints

We have established that no realistic model of reasoning, including therefore the relevance theoretic comprehension procedure, can be an example of unbounded rationality. Next I will consider optimization under constraints. This vision of rationality 'holds that the mind should calculate the benefits and costs of searching for each further piece of information and stop search as soon as the costs outweigh the benefits.' (Gigerenzer and Todd, 1999, section 2.2)

Therefore a solution is reached without consulting all of the evidence, in contrast to the way models of unbounded rationality work, using all possible information. However the requirement that at each stage the costs and benefits of containing the search be calculated leads to a computational explosion: 'the paradoxical approach is to model "limited" search by assuming that the mind has essentially unlimited time and knowledge with which to evaluate the costs and benefits of future information search.' (Gigerenzer and Todd, 1999, section 2.2) This means that 'optimization under constraints can require even more knowledge and computation than unbounded rationality.' (op cit, section 2.2, referring to work by Vriend, 1996; Winter 1975)

As a realistic model of cognition, then, relevance theory cannot rely on optimization under constraints; indeed it does not, but there are two reasons why someone might suppose that it does. First, according to the communicative principle of relevance, 'Every ostensive stimulus conveys a presumption of its own optimal relevance' (Sperber and Wilson 1986/95, p 158). This means that the hearer is licensed to search for an optimally relevant interpretation: 'An ostensive stimulus is optimally relevant to an audience [if] it is the most relevant one compatible with the communicator's abilities and preference.' (Wilson and Sperber, 2002, section 3) So the relevance theoretic comprehension procedure looks for the optimal solution, given a particular stimulus in a particular context.

Thus it appears to be an optimizing procedure, but as I shall argue, it is not a kind of optimization under constraints.

Secondly, as previously noted, relevance is a matter of effort and effects, so the generalization 'stop when your expectations of [optimal] relevance are satisfied' (Wilson and Sperber, 2002, section 3) may seem to be an injunction to calculate at each stage the costs and benefits of continuing with the search and to stop when the projected costs in effort outweigh the prospective benefits in cognitive effects.

This is a misinterpretation, however, for two reasons. First, the hearer is licensed to 'stop at the first interpretation that satisfies his expectations of relevance, because there should never be more then one.' (Wilson and Sperber, 2002, section 3) This follows from the special nature of ostensive-inferential communication: the speaker 'wants her utterance to be as easy as possible to understand so that the first interpretation to satisfy the hearer's expectations of relevance is the one she intended to convey.' (op cit. , section 3) This means that there is no need to calculate the costs and benefits of continuing the search: what is at issue is rather whether the cognitive effects are (more than) enough at some time, t, to justify the processing effort incurred from the beginning of the search to that time.

The second reason why the relevance theoretic comprehension procedure could not be a species of optimization under constraints is that from the beginning Sperber and Wilson have been clear that 'contextual effects and processing effort are non-representational dimensions of mental processes' (1986/95, pp 131). We may sometimes have intuitions about degrees of effort and effect but efforts and effects - and therefore relevance - are not generally mentally represented and therefore cannot be used in computations. Thus there is no possibility that future effort and effects could in general be summed and weighed up against each other as optimization under constraints requires.

5.3 Satisficing

Having eliminated daimonic visions of rationality we are left with the conclusion that the relevance theoretic comprehension procedure is a species of bounded rationality. According to Gigerenzer there are (at least) two types of bounded rationality: satisficing and fast and frugal heuristics. I will argue that the relevance theoretic comprehension procedure falls into both of these categories, that is that it is a fast and frugal satisficing heuristic.

The argument that the relevance theoretic comprehension procedure is a type of satisficing procedure comes from comparing Gigerenzer's definition, 'satisficing (sets) an aspiration level and ends the search for alternatives as soon as one is found that exceeds the aspiration level' (Gigerenzer and Todd, 1999, section 2.3, referring to work by Simon, 1956, 1990) with the specification of the relevance theory comprehension procedure:

(a) Follow a path of least effort...(and)
(b) Stop when your expectations of relevance are satisfied. (Wilson and Sperber, 2002 section 3)

Since relevance is a positive function of cognitive effects and a negative function of processing effort, this amounts to processing until the effects derived are worth the effort that has been expended. Thus the amount of effort involved is a kind of dynamic aspiration level, and given that satisficing can involve either fixed or dynamic aspiration levels (Gigerenzer, 2002) the relevance theory comprehension procedure seems to fall under the definition of satisficing.[4]

This may seem unsatisfactory as a conclusion because of the connotations of the term 'satisficing'. It seems to imply that a solution is found that is satisfactory or just good enough and this seems to be opposed to the search for optimal relevance, the best possible solution. I would like to argue that in this case (and probably in no other case) we are dealing with a satisficing procedure which nonetheless seeks an optimal solution. How can a satisficing procedure arrive at an optimal result? The answer I propose is that the relevance theoretic comprehension procedure should be seen, as Sperber and Wilson have stressed, as an adapted piece of mental equipment which exploits a unique environmental regularity. As Gigerenzer and Todd stress, 'the second component of Simon's view of adapted rationality, environmental structure, is of crucial importance because it can explain when and why simple heuristics perform well.' (Gigerenzer and Todd, 1999, section 2.3) Evolved mental abilities incorporate assumptions about the data that they are presented with, so they can rapidly move to correct conclusions, as long as the data really do have those regularities. (See Gigerenzer et al. 2002, pp 161-3 for examples and discussion.) In the case of ostensive-inferential communication, 'relevance theory claims that use of an ostensive stimulus may create precise and predictable expectations of relevance not raised by other stimuli.' (Wilson and Sperber, 2002, section 3) This is because 'an ostensive stimulus is designed to attract the audience's attention. Given the universal tendency to maximise relevance an audience will only pay attention to a stimulus that seems relevant enough. By producing an ostensive stimulus the communicator therefore encourages her audience to presume that it is relevant enough to be worth processing.' (Wilson and Sperber, 2002, section 3)

From these considerations it follows that in the environment of ostensive stimuli a comprehension process can expect to find an optimal solution. Now according to

[4] A possible difference, however is that there may be cases where '…the most accessible interpretation overshoots the mark, and is accepted by Gigerenzer's definition of a satisficing procedure but ruled out by our clause (b) [of the definition of optimal relevance].' (p.c., Deirdre Wilson.)

Gigerenzer and Todd, a 'heuristic is ecologically rational to the degree that it is adapted to the structure of the environment' (1999, section 2.3). If we assume that humans are well adapted to ostensive stimuli, then ecological rationality demands that they must have evolved to look for optimally relevant interpretations. Putting it the other way around, evidence that humans seek optimal relevance from ostensive stimuli is evidence that we are ecologically rational. (See Sperber and Wilson, 2002 for further discussion of the ecological rationality of the pragmatics module.)

5.4 Fast and frugal heuristics

Finally, having discussed the relevance theory comprehension procedure as a satisficing procedure, I want to argue that it is also fast and frugal. Contrary to the implication of mutual exclusivity that figure 3 may give, 'satisficing and fast and frugal heuristics (are) two overlapping but different categories of bounded rationality: there are some forms of satisficing that are fast and frugal and others that are computationally unreasonable.' (Gigerenzer and Todd, 1999, section 2.4)

It is empirically observable that interpretation of ostensive stimuli takes place fast; given human limitations this implies that computation is minimised, although the procedures involved can hardly be trivial. A successful account of the ostensive-inferential device should therefore model it as fast and frugal. In a recent paper Wilson and Sperber (2002) have claimed that the relevance theoretic comprehension procedure can be seen as a fast and frugal heuristic in Gigerenzer's sense since it employs a minimum of time, knowledge and computation. This seems to follow straightforwardly from comparison of the procedure with other possibilities. Within the framework of computation over mental representations it is hard to see what could be computationally simpler or use fewer resources than following a path of least effort, testing potential interpretations in order of accessibility and taking the first one which matches or exceeds expectations[5].

It is an empirical question whether comprehension works this way; as Wilson and Sperber point out, the selection task results and those of other recent experiments provide corroboration for the view that it does.

[5] See footnote 4, above, however, for a way in which there may be cases where the relevance theoretic comprehension procedure yields results that a satisficing heuristic would not.

6 Summary

I have discussed the implications for human rationality of the possibility that considerations of relevance may drive performance in experimental reasoning tasks. The relevance theoretic comprehension procedure does not generally reproduce the normative results expected by the 'standard model' of rationality. Instead it allows us to predict when these results will arise and when subjects will give other answers.

I have argued that the standard model and other models which value consistency of belief above other criteria are unrealistic models of human reasoning and should not be used to evaluate the rationality of reasoning procedures. Instead, naturalized visions of rationality are to be preferred.

Naturalized rationality must take into account human finiteness and the evolved nature of human cognition. Models which respect the first criterion fall under the category of bounded rationality; adaptive rationality is an attribute of models which describe reasoning as well-adapted to the human environment, exploiting its regularities.

I have argued that the relevance theoretic comprehension procedure is rational according to both of these criteria. In Gigerenzer's terms it is fast and frugal, since it is computationally efficient. It also falls under his definition of satisficing, since it stops search when it reaches a (dynamic) aspiration level. Uniquely among satisficing procedures it seeks an optimal solution given a particular stimulus in a context. It is able to do so because it exploits the environmental regularity proposed by Sperber and Wilson, that an ostensive stimulus conveys an expectation of its own optimal relevance. In making use of this regularity, the relevance theoretic comprehension procedure is an excellent example of adaptive rationality.

References

Cosmides, L. (1989) The logic of social exchange: has natural selection shaped how humans reason? Cognition, 31, 187-276.

Cosmides, L. & Tooby, J. (1996) Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgement under uncertainty. Cognition 58: 1-73.

Cheng, P. W., & Holyoak, K. J. (1985) Pragmatic reasoning schemas. Cognition, 17, 391-417.

Evans, J. St. B. T. (1984) Heuristic and analytic processes in reasoning. British Journal of Psychology, 75, 451-468.

Elio, R. (ed.) (2002) Common Sense, Reasoning and Rationality. New York: Oxford University Press.

Elster, J. (1983) Sour Grapes: Studies in the Subversion of Rationality. Cambridge: Cambridge University Press.

Fodor, J (2000) The Mind Doesn't Work That Way. Cambridge, Mass.: MIT Press.

Gigerenzer, G. (2002) Rationality the Fast and Frugal Way. Public lecture at London School of Economics, 8th July.

Gigerenzer, G., Czerlinsky, J. & Martignon, L. (2002) How good are fast and frugal heuristics? In Elio (ed.) 148-73.

Gigerenzer, G. & Todd, P.M. (1999) Precis of Simple Heuristics that Make Us Smart. In Behavioural and Brain Sciences 22:5. (References are by section number to the version available at www.bbsonline/documents/a/00/00/04/69/bbs00000469-00/bbs.todd.html)

Girotto, V., Kemmelmeier, M., Sperber, D., van der Henst, J-B. (2001) Inept reasoners or pragmatic virtuosos? Cognition, 81, 69-76.

Griggs, R. A., & Cox, J.R. (1982) The elusive thematic materials effect in Wason's selection task. British Journal of Psychology, 73, 407-420.

Kahneman, D. & Tversky, A. (1973) On the psychology of prediction. Psychological Review 80: 237 51. Reprinted in Kahneman et al. (eds.) (1982)

Kahneman, D., Slovic, P. & Tversky, A. (eds.) (1982) Judgement Under Uncertainty: Heuristics and Biases. Cambridge: Cambridge University Press.

Kitamura, S. (1995) Sheep in Wolves' Clothing. Great Britain: Andersen Press.

Nisbett, R. & Borgida, E. (1975) Attribution and the social psychology of prediction. Journal of Personality and Social Psychology 32: 932-943.

Samuels, S., Stich, S. & Bishop M. (2002) Ending the rationality wars: How to make disputes about human rationality disappear. In Elio (ed.) 236-68

Simon, H. A. (1956a) Rational choice and the structure of environments. Psychological Review 63: 129-138.

Simon, H. A. (1956b) Dynamic programming under uncertainty with a quadratic criterion function. Econometrica, 24, 19-33.

Simon, H. A. (1990) Invariants of human behaviour. Annual Review of Psychology, 41: 1-19.

Sperber, D., Cara, F., & Girotto, V. (1995) Relevance theory explains the selection task. Cognition, 52, 3-39.

Sperber, D. & Wilson, D. (1986/95) Relevance: Communication and Cognition. Second Edition. Oxford: Blackwell.

Sperber, D. & Wilson, D. (1996) Fodor's frame problem and relevance theory. In Behavioural and Brain Sciences 19:3, 530-532.

Sperber, D. & Wilson, D. (2002) Pragmatics, Modularity and Mind-reading. In Mind and Language 17(1):3-33.

Stein, E. (1996) Without Good Reason: The Rationality Debate in Philosophy and Cognitive Science. Oxford: Clarendon Press.

Tversky, A. & Kahneman, D. (1983) Extensional versus intuitive reasoning: The conjunction fallacy in probability judgements. Psychological Review 90: 293-315.

Van der Henst, J-B., Politzer, G. & Sperber, D. (2002) When is a conclusion worth deriving? A relevance-based analysis of indeterminate relational problems. In Thinking and Reasoning 8:1, 1-20.

Vriend, N. J. (1996) Rational Behaviour and Economic Theory. Journal of Economic Behaviour and Organisation, 29, 263-285.

Wason, P. C. (1966) Reasoning. In B. M. Foss, New horizons in psychology. Penguin.

Wilson, D. & Sperber, D. (2002/ in press) Relevance Theory. To appear in Ward, G. & Horn, L. (eds.) Handbook of Pragmatics. Oxford: Blackwell.

Winter, S. G. (1975) Optimization and evolution in the theory of the firm. In Day, R. H. & Groves, T. (eds.) Adaptive economic models. New York: Academic Press.