The Fundamental Computational Biases of Human Cognition:

Heuristics that (Sometimes) Impair Decision Making and Problem Solving

Keith E. Stanovich

University of Toronto

Stanovich, K. E. (2003). The fundamental computational biases of human cognition: Heuristics that (sometimes) impair decision making and problem solving. In J. E. Davidson & R. J. Sternberg (Eds.), The psychology of problem solving (pp. 291-342). New York: Cambridge University Press.

The Fundamental Computational Biases of Human Cognition: Heuristics that (Sometimes) Impair Decision Making and Problem Solving

Consider the following syllogism. Ask yourself if it is valid--whether the conclusion follows logically from the two premises:

Premise 1: All living things need water

Premise 2: Roses need water

Therefore, Roses are living things

What do you think? Judge the conclusion either logically valid or invalid before reading on.

If you are like about 70% of the university students who have been given this problem, you will think that the conclusion is valid. And if you did think that it was valid, like 70% of university students who have been given this problem, you would be wrong (Markovits & Nantel, 1989; Sá, West, & Stanovich, 1999; Stanovich & West, 1998c). Premise 1 says that all living things need water, not that all things that need water, are living things. So just because roses need water it doesn't follow from Premise 1 that they are living things. If that is still not clear, it probably will be after you consider the following syllogism with exactly the same structure:

Premise 1: All insects need oxygen

Premise 2: Mice need oxygen

Therefore, Mice are insects

Now it seems pretty clear that the conclusion does not follow from the premises.

If the logically equivalent "mice" syllogism is solved so easily, why is the "rose" problem so hard? Well for one thing, the conclusion (roses are living things) seems so reasonable and you know it to be true in the real world. And that is the rub. Logical validity is not about the believability of the conclusion--it is about whether the conclusion necessarily follows from the premises. The same thing that made the rose problem so hard made the mice problem easy. The fact that "mice are insects" is not definitionally true in the world we live in might have made it easier to see that the conclusion did not follow logically from the two premises.

In both of these problems, prior knowledge about the nature of the world (that roses are living things and that mice are not insects) was becoming implicated in a type of judgment (judgments of logical validity) that is supposed to be independent of content. In the rose problem prior knowledge was interfering and in the mice problem prior knowledge was facilitative. In fact, if we really wanted to test a person's ability to process the relationships in this syllogism, we might have used totally unfamiliar material. For example, we might have told you to imagine you were visiting another planet and that you found out the following two facts:

All animals of the hudon class are ferocious.
Wampets are ferocious.

We might then ask you to evaluate whether it logically follows that: Wampets are animals of the hudon class. We can see here that the conclusion does not follow. Research has shown that it is easier to see that the conclusion lacks validity in this unfamiliar version than it is in the rose version, but it is harder to see that the conclusion does not follow in the unfamiliar version than it is in the mice version (Markovits & Nantel, 1989; Sá et al., 1999). These differences prove that factual knowledge is becoming implicated in both the rose and mice problems--even though the content of syllogisms should have no impact on their logical validity. The effect on the rose problem is large. Only about 32% of university students solve it (Sá et al., 1999) whereas the same participants respond correctly 78% of the time on logically equivalent versions with unfamiliar material (versions where prior knowledge does not get in the way).

The rose problem illustrates one of the fundamental computational biases of human cognition--the tendency to automatically bring prior knowledge to bear when solving problems. That prior knowledge is implicated in performance on this problem even when the person is explicitly told to ignore the real-world believability of the conclusion, illustrates that this tendency toward contextualizing problems with prior knowledge is so ubiquitous that it cannot easily be turned off--hence its characterization here as a fundamental computational bias[1] (one that pervades virtually all thinking whether we like it or not). Of course, the tendency to use prior knowledge to supplement problem solving is more often a help than a hindrance. Nevertheless, it will be argued below that there are certain improbable but important situations in modern life in which the fundamental computational biases must be overridden, and that failure to do so can have negative real-life consequences. As research summarized in Baron (1998), Belsky and Gilovich (1999), Dawes (1988), Sutherland (1992), and Thaler (1992) has shown, because of the failure to override fundamental computational biases physicians choose less effective medical treatments; people fail to accurately assess risks in their environment; information is misused in legal proceedings; millions of dollars are spent on unneeded projects by government and private industry; parents fail to vaccinate their children; unnecessary surgery is performed; animals are hunted to extinction; billions of dollars are wasted on quack medical remedies; and costly financial misjudgments are made.

In the remainder of this chapter, I will describe a collection of related processing styles that I have termed the fundamental computational biases of human cognition (Stanovich, 1999). It will be argued, consistent with arguments in evolutionary psychology (e.g., Badcock, 2000; Barkow, Cosmides, & Tooby, 1992; Buss, 1999, 2000; Cartwright, 2000; Cosmides & Tooby, 1994; Pinker, 1997; Plotkin, 1998; Tooby & Cosmides, 1992), that these fundamental computational biases are resident in the brain because they were adaptive in the so-called environment of evolutionary adaptedness (EEA) that existed throughout the Pleistocene (Buss, 1999). In short, it will be argued that these computational biases make evolutionary sense. Nevertheless, it will also be argued that despite their usefulness in the EEA, and despite the fact that even in the present environment they are more useful than not, the modern world presents situations in which the type of contextualization rendered by the fundamental computational biases proves extremely problematic. Such situations are numerically minority situations, but they tend to be ones where a misjudgment tends to have disproportionately large consequences for a person's future utility maximization--for the future fulfillment of the person's life's goals, whatever those goals may be (see Baron, 1993, 1994). In these situations, in order for people to maximize personal utility, they will need to override the fundamental computational biases. We will see that humans have available other cognitive structures (with somewhat different processing architectures) to use in such situations--structures that have the capability of overriding the fundamental computational biases.

In short, in situations where the present human environment is similar to the EEA, the human brain is characterized by fundamental computational biases that bring massive amounts of stored contextual information to bear on the problem. However, when technological societies throw up new problems that confound these evolutionarily adapted mechanisms, humans must use cognitive mechanisms that are in part cultural inventions (see Dennett, 1991) to override the fundamental computational biases that, in these situations, will prime the wrong response. These culturally-induced processing modes more closely resemble the abstract, rule-based, serial processes in many more traditional models of problem solving (Rips, 1994; Sloman, 1996).

The Fundamental Computational Biases

The fundamental computational biases of human cognition work separately and sometimes in combination to ensure that problems with which we are faced are heavily contextualized. They are part of the automatic inferential machinery of the brain that supplements problem solving with stored declarative knowledge, linguistic information, and social knowledge. These processes provide rich supplemental knowledge to augment the sometimes fragmentary and incomplete information we receive when faced with a real-world problem. The four interrelated biases that I will introduce here are: 1) the tendency to contextualize a problem with as much prior knowledge as is easily accessible, even when the problem is formal and the only solution is a content-free rule; 2) the tendency to "socialize" problems, even in situations where interpersonal cues are few; 3) the tendency to see deliberative design and pattern in situations that lack intentional design and pattern; 4) the tendency toward a narrative mode of thought.

Automatic Contextualization: The Use of Prior Knowledge and Context

The property of cognition illustrated in the rose syllogism problem is sometimes termed knowledge projection (Stanovich, 1999). In the rose example, we can see why the tendency to supplement problems with prior knowledge is deemed an automatic tendency. It cannot be "shut off"--even in situations (like that of judging logical validity) where knowledge gets in the way and the problem solver is actively trying to suppress it. The automatic activation of prior knowledge is not limited to syllogistic reasoning problems. Experiments have shown it to operate in several different problem solving domains and in several different paradigms (Banaji & Greenwald, 1995; Evans, Over, & Manktelow, 1993; Greenwald, McGhee, & Schwartz, 1998; Kunda, 1999; Nickerson, 1999; Nickerson, Baddeley, & Freeman, 1987; Sá & Stanovich, 2001; Sá et al., 1999; Stanovich, 1999; Stanovich & West, 1997).

Before looking at some of these additional domains, it should be noted that sometimes the knowledge that is brought to bear on a problem is not accurate. This does not change the nature or logic of the phenomenon. In fact, what often is projected onto a problem is not knowledge at all but instead inadequately substantiated opinion. Thus, since philosophers often define knowledge as justified true belief, in some paradigms in cognitive psychology, strictly speaking, what we are looking at is not knowledge projection--because the belief may not be justified or true. In some cases, we should be talking about belief projection rather than knowledge projection. Nevertheless, we will not distinguish the two here, because the line between them can be very fuzzy and because they appear to reflect an underlying phenomenon that is very similar (Sá & Stanovich, 2001).

One belief projection paradigm that has been extensively studied in the reasoning literature: the evaluation of information in 2 x 2 contingency tables (Levin, Wasserman, & Kao, 1993; Schustack & Sternberg, 1981; Stanovich & West, 1998d). For example, in one such paradigm, participants are asked to evaluate the efficacy of a drug based on a hypothetical well-designed scientific experiment. They are told that:

150 people received the drug and were not cured

150 people received the drug and were cured

75 people did not receive the drug and were not cured

300 people did not receive the drug and were cured

They are asked to evaluate the effectiveness of the drug based on this information. In this case, they have to detect that the drug is ineffective. In fact, not only is it ineffective, it is positively harmful. Only 50% of the people who received the drug were cured (150 out of 300), but 80% of those who did not receive the drug were cured (300 out of 375).

The drug context of this problem is fairly neutral to most participants. But it is easy to trigger prior knowledge and belief by using problems that have more content. For example, in one study, we (see Stanovich & West, 1998d) asked participants to evaluate the outcome of an experiment to test whether having siblings is associated with sociability. The association presented was the same as in the drug experiment:

150 children had siblings and were not sociable

150 children had siblings and were sociable

75 children did not have siblings and were not sociable

300 children did not have siblings and were sociable

Now, however, it was more difficult for our participants (who, as a group, did think that sociability was positively associated with having siblings) to see that, in these data, having siblings was negatively associated with sociability. As in the rose syllogism problem, prior knowledge/belief automatically colored the evaluation of the data. The fact that the numerical paradigm here was quite different from the verbal reasoning domain of syllogistic reasoning indicates the generality of the phenomenon.

Controlled studies (e.g., Broniarczyk & Alba, 1994; King & Koehler, 2000; Levin et al., 1993; Nisbett & Ross, 1980) have demonstrated that when people have a prior belief that two variables are connected, they tend to see that connection even in data in which the two variables are totally unconnected. Unfortunately this finding generalizes to some real-world situations that adversely affect people's lives. For example, many psychological practitioners continue to believe in the efficacy of the Rorschach Test. This is the famous inkblot test in which the client responds to blotches on a white paper. Because the inkblots lack structure, the theory is that people will respond to them in the same style that they typically respond to ambiguity and thus reveal "hidden" psychological traits. The test is called <u>projective</u> because the clients presumably “project” unconscious psychological thoughts and feelings in their responses to the inkblots. The problem with all of this is that there is no evidence that the Rorschach Test test provides any additional diagnostic utility when used as a “projective” test (Dawes, 1994; Garb, Florio, & Grove, 1998; Lilienfeld, 1999; Shontz & Green, 1992; Widiger & Schilling, 1980; Wood, Nezworski, & Stejskal, 1996). Belief in the Rorschach Test arises from the phenomenon of illusory correlation. Clinicians see relationships in response patterns because they believe they are there, not because they are actually present in the pattern of responses being observed (Broniarczyk & Alba, 1994; Chapman & Chapman, 1967, 1969; King & Koehler, 2000).

One final paradigm will illustrate the range of tasks in which knowledge/belief projection has been studied. The paradigm derives from the informal reasoning and problem solving research tradition

(Baron, 1995; Klaczynski, Gordon, & Fauth, 1997; Kuhn, 1991, 1993, 1996; Perkins, Farady, & Bushey, 1991; Voss, Perkins, & Segal, 1991; Sternberg & Wagner, 1986). Unlike deductive reasoning paradigms where there is a single right answer to a problem, research on informal reasoning tries to mimic the type of argumentation and reasoning that goes on in everyday life. This type of argumentation relies on inductive arguments more than deductive arguments. Good arguments in these more realistic contexts tend to be probabilistic and they tend to be what philosophers call defeasible--unlike deductive arguments, they can be defeated by additional information that can be brought to bear (see Hilton, 1995; Johnson-Laird, 1999; Oaksford & Chater, 1995; Pollock, 1995; Stevenson & Over, 1995).

Richard West and I (see Stanovich & West, 1997) developed an argument evaluation test of informal reasoning ability and we can use it to illustrate how belief projection operates in this domain by considering one problem from it. In this particular problem, the respondents are presented with a protagonist, Dale, who believes a particular proposition. In this case, the proposition is that "Students, not the general public, should have the ultimate voice in setting university policies." Dale is said to justify the belief with the following argument: "Because students are the ones who must pay the costs of running the university through tuition, they should have the ultimate voice in setting university policies." A critic then attacks Dale's justification with the following counter-argument: "Tuition covers less than one half the cost of an education at most universities" (the respondent is told to assume that this is factually correct). Finally, Dale rebuts the critic with the following argument: "Because it is the students who are directly influenced by university policies, they are the ones who should make the ultimate decisions." Respondent's are asked to rate the quality of Dale's rebuttal argument, and they are specifically focussed on Dale's argument as a rebuttal of the counter-argument made by the critic. Furthermore, the respondents are very specifically reminded to evaluate the quality of the rebuttal independently of their feelings about Dale's original belief.

This problem was given to eight university professors of philosophy and psychology and all but one of them agreed that the rebuttal argument was weak because it lacked relevance as a refutation of the critic's counter-argument. However, a group of university students (who had a mean prior belief about the matter that corresponded to "agree" on the response scale) gave the rebuttal argument a mean rating of "strong" on the scale. The influence of prior belief was confirmed more directly by the finding that the students who agreed with the original proposition rated Dale's rebuttal significantly more highly than those who disagreed with the original proposition put forward by Dale. In our subsequent work with the argument evaluation test we have found that belief projection operates in a variety of informal reasoning situations and across a host of topics including taxes, the legal drinking age, car safety, the validity of interviews, fairness of Social Security, child care, welfare payments, prison sentencing, the speed limit, the death penalty, the voting age, labor unions, and the effects of secondhand smoke.

The three tasks discussed here--syllogistic reasoning, reasoning about numerical covariation data, and informal reasoning--are just a small sampling from a large number of paradigms that have demonstrated the tendency to contextualize a problem solving situation with prior knowledge and prior belief, even when the problem requires content-free inference and evaluation procedures. These findings confirm the observation of Evans, Barston, and Pollard (1983) that "specific features of problem content, and their semantic associations, constitute the dominant influence on thought" (p. 295).

The tendency to supplement formal, decontextualized problems with prior knowledge is often commented upon in both the problem solving and decision-making literatures. It seems that even in problems that are viewed purely formally by the experimenter, the slightest bit of real-world information seems to provoke contextual inferences. For example, Doherty and Mynatt (1990) used a relatively sparse toy problem to study Bayesian reasoning:

> Imagine you are a doctor. A patient comes to you with a red rash on his fingers. What information would you want in order to diagnose whether the patient has the disease Digirosa? Below are four pieces of information that may or may not be relevant to the diagnosis. Please indicate <u>all</u> of the pieces of information that are necessary to make the diagnosis, but <u>only</u> those pieces of information that are necessary to do so.

Participants then chose from the alternatives listed in the order: % of people without Digirosa who have a red rash, % of people with Digirosa, % of people without Digirosa, and % of people with Digirosa who have a red rash. These alternatives represented the choices of P(D/~H), P(H), P(~H), and P(D/H), respectively, from Bayes' rule for conditioning hypotheses based on data. Our concern here is not with Bayes' rule per se, but the fact that even in a problem as stripped down as this one, participants tended to import contextual information based on their prior knowledge (contextual information that in this case could only be disruptive because all necessary information is already in the problem). Doherty and Mynatt (1990) observed that many participants "brought real-world knowledge to the task, knowledge about the relations between symptoms and diseases in general and knowledge about rashes in particular" (p. 8). These participant responses exemplify the tendency toward automatic contextualization which Evans, Over, and Manktelow (1993) stress is often a generally adaptive characteristic of thought: "The purpose of reasoning is best served by drawing inferences from all our beliefs, not just from an artificially restricted set of premises" (p. 175).

Philosopher Nicholas Rescher (1988) emphasizes how contextualization with prior knowledge beyond the explicitly stated is a pervasive aspect of cognitive life. For example, Rescher (1988) draws attention to the enthymematic character of much human reasoning and problem solving. In logic, a syllogism with an unstated premise is called an enthymeme. In his logic text, Kelley (1990) provides the example "John will support the gun control law because he's a liberal" and notes that the implicit argument is: a. All liberals support gun control; b. John is a liberal; c. John will support a gun control law. But (a) is unstated, so the argument is enthymematic. Rescher argues that "we frequently make substantive assumptions about how things stand in the world on the basis of experience or inculcation, and the incorrect conclusions people draw can stem from these assumptions rather than from any error of inferential reasoning" (pp. 195-196). In reviewing the earlier literature on content effects in reasoning and problem solving, Evans (1982) refers to how it is a ubiquitous finding that "in effect, the subject reasons with an augmented problem space, enriched by relevant personal experience" (p. 225). The enthymematic reasoning styles that form one aspect of the fundamental computational bias are thus natural and nearly universal reasoning styles of human beings----perhaps because of their evolutionary adaptiveness (Cosmides & Tooby, 1992; Pinker, 1997). They will thus facilitate reasoning when context reinforces the conclusion of any explicitly stated information (see Stevenson, 1993). But equally obviously they will be maladaptive when the situation demands a nonenthymematic reasoning style (one where context must be ignored or decoupled).

The enthymematic aspect of the fundamental computational bias thus guarantees that when the environment calls on a person to fix beliefs via a content-free construal--without supplementing with additional inferred information--then they will reason poorly. An example is provided by the most investigated task in the entire reasoning and problem solving literature--Wason's (1966) selection task. The participant is shown four cards lying on a table showing two letters and two numbers (A, D, 3, 8). They are told that each card has a number on one side and a letter on the other and that the experimenter has the following rule (of the if P, then Q type) in mind with respect to the four cards: "If there is a vowel on one side of the card, then there is an even number on the other side". The participant is then told that he/she must turn over whichever cards are necessary to determine whether the experimenter's rule is true or false. Performance on such abstract versions of the selection task is extremely low (Evans, Newstead, & Byrne, 1993; Manktelow, 1999; Newstead & Evans, 1995). Typically, less than 10% of participants make the correct selections of the A card (P) and 3 card (not-Q)--the only two cards that could falsify the rule. The most common incorrect choices made by participants are the A card and the 8 card (P and Q) or the selection of the A card only (P).

Numerous alternative explanations for the preponderance of incorrect PQ and P responses have been given (see Evans, Newstead, & Byrne, 1993; Hardman, 1998; Johnson-Laird, 1999; Liberman & Klar, 1996; Margolis, 1987; Newstead & Evans, 1995; Oaksford & Chater, 1994; Sperber, Cara, & Girotto, 1995; Stanovich & West, 1998a). Notably, several of these alternative explanations involve the assumption that the participant is reading more into the instructions than is actually there. For example, Oaksford and Chater's (1994) analysis of the selection task assumes that participants approach the task as an inductive problem in data selection with assumptions about the relative rarity of the various <u>classes</u> (vowels, odd numbers) of cards. That is, despite the fact the instructions refer to <u>four cards only</u>, it is proposed that the participant is thinking that they are sampling from four <u>classes</u> (a bunch of vowels, a bunch of consonants, etc.). Now imagine that you are verifying the statement "if you eat tripe you will get sick" in the real world. Of course you would sample from the class "people who eat tripe" to verify or falsify it. However, would you sample from the class of "people who are not sick?" Probably not, because this class is too big. But you might well sample from the class "people who are sick" to see if any of them have eaten tripe.

Such an approach to the problem is entirely enthymematic. The participant is assumed to be adding details and context to the problem--none of which are present in the actual instructions. Nothing at all has been said in the instructions about <u>classes</u>--the instructions refer to four cards only. But the alternative explanation of Oaksford and Chater (1994) assumes that the participants are thinking in terms of sampling from classes of cards. Participants even have implicit hypotheses about the relative rarity of these classes according to the particular model of performance championed by Oaksford and Chater (1994). Also, despite the fact that the instructions speak in terms of determining truth and falsity, most participants are thought to ignore this and instead to think in terms of inductive probabilities. In short, the deductive requirements of the instructions violate the fundamental computational bias of most participants who proceed instead to solve a contextualized, familiar, inductive version of the problem.

Margolis (1987) has proposed another highly enthymematic interpretation on the part of participants to explain the incorrect choice--one that again involves the participant thinking about categories even when nothing about categories has been mentioned in the instructions. He argues that some individuals develop an open reading of the task (in terms of choosing entire categories) rather than a closed reading (in terms of choosing individual cards). Margolis (1987) demonstrates the distinction by suggesting a selection task in which participants are given the rule: "If it says swan on one side of a card, it must say white on the other." Given the four categories of cards--swan, raven, black, white--the participants are asked which categories they would need to examine exhaustively to test whether there has been a violation of the rule. Many people familiar with selection task do not recognize that this is not a selection task in the traditional sense. The correct answer is not that categories P and not-Q should be picked. In this open scenario, only P or not-Q need be chosen, but not both. Examining all swans will reveal any violations of the rule--as will examining all black things. But examining both only provides a redundant opportunity to see a violation. This open scenario (where one chooses categories) is different from the closed selection task scenario where one must check designated exemplars of the stated rule. If the participant views the task from an open scenario then P-only is a reasonable choice. If the participant adopts a biconditional interpretation, then they should choose categories P and Q--the modal response in the selection task.

There are continuing disputes about whether either the Oaksford and Chater (1994) or the Margolis (1987) interpretations of the selection task is the correct one (see Gebauer & Laming, 1997; Griggs, 1989; Hardman, 1998; Oaksford & Chater, 1995, 1996; Sperber, Cara, & Girotto, 1995; Stanovich & West, 1998a). It looks instead like there a variety of different interpretations that could lead to the PQ response and that the task is probably characterized by multiple interpretations (see Gebauer & Laming, 1997). The key point is that both of these interpretations posit the recruitment of knowledge and scenarios <u>extrinsic</u> to the instructions--an enthymematic style of reasoning that adds unstated information to further contextualize the problem.

<u>The Tendency to "Socialize" Abstract Problems</u>

Hilton (1995) describes the following conversational exchange:

> Q: How is your family?
>
> A: Fairly well, thank you.

Then he asks us to consider the same man who answered the above question engaging in the next exchange with a different person--perhaps even on the same day:

> Q: How is your wife?
>
> A. Not too good I'm afraid
>
> Q: And how is your family?
>
> A: Extremely well, thank you.

The man has answered the exact same question ("how is your family?") in two different ways on the same day. Should we consider the man's responses irrational? Hilton (1995) argues that of course we would not. Context is everything here. What if the man's wife had recently lost a close friend but the man's two children were doing just great. That might have provoked the "fairly" well in the first exchange because the man took "family" to mean wife and kids. But in the second instance the man had already provided information about his wife so he feels bound to interpret "family" as "the kids only" so that he does not burden the questioner with knowledge that the questioner already has.

Hilton (1995) analyzes this case in terms of Grice's (1975) norms of rational communication (see Hilton & Slugoski, 2000; Sperber & Wilson, 1986; Sperber et al., 1995) which require that the speaker be cooperative with the listener--and one of the primary ways that speakers attempt to be cooperative is by not being redundant. The key to understanding the so-called Gricean maxims of communication is to see that to understand a speaker's meaning the listener must comprehend not only the meaning of what is spoken but also what is implicated in a given context assuming that the speaker intends to be cooperative. So, for example, Hilton (1995) notes that the statement "I went to the cinema last night" is taken to imply that you saw a film even though that was not explicitly stated. Hilton (1995) points out that the "default assumption of conversational rationality enjoins...[a person] to go beyond the information given in making inferences about what is required of them" (p. 257). These additional pieces of information are termed conversational implicatures. Hilton (1995) illustrates that these implicatures share properties with inductive inference. For example, they are ampliative--they are conclusions that contain more information than the premises (see also, Levinson, 1995).

Grice (1975; see also Hilton, 1995; Hilton & Slugoski, 2000; Levinson, 1983, 1995; Schwarz, 1996; Sperber & Wilson, 1986) embeds under his superordinate principle of communicative cooperativeness four maxims: quality (try to make sure that your contribution is true); quantity (make your contribution informative but do not make it more informative than is required); relation (make your contribution relevant); and manner (do not be unnecessarily ambiguous or obscure). Hilton (1995) stresses that the Gricean maxims apply not just to the production of speech but to our comprehension and interpretation as well. Indeed, he develops the idea that humans automatically apply them--even in situations that are not characterized by interactive exchange. Hilton (1995) argues that people treat even decontextualized and depersonalized situations as conversational exchanges with an active interlocutor--"that no utterance is depersonalized, all messages have a source, and that reasoning and inference processes typically operate on socially communicated information" (p. 267). In short, people tend to treat even depersonalized communications as quasi-personal linguistic encounters in a social context. As in our previous discussion of the automatic recruitment of prior knowledge, this too is a type of automatic use of stored information. But in the present case, the additional information that is recruited is information about pragmatic expectancies in conversational exchanges. As with the use of prior knowledge, the linguistic socialization of situations is most often a useful cognitive skill (indeed a necessary and quite astounding one, see Levinson, 1995). However, as will be argued in a later section, it can create difficulties for problem solving in a modern technological society with its high degree of decontextualization.

It has not proven difficult to find experimental situations where Gricean implicatures were operating to make the solution of a decontextualized problem difficult. Consider the demonstration by Tversky and Kahneman (1983) that people show a conjunction fallacy in probabilistic reasoning. An example is provided by the much-investigated Linda Problem (Tversky & Kahneman, 1983):

> 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. Please rank the following statements by their probability, using 1 for the most probable and 8 for the least probable.
>
> a. Linda is a teacher in an elementary school
>
> b. Linda works in a bookstore and takes Yoga classes
>
> c. Linda is active in the feminist movement
>
> d. Linda is a psychiatric social worker
>
> e. Linda is a member of the League of Women Voters
>
> f. Linda is a bank teller
>
> g. Linda is an insurance salesperson
>
> h. Linda is a bank teller and is active in the feminist movement

Because alternative h (Linda is a bank teller and is active in the feminist movement) is the conjunction of alternatives c and f, the probability of h cannot be higher than that of either c (Linda is active in the feminist movement) or f (Linda is a bank teller), yet 85% of the participants in Tversky and Kahneman's (1983) study rated alternative h as more probable than f, thus displaying the conjunction fallacy. Those investigators argued that logical reasoning on the problem (all feminist bank tellers are also bank tellers, so h cannot be more probable than f) was trumped by a heuristic based on so-called representativeness that primes answers to problems based on an assessment of similarity (a feminist bank teller seems to overlap more with the description of Linda than does the alternative "bank teller"). Of course, logic dictates that the subset (feminist bank teller)--superset (bank teller) relationship should trump assessments of representativeness when judgments of probability are at issue.

A large literature on the conjunction effect has established that representativeness is not the only reason that the conjunction effect occurs (Hertwig & Gigerenzer, 1999). However, our interest here is in several explanations that posit that participants are imposing on the problem social assumptions of conversational relevance (Adler, 1984, 1991; Dulany & Hilton, 1991; Hertwig & Gigerenzer, 1999; Hilton, 1995; Hilton & Slugoski, 2000; Macdonald & Gilhooly, 1990; Slugoski & Wilson, 1998). Hilton (1995), for example, provides a Gricean explanation of participants' behavior on the Linda Problem. Under the assumption that the detailed information given about the target means that the experimenter knows a considerable amount about Linda, then it is reasonable to think that the phrase "Linda is a bank teller" in the list of alternatives does not contain the phrase "and is not active in the feminist movement" because the experimenter already knows this to be the case. If "Linda is a bank teller" is interpreted in this way, then rating h as more probable than f no longer represents a conjunction fallacy. Indeed, there is some evidence that participants are reading this into the problem (Dulany & Hilton, 1991; Hertwig & Gigerenzer, 1999; Messer & Griggs, 1993; Politzer & Noveck, 1991; Tversky & Kahneman, 1983).

Schwarz (1996) discusses a host of demonstrations like this--situations in the reasoning and problem solving literature where participants interpret depersonalized and decontextualized communications as if they were personalized communications from someone engaged in an ongoing interaction with them. The phenomenon is real. But this leaves the question of why would someone interpret the words on a sheet of paper from an experimenter or survey researcher (who is often not even physically present) as an exchange in an ongoing personalized discussion? Recent speculative, interdisciplinary theories about the origins of human intelligence may provide the answer. These theories, although varied in their details, all posit that much of human intelligence has foundations in social interaction (Baldwin, 2000; Barton & Dunbar, 1997; Blackmore, 1999; Brothers, 1990; Byrne & Whiten, 1988; Bugental, 2000; Caporael, 1997; Cosmides, 1989; Cosmides & Tooby, 1992; Cummins, 1996; Dunbar, 1998; Gibbard, 1990; Gigerenzer, 1996b; Goody, 1995; Humphrey, 1976, 1986; Jolly, 1966; Kummer, Daston, Gigerenzer, & Silk, 1997; Levinson, 1995; Mithen, 1996; Tomasello, 1998, 1999; Whiten & Byrne, 1997).

In a seminal essay that set the stage for this hypothesis, Nicholas Humphrey (1976) argued that the impetus for the development of primate intelligence was the need to master the social world moreso than the physical world. Based on his observation of nonhuman primates, Humphrey (1976) concluded that the knowledge and information processing necessary to engage efficiently with the physical world seemed modest compared to the rapidly changing demands of the social world with its everchanging demands to engage in inter-personal relationships in a changing social milieu. Humphrey (1976) posited that the latter was the key aspect of the environment that began to bootstrap higher intelligence in all primates. Humphrey's (1976) hypotheses (and other related speculations, see Jolly, 1966) in part helped to spawn a host of research on so-called Machiavellian intelligence (Byrne & Whiten, 1988; Whiten & Byrne, 1997)--the ability to engage in multiple and rapidly changing cooperative and competitive interchanges with a host of conspecifics and maintain these relationships over a long period of time.

The key aspects of Humphrey's (1976) speculations that are critical for the present discussion are that all further advances in cognitive capabilities were built upon a social substrate and that these social mechanisms are still present in our brains and pervasively color all of our cognition. They were not replaced by the more analytic intelligence developed in humans. Thus, a social orientation toward problems is always available as a default processing mode when computational demands become onerous. The cognitive architecture is one where analytic cognition was laid down on top of the modules for social cognition that were already in place. In this feature, the architecture proposed is like Dennett's (1996) notion of four different "kinds of minds" which all overlap with each other in the brain and which all are simultaneously active in controlling behavior (see Stanovich & West, 2000).

Because the social exchanges that provided the environment of evolutionary adaptation for social intelligence involved fluid transactional exchanges with partners engaged in their own dynamic problem solving, Humphrey (1976) argues that the thinking styles involved became the dominant or default ways of approaching problems. Humphrey (1976) felt that this processing bias--the "predisposition among people to try to fit nonsocial material into a social mould" (p. 312)--would always compete with subsequent types of intelligence ("kinds of minds") that might develop. The cognitive illusions demonstrated by three decades of work in problem solving, reasoning, and decision making (Evans, 1989; Kahneman, Slovic, & Tversky, 1982; Kahneman & Tversky, 1996, 2000; Stanovich, 1999) seem to bear this out. As in the Linda Problem and four-card selection task discussed above, the literature is full of problems where an abstract, decontextualized--but computationally expensive--approach is required for the normatively appropriate answer. However, often, alongside such a solution, resides a tempting social approach ("oh, yeah, the author of this knows a lot about Linda") that with little computational effort will prime a response.

In short, the social intelligence hypothesis posits that evolutionary pressures were focused more on negotiating cooperative mutual intersubjectivity than on understanding the natural world. Having as its goals the ability to model other minds in order to read intention and to make rapid interactional moves based on those modeled intentions, interactional intelligence (to use Levinson's, 1995, term) is composed of the mechanisms that support a Gricean theory of communication that relies on intention-attribution. This social, or interactional intelligence, forms that substrate upon which all future evolutionary and cultural developments in modes of thought are overlaid.

Thus, as Humphrey (1976) puts it, "for better or worse, styles of thinking which are primarily suited to social problem-solving colour the behaviour of man and other primates even toward the inanimate world" (p. 316). This social orientation toward the inanimate world leads to one of the other fundamental computational biases that is both a strength of cognition and also sometimes a source of error when dealing the abstract aspects of the modern technological world--the tendency to see deliberate design and pattern in situations that lack intentional design.

<u>Seeing Intentional Design in Random Events</u>

The tendency of a socially-based human intelligence to respond as if in a social dynamic even when faced with the impersonal or random has other consequences as well--particularly when faced with unpatterned events in the world. As Humphrey (1976) noted, "thus the gambler at the roulette table, who continues to bet on the red square precisely because he has already lost on red repeatedly, is behaving as though he expects the behaviour of the roulette wheel to respond eventually to his persistent overtures" (p. 313). Levinson (1995) proposes that the interactional intelligence behind conversational understanding operates with an important default--that the conversational puzzles it is trying to solve were intentionally "<u>designed</u> to be solved and the clues have been designed to be sufficient to yield a determinate solution" (p. 238). This makes it hard to fathom the concept of a "design without a designer" that is at the heart of the concept of evolution (Dawkins, 1986; Dennett, 1995). As Denzau and North (1994) have argued, "it may be an evolutionarily superior survival trait to have explanations for inexplicable phenomena; or this may just be a by-product of the curiosity which helps make humans model builders" (pp. 12-13).

Levinson argues that there are important "spill-over" problems when interactional intelligence--rather than analytic intelligence--is used to decode the structure of the natural world. As a result:

> “we see design in randomness, think we can detect signals from outer space in stellar X-rays, suspect some doodles on archaeological artifacts to constitute an undiscovered code, detect hidden structures in Amazonian myths. If we are attuned to think

> that way, then that is perhaps further evidence for the biases of interactional intelligence: in the interactional arena, we must take all behaviour to be specifically designed to reveal its intentional source" (p. 245).

There is plenty of evidence for the "spill-over" effects that Levinson posits--instances where we carry over assumptions about human design into situations where it is absent. What in particular confounds our quest for structure are situations infused with chance and randomness. The social intelligence hypothesis posits that our voracious search for the purposive meaning behind communicative events--our relentless search for an intentional explanation (see Levinson, 1995)--was an adaptive characteristic that allowed us to predict important events in the social world. But we impose this relentless search for patterns on <u>everything</u> in the world--and sometimes this extremely adaptive aspect of human cognition backfires on us.

The quest for conceptual understanding is maladaptive when it takes place in an environment in which there is nothing to conceptualize. Humphrey (1976), in a passage cited above, mentioned the gambler's fallacy--the tendency for people to see links between events in the past and events in the future when the two are really independent. For example, the number that comes up on a roulette wheel is independent of the outcome that preceded it. Yet after five or six consecutive reds, many bettors switch to black, thinking that it is now more likely to come up. However, the roulette wheel has no memory of what has happened previously. The probability of red and black is still equal.

This pattern--the so-called gambler's fallacy--is not restricted to the inexperienced or novice gambler. Research has shown that even habitual gamblers, who play games of chance over 20 hours a week, still display belief in the gambler's fallacy (Wagenaar, 1988). Also, it is important to realize that the gambler's fallacy is not restricted to games of chance. The gambler's fallacy operates in any domain that has a chance component, such as sporting events and stock markets (see Andreassen, 1987). It operates in any domain in which chance plays a substantial role (that is, in almost everything). The genetic makeup of babies is an example. Psychologists, physicians, and marriage counselors often see couples who, after having two female children, are planning a third child because "We want a boy, and it's bound to be a boy this time." This, of course, is the gambler's fallacy. The probability of having a boy (approximately 50 percent) is exactly the same after having two girls as it was in the beginning.

The tendency to see pattern and design in randomness is especially characteristic of people who interact with the financial markets that play such a large role in modern technological societies. Financial analysts routinely concoct elaborate explanations for every little fluctuation in stock market prices. In fact, much of this variability is simply random fluctuation (Malkiel, 1999; Shefrin & Statman, 1986; Shiller, 1987). Nevertheless, stock market "analysts" continue to imply to their customers (and perhaps believe themselves) that they can "beat the market" when there is voluminous evidence that the vast majority of them can do no such thing. If you had bought all of the 500 stocks in the Standard and Poor's Index and simply held them throughout the 1970s (what we might call a no-brain strategy), you would have had higher returns than 80 percent of the money managers on Wall Street (Malkiel, 1999). If you had done the same thing in the 1980s and 1990s, you would have beaten two-thirds of the money managers on Wall Street (Malkiel, 1999; Updegrave, 1995). You would also have beaten 80 percent of the financial newsletters which subscribers buy at rates of up to $500 per year (Kim, 1994).

The tendency to see design in every random happenstance--especially when the happenstance is a salient event that happens to us--is illustrated in a humorous anecdote related by Dawes (1991). He wrote about meeting an inmate from a prison education program run by the university for which Dawes was working. The inmate had been a bartender in a tavern and had also had a side job collecting gambling debts (at gunpoint if necessary). One day, he was sent to another state, Oregon, to collect a sum of money. When he got to Oregon, he cornered the man owing the money and drew out his gun. Just then, the police jumped from hiding and arrested him on the spot. After telling Dawes the story, the inmate declared confidently that he would never go to jail again. When Dawes asked how he knew that he would never go to jail again, the inmate replied, “Because I’ll never go to Oregon again!”

Now what was wrong with the inmate’s thinking? If we ignore its context and accept it on its own terms, the inmate’s deduction isn’t bad. As Dawes (1991) wrote:

> "People laugh when I tell that story, but viewed in non-probabilistic terms, the enforcer's inference isn't that bad. He had collected debts at gunpoint many times without being arrested, and he had been in Oregon only once. To believe that there was something special about Oregon that 'caused' his arrest is compatible with the canons of inductive inference". (p. 245)

The problem here is not the inference per se, but the focus on trying to explain a particular instance. The inmate was trying to explain why he was arrested in <u>Oregon</u>—why he was arrested that particular time. And it was this focus on the instance--<u>this particular case</u>--that was leading the inmate to an absurd conclusion. Instead, the inmate should have been thinking probabilistically. There was some probability of his being arrested each time he collected debts at gunpoint (perhaps three times in 100). The inmate had done it many times. This time he got arrested. It was simply one of those three times. The odds simply caught up with him. A general trend (such-and-such a probability of being arrested each time he collected a debt at gunpoint) explains why he was arrested. There was probably nothing unique about the fact that it happened to be in Oregon. But instead, the prisoner saw some special "design" specific to Oregon.

Computer scientist Marvin Minsky (1985) has said that "whatever happens, where or when, we're prone to wonder who or what's responsible. This leads us to discover explanations that we might not otherwise imagine, and that help us predict and control not only what happens in the world, but also what happens in our minds. But what if those same tendencies should lead us to imagine things and causes that do not exist? Then we'll invent false gods and superstitions and see their hand in every chance coincidence" (p. 232). Philosopher Daniel Dennett (1995) alludes to this characteristic of thought when discussing the trouble people have in understanding how the process of evolution produces adaptive design through random variation, differential survival, and heredity. A perplexing question that evolution presents us with is "Could something exist for a reason without its being <u>somebody's</u> reason?" (p. 25).

Writer Annie Dillard (1999) alludes to how the tendency to see design often proves embarrassing to the "keenest thinkers" in every religion who are faced with explaining to skeptical rationalists the "presenting face of religion" in the form of its mass superstitions: "In New Mexico in 1978 the face of Jesus arose in a tortilla. 'I was just rolling out my husband's burrito...,' the witness began her account. An auto parts store in Progresso, Texas, attracted crowds when an oil stain on its floor resembled the Virgin Mary. Another virgin appeared in 1998 in Colma, California, in hardened sap on a pine trunk. At a Nashville coffee shop named Bongo Java, a cinnamon bun came out of the oven looking like Mother Teresa--the nun bun, papers called it. In 1996 in Leicester, England, the name of Allah appeared in a halved eggplant" (pp. 76-77). In short, it is not difficult to enumerate many examples of our social intelligence defaulting to the assumption that all fluctuations that we see are due to intentional design. The world that determined a critical substrate of our intelligence was a world populated with other organisms who were transforming the environment with their own plans (see Humphrey, 1976, p. 45). Humphrey (1976) and others (e.g., Levinson, 1995) argue that we evolved to decipher such plans and the social decipherer, the intentional interpreter in our brains (Wolford, Miller, & Gazzaniga, 2000), does not automatically decouple itself from problems when it is not needed. Chance fluctuations will thus routinely confound such a mechanism if it is not decoupled. And of course, chance and randomness infuse many events in our global, interactive modern world.

<u>The Narrative Mode of Thought</u>

The fourth fundamental computational bias is once again a bias that leads to the contextualization of problem solving. It will be mentioned here only briefly because it has been extensively discussed in the literature. This fourth bias is the tendency toward a narrative mode of thought.

Perhaps Bruner's (1986, 1990, 1996) treatments of this mode of thought are the most well-known in psychology, but a number of authors from numerous cross-disciplinary perspectives have reiterated related themes (e.g., Carrithers, 1995; Dennett, 1991; Goody, 1977; Margolis, 1987; Oatley, 1992; 1996, 1999; Schank, 1991). Many characteristics of the narrative mode of thought are closely related to the contextualizing features of the fundamental computational biases discussed previously. For example, people thinking in the narrative mode are biased toward interpreting problems in terms of stories involving agents acting instrumentally to fulfill their goals (Dennett, 1987).

Dennett (1991) has emphasized the role of the narrative in forming our sense of self. In his Multiple Drafts theory of consciousness and its evolution, the self is viewed as a narrative center of gravity--the result of distributed brain centers constantly churning narrative-like language representations in order to aid in behavioral control. Narratives are the main medium by which we present ourselves to ourselves--thus constructing, from our constant storytelling with ourselves as the main protagonist, a sense of self that then becomes the main story we tell to others. Dennett (1991) emphasizes the naturalness and automaticity of the narratives about the self: "Our fundamental tactic of self-protection, self-control, and self-definition is not spinning webs or building dams, but telling stories, and more particularly concocting and controlling the story we tell others--and ourselves--about who we are. And just as spiders don't have to think, consciously and deliberately, about how to spin their webs, and just as beavers, unlike professional human engineers, do not consciously and deliberately plan the structures they build, we (unlike <u>professional</u> human storytellers) do not consciously and deliberately figure out what narratives to tell and how to tell them. Our tales are spun, but for the most part we don't spin them; they spin us. Our human consciousness, and our narrative selfhood, is their product, not their source" (p. 418).

Flanagan (1996) emphasizes the naturalness--and the automatic properties--of our narrative tendencies in his discussion of dreams as epiphenoma of evolutionarily useful tendency// to make sense of stimuli by putting them into narrative structures. He argues that "it did us a lot of good to develop a cortex that makes sense out of experience while awake, and the design is such that there are no costs to this sense-maker always being ready to do its job" (p. 36). The result is the bizarre weaving together of the random cascades of neuronal firing that occur during sleep. The narrative-maker is working as ever--but during sleep is working with degraded and chaotic stimuli. It grinds away nonetheless. The fragmented semi-narrative of dreams thus becomes further evidence of the modes of thought that are our defaults in our waking hours as well.

Perhaps Carrithers' (1995) discussion of the narrative mode has the closest affinity with the social intelligence hypothesis of Humphrey (1976) and others. He notes that stories "have the capacity to frame a markedly intricate and elaborate flow of social events, indeed just the sort of flow that seems even more characteristic of human than of other social primate societies" (p. 261). Carrithers (1995) emphasizes the binding function of narratives--that they tie together many-sided interactions over a considerable period of time. He shares Bruner's (1986) emphasis on the narrative as sequencing the consequences of intentional action, as well as structuring and exploring the consequences of emotional reactions (see Oatley, 1992, 1996, 1999), but emphasizes the role of narrative in solving the evolutionary problem of social coordination (see Levinson, 1995).

Margolis (1987) argues that people have a tendency to fill in contextual information when a problem is not "narrativized" or does not contain a schema that is familiar to the participant: "In the impoverished environment of the set-piece puzzle, therefore, we may impute a wider context to the problem that is not only not there at all but perhaps is flatly inappropriate” (pp. 143-144). The tendency to supplement purely formal problems with prior knowledge often takes the form of a constructed narrative that is largely free of the constraints of information actually presented in the problem[2]. Margolis (1987) has pointed out that it is not uncommon for a participant to concoct a task construal that is so discrepant from anything in the problem as set (even if narratively coherent in itself) that it represents a serious cognitive error: "an anomalous response will almost always in fact be a reasonably logical response to another question (as Henle has claimed), and in particular to a question that means something in the life experience of the individual giving the response. But the other question will often turn out to be a logically irrelevant or absurd interpretation of the context that actually prompted the response" (p. 6).

Many authors emphasizes the "naturalness" of the narrative mode for most real-world situations--and the corresponding unnaturalness of the analytic (scientific and abstract) mode of thinking. For example, Oatley (1996) notes that "one of the properties of the narrative mode is that objects expressed in this mode, that is to say, stories about agents, slip easily into the mind....[in contrast] the mind is more resistant to objects based on the paradigmatic [analytic] mode. At least such objects need elaborate cultural assistance to allow them to enter the mind, for example, knowledge about how to reason mathematically, how to understand statistical data presented in tables and diagrams, or how to draw inferences validly from scientific experiments" (p. 123). Likewise, Carrithers (1995) emphasizes that the narrative mode is not well suited to generalization but works more on the basis of particular to particular. If the former is the focus, then abstract, essayist-styled (Olson, 1977, 1994), scientific thought is the most efficient mode. This contrast between the analytic scientific styles of thought and those reflected in the fundamental computation biases will be the focus of later sections of this chapter.

The Evolutionary Adaptiveness of the Fundamental Computational Biases

Each of the fundamental computational biases discussed previously is a functional aspect of human cognition. Indeed, they are fundamental precisely because they are basic information-processing adaptations that arose in our evolutionary history probablylong before the more abstract features of analytic intelligence (Mithen, 1996). Many investigators have painted compelling theoretically- and empirically-based explanations of why these computational biases developed in the course of human evolution (Cosmides & Tooby, 1992; Humphrey, 1976, 1984; Mithen, 1996; Pinker, 1997). The socialization of problems and the tendency to see deliberate design in undesigned parts of the environment follow from the evolutionary assumptions behind the social intelligence hypothesis--that attributing intentionality in order to predict the behavior of conspecifics and to coordinate behavior with them (see Gibbard, 1990) was a major evolutionary hurdle facing the social primates, in many cases more computationally complex than mastering the physical environment (Humphrey, 1976). The tendency to see design may have other evolutionary sources. For example, Dennett (1991) discusses the "innate tendency to treat every changing thing at first as if it had a soul" (p. 32). He speculates that this tendency is innate because it is an evolutionary design trick that is "a shortcut for helping our time-pressured brains organize and think about the things that need thinking about if we are to survive" (p. 32).

The ubiquitous tendency to adopt what Dennett (1978, 1987) calls the intentional stance underlies many of the fundamental computational biases (particularly the tendency to see human design in the world and to socialize problems). There appear to be biologically-based brain structures devoted to supporting the intentional stance toward other animate beings (Baron-Cohen, 1995). However, these mechanisms do not appear to be modular in Fodor's (1983) strict sense, because they are not informationally encapsulated--they do draw on information from domain-general parts of the brain (Baron-Cohen, 1998; Thomas & Karmiloff-Smith, 1998; Tsimpli & Smith, 1998). Instead, they reflect a processing style that infuses much of cognition because they were early developing aspects of intelligence. Evolutionarily later aspects of analytic cognition (see Dennett, 1991, 1996; Mithen, 1996) did not replace these older socially-based mechanisms but were built on top of them. Thus, aspects of social intelligence infuse even abstract problems that are best solved with later developing (see Mithen, 1996; Reber, 1992a, 1992b) analytic intelligence.

Finally, there exist many theoretical arguments for why the automatic contextualization of problems with prior knowledge might be adaptive (see Stanovich, 1999). For example, Evans and Over (1996) provide arguments in favor of the adaptiveness of contextualization and the nonoptimality of always decoupling prior beliefs from problem situations ("beliefs that have served us well are not lightly to be abandoned," p. 114).

Their argument parallels the reasons that philosophy of science has moved beyond naive falsificationism (see Howson & Urbach, 1993). Scientists do not abandon a richly confirmed and well integrated theory at the first little bit of falsifying evidence, because abandoning the theory might actually decrease explanatory coherence (Thagard, 1992). Similarly, Evans and Over (1996) argue that beliefs that have served us well in the past should be hard to dislodge, and projecting them on to new information--because of their past efficacy--might help in assimilating the new information.

This argument for the adaptiveness of contextualization was termed the knowledge projection argument by Stanovich (1999) because it reappears in a remarkably diverse set of disciplines and specialties within cognitive science. For example, philosopher Hilary Kornblith (1993) stresses that "mistaken beliefs will, as a result of belief perseverance, taint our perception of new data. By the same token, however, belief perseverance will serve to color our perception of new data when our preexisting beliefs are accurate....If, overall, our belief-generating mechanisms give us a fairly accurate picture of the world, then the phenomenon of belief perseverance may do more to inform our understanding than it does to distort it" (p. 105).

This argument--that in a natural ecology where most of our prior beliefs are true, projecting our beliefs on to new data will lead to faster accumulation of knowledge--has been used to explain the false consensus effect in social psychology (Dawes, 1989, 1990; Hoch, 1987; Krueger & Zeiger, 1993), findings on expectancy effects in learning (Alloy & Tabachnik, 1984), biases in evaluating scientific evidence (Koehler, 1993), realism effects in attributing intentionality (Mitchell, Robinson, Isaacs, & Nye, 1996), syllogistic reasoning (Evans, Over, & Manktelow, 1993), and informal reasoning (Edwards & Smith, 1996). Alloy and Tabachnik (1984) summarize the generic case for knowledge projection, arguing for the general adaptiveness of such contextualization because "covariation information provided in an experiment may represent only one piece of conflicting evidence against the background of the large body of data about event covariations summarized by an expectation, it would be normatively appropriate for organisms to weight their expectations more heavily than situational information in the covariation judgment process" (p. 140). Of course, Alloy and Tabachnik (1984) emphasize that we must project from a largely accurate set of beliefs in order to obtain the benefit of knowledge projection. In a sea of inaccurate beliefs, the situation is quite different (see Stanovich, 1999, Chapter 8 for a discussion).

Facilitating Reasoning by Fitting Problems to the Fundamental Computational Biases

Evolutionary psychologists have shown that many of the problems that are difficult for people in their abstract forms can be made easier to solve if they are contextualized--particularly if they are contextualized in ways that are compatible with the representations used by specific evolutionarily adapted modules (see Cosmides & Tooby, 1992, 1996; Gigerenzer & Hoffrage, 1995). The most famous demonstration involves Wason's (1966) four-card selection task described above. The abstract rule (if there is a vowel on one side of the card, then there is an even number on the other side) is notoriously difficult, and this has been known for some time (Evans, Newstead, & Byrne, 1993). At first it was thought that the abstract content of the vowel/number rule made the problem hard for people and that more real-life or "thematic," nonabstract problems would raise performance markedly. Investigators tried examples like the following "Destination Rule": "If 'Baltimore' is on one side of the ticket, then 'plane' is on the other side of the ticket." The fours cards facing the participant said:

Destination: **Baltimore**	Destination: **Washington**	Mode of Travel: **Plane**	Mode of Travel: **Train**

Surprisingly, this type of content did not improve performance at all (Cummins, 1996; Manktelow & Evans, 1979; Newstead & Evans, 1995; Stanovich & West, 1998a). Most participants still picked either P and Q or the P card only. The correct P, not-Q solution escaped the vast majority.

However, Griggs and Cox (1982) were the first to use a thematic version of the task that did markedly improve performance in their experiment and in many subsequent experiments by other investigators (Cummins, 1996; Dominowski, 1995; Newstead & Evans, 1995; Pollard & Evans, 1987). Here is a particularly easy version of the Griggs and Cox (1982) rule used by my research group (Stanovich & West, 1998a). Do the problem and experience for yourself how easy it is:

> Imagine that you are a police officer on duty, walking through a local bar. It is your job to ensure that the drinking laws are in effect in this bar. When you see a person engaging in certain activities, the laws specify that certain conditions must first be met. One such law is "If a person is drinking beer then the person must be over 21 years of age." Each of the boxes below represents a card lying on a table. There are two pieces of information about a person on each card. Whether or not the person is drinking beer is on one side of the card and the person's age is on the other side. For two of the people, you can see their age, but you cannot see what they are drinking. For

the other two people, you can see what they are drinking, but you cannot see their age. Your task is to decide whether or not this law is being broken in the bar. Circle the card or cards you would definitely need to turn over to decide whether or not the law is being broken. You may select any or all of the cards.

Age: **22**	Age: **18**	Drink: **Beer**	Drink: **Coke**

Many people answer the Drinking-Age problem correctly, including many people who answer the abstract version incorrectly. This is true even though the underlying logical structure of the two problems are seemingly the same. The answer to both is to pick P and not-Q--in this problem, Beer and Age 18.

With the invention of the Drinking-Age problem, researchers had finally found a way to get participants to give the right answer to the Wason four-card selection task after 15 years of research (at the time of the Griggs & Cox, 1982 report). Joy was short-lived however, because researchers immediately began to doubt whether the reasoning process leading to correct responses on the abstract version was anything like the reasoning process leading to correct responding on the Drinking-Age version. That is, despite the surface similarity of the two rules, investigators began to think that they were actually tapping fundamentally different reasoning mechanisms. The Destination rule, for example, is what is called an indicative rule--a rule concerning the truth status of a statement about the world. In contrast, the Drinking-Age rule is a so-called deontic rule. Deontic reasoning problem solving concerns thinking about the rules used to guide human behavior--about what “ought to” or "must" be done. Cummins (1996) terms a deontic rule a rule about "what one may, ought, or must not do in a given set of circumstances" (p. 161; see also, Manktelow, 1999; Manktelow & Over, 1991). A number of theorists have argued that deontic rules and indicative rules engage different types of mental mechanisms.

The most famous of these proposals was in a highly influential paper by Cosmides (1989)--one of the leading figures in the move to ground psychology in evolutionary theory that swept through psychology in the 1990s (for summaries, see Buss, 1999; Cosmides & Tooby, 1992, 1994, 1996; Geary & Bjorklund, 2000; Pinker, 1997; Plotkin, 1998). She proposed that evolution has built processing systems (what she termed Darwinian algorithms) exclusively concerned with social exchange in human interactions. These algorithms embody the basic rule "if you take a benefit, then you must pay a cost" and are extremely sensitive "cheater detectors"--they react strongly to instances where an individual takes a benefit without paying the cost. In the Drinking-Age problem, an individual underage and drinking beer is just that--a cheater. Thus, with this rule, the possibility of an 18-year-old drinking beer (the P and not-Q case) becomes very salient because the rule automatically triggers an evolutionary algorithm specifically concerned with detecting card selections that happen to be correct. The indicative rule of course does not trigger such a Darwinian algorithm. Evolution has provided no special module in the brain for solving indicative problems. The tools for solving such problems are largely cultural inventions (Dennett, 1991; Jepson, Krantz, & Nisbett, 1983; Krantz, 1981; McCain, 1991; Stanovich, 1999; Thagard & Nisbett, 1983) and the brain processes supporting them are fragile because they demand much computational capacity (Dennett, 1991;

Evans & Over, 1996, 1999; Stanovich & West, 2000).

Cosmides' (1989) hypothesis is not the only explanation for the superior performance in the Drinking-Age problem (see Cheng & Holyoak, 1985; Evans & Over, 1996; Kirby, 1994; Manktelow, 1999; Manktelow & Over, 1991; Manktelow, Sutherland, & Over, 1995; Oaksford & Chater, 1994, 1996). Other theorists have taken issue with her emphasis on a domain-specific and informationally encapsulated modules for regulating social exchange. However, the alternative explanations do share some family resemblances with Cosmides' (1989) account. They tend to view the Drinking-Age task as a problem of pragmatic rationality rather than epistemic rationality. Indicative selection tasks tap epistemic rationality--they probe how people test hypotheses about the nature of the world. In contrast, deontic tasks tap pragmatic rationality--they concern how actions should be regulated and what people should <u>do</u> in certain situations. Given the arguments of the social intelligence theorists, we again might expect the mechanisms for the the former to be evolutionarily younger (see Mithen, 1996; Reber, 1992a, 1992b) and more computationally cumbersome (Dennett, 1991; Evans & Over, 1996, 1999; Stanovich & West, 2000). If participants do not use the computationally expensive processes of analytic intelligence (Dennett, 1991; Stanovich, 1999) to solve the indicative task, then they must rely on automatic heuristics such as Evans' (1996) relevance heuristic. This heuristic generates a primitve, so-called "matching" response of P and Q--the participant is primed to choose just the two cards mentioned in the rule. The problem is that, in the indicative task, these heuristics do not lead to the correct response.

Consistent with this interpretation are some individual difference data from my own lab (Stanovich & West, 1998a). Richard West and I have found that there were large intelligence differences in abstract selection task problems. The minority of individuals who answered correctly had significantly higher SAT scores than those that did not, and the difference was quite large in magnitude (effect sizes of roughly .500 to .800). In contrast, cognitive ability differences between those who answered deontic problems such as the Drinking-Age problem correctly and incorrectly are considerably smaller (effect sizes of roughly .050 to .400).

A similar pattern is apparent on other tasks where the fundamental computational biases lead to overall poor performance on noncontextualized task versions. Take the conjunction fallacy illustrated above (the Linda problem). Evolutionary psychologists have argued that the human cognitive apparatus is more adapted to dealing with frequencies than with probabilities (Brase, Cosmides, & Tooby, 1998; Cosmides & Tooby, 1996; Gigerenzer & Hoffrage, 1995). It has been found that when tasks such as the Linda problem are revised in terms of estimating the frequency of categories rather than judging probabilities that performance is improved (see Fieldler, 1988; Gigerenzer, 1991b, 1993). My research group has replicated this finding, but we again found that cognitive ability differences are much smaller in the frequentist versions of these problems (Stanovich & West, 1998b).

Stanovich (1999; Stanovich & West, 2000) reports other instances of this trend. Specifically, the abstract, noncontextualized versions of many problem solving and reasoning problems usually produces large cognitive ability differences. Versions of many of these problems designed with considerations of evolutionary psychology in mind have indeed produced vastly superior performance overall. However, these same versions often attenuate individual differences in analytic intelligence. I have argued (Stanovich, 1999) that these findings are reconciled by clearly distinguishing evolutionary adaptation from normative (or individual) rationality. In this distinction lies a possible rapprochement between the researchers who have emphasized the flaws in human cognition (Kahneman & Tversky, 1973, 1996, 2000) and the evolutionary psychologists who have emphasized the optimality of human cognition (Cosmides & Tooby, 1994, 1996; Gigerenzer & Todd, 1999). What is useful here is to use Dawkins' (1976) replicator/vehicle terminology to distinguish between evolutionary adaptation at the level of the gene and instrumental rationality (utility maximization given goals and beliefs) at the level of the individual person.

Distinguishing optimization at the level of the replicator from optimization at the level of the vehicle can reconcile both the impressive record of descriptive accuracy enjoyed by a variety of adaptationist models (Anderson, 1990, 1991; Oaksford & Chater, 1994, 1996, 1998) with the fact that cognitive ability sometimes dissociates from the response deemed optimal on an adaptationist analysis (Stanovich & West, 2000). For example, Oaksford and Chater (1994) have had considerable success in modeling the abstract selection task as an inductive problem in which optimal data selection is assumed (see also, Oaksford, Chater, Grainger, & Larkin, 1997). Their model predicts the modal response of P and Q and the corresponding dearth of P and not-Q choosers. Similarly, Anderson (1990, p. 157-160) models the 2 x 2 contingency assessment experiment using a model of optimally adapted information processing and shows how it can predict the much-replicated finding that the D cell (cause absent and effect absent) is vastly underweighted (see also Friedrich, 1993;

Klayman & Ha, 1987). Finally, a host of investigators (Adler, 1984, 1991; Dulany & Hilton, 1991; Hilton, 1995; Levinson, 1995; Hilton & Slugoski, 2000) have stressed how a model of rational conversational implicature predicts that violating the conjunction rule in the Linda Problem reflects the adaptive properties of interactional intelligence.

Yet in all three of these cases--despite the fact that the adaptationist models predict the modal response quite well--individual differences analyses demonstrate associations that also must be accounted for. Correct responders on the abstract selection task (P and not-Q choosers--not those choosing P and Q) are higher in cognitive ability. In the 2 x 2 covariation detection experiment, it is those participants weighting cell D more equally (not those underweighting the cell in the way that the adaptationist model dictates) who are higher in cognitive ability (Stanovich & West, 1998d). Finally, despite conversational implicatures indicating the opposite, individuals of higher cognitive ability disproportionately tend to adhere to the conjunction rule. These patterns make sense if it is assumed: 1) that there are two systems of processing--sometimes labelled heuristic and analytic (see Evans & Over, 1996; Sloman, 1996; Stanovich, 1999); 2) that the two systems of processing are optimized for different situations and different goals; 3) that in individuals of higher cognitive ability there is a greater probability that the analytic system will override the heuristic system.

This differential proclivity for override could become important in situations where the two systems compute different responses. Stanovich and West (2000; Stanovich, 1999) have argued that this is more likely to happen in situations where the human cognitive system has instantiated conflicting goals. Differing goals might characterize different levels of computational complexity in a "tower of intellect" model of the type popularized by Dennett (1991, 1995, 1996) and illustrated by the title of his book Kinds of Minds. In such a conception, more computationally complex cognitive structures do not replace simpler ones, but are posited to operate in parallel. Specifically, what we must be concerned about are situations where the evolutionary adapted goals (instantiated in the evolutionarily older heuristic mechanisms) do not coincide with personal goals in the current environment which are more likely to be tracked by systems displaying a more flexible, analytic intelligence (Sternberg, 1997).

In short, the critical situations are those where the the interests of the replicators and the vehicle do not coincide (again, to use Dawkins' 1976, terms). Evolutionary psychologists are prone to emphasize situations where genetic goals and personal goals coincide. They are not wrong to do so because this is often the case. Accurately navigating around objects in the natural world was adaptive during the EEA, and it similarly serves our personal goals as we carry out our lives in the modern world. Likewise, with other evolutionary adaptations: It is a marvel that humans are exquisite frequency detectors (Hasher & Zacks, 1979), that they infer intentionality with almost supernatural ease (Levinson, 1995), and that they acquire a complex language code from impoverished input (Pinker, 1994)--and all of these mechanisms served personal goal fulfillment in the modern world. But none of this means that the overlap is necessarily one-hundred percent.

First, evolutionary biology is full of examples where the genes' instantiate strategies that necessitate sacrificing the vehicle. Dawkins' book (1982) The Extended Phenotype contains many such examples. For instance, the genes have little interest in the vehicle they build once it is beyond its reproductive years--which is why many creatures (like salmon) die immediately after reproducing. But as humans, we <u>are</u> interested in our postreproductive longevity--so we have a clear example of genetic goals and individual goals coming apart. Skyrms (1996) devotes an entire book on evolutionary game theory to showing that instrumental goal optimization for an individual organism might not coincide with adaptive fitness. He concludes that "if evolutionary game theory is generalized to allow for correlation of encounters between players and like-minded players, then strongly dominated strategies--at variance with both rational decision and game theory--can take over the population....When I contrast the results of the evolutionary account with those of rational decision theory, I am not criticizing the normative force of the latter. I am just emphasizing the fact that the different questions asked by the two traditions may have different answers" (pp. x-xi). Skyrms' (1996) book articulates the environmental and population parameters under which "rational choice theory completely parts ways with evolutionary theory" (p. 106; see also Cooper, 1989). The point is that local maximization in the sense of genetic fitness is not the same as the maximization of expected utility for the individual.

Unfortunately, the modern world tends to create situations where some of the default values of evolutionarily adapted cognitive systems are not optimal. Many of these situations implicate the fundamental computational biases discussed previously. These biases serve to radically contextualize problem-solving situations. In contrast, modern technological societies continually spawn situations where humans must decontextualize information--where they must deal abstractly (Adler, 1984) and in a depersonalized manner with information. Such situations require the active suppression of the personalizing and contextualizing styles that characterize the fundamental computational biases. These situations may not be numerous, but they tend to be in particularly important domains of modern life--indeed, they in part define modern life in postindustrial knowledge-based societies.

The Fundamental Computational Biases and the Problems of Modern Society

Mechanisms designed for survival in preindustrial times are clearly sometimes maladaptive in a technological culture. Our mechanisms for storing and utilizing energy evolved in times when fat preservation was efficacious. These mechanisms no longer serve the goals of people in a technological society where a Burger King is on every corner. Likewise, the cognitive mechanisms that lead us to stray from normative models that would maximize utility are probably mechanisms that once were fitness enhancing but now serve to thwart our goals (see Baron, 1993, 1994, 1998; Stanovich, 1999).

It will be argued here that many of the fundamental computational biases are now playing this role. Such biases directly conflict with the demands for decontextualization that a highly bureaucratized society puts on its citizens. Indeed, this is often why schools have to explicitly teach such skills of cognitive decontextualization. Donaldson (1978, 1993) views this as one of the primary tasks of schooling (see Anderson, Reder, & Simon, 1996). She argues that “what is involved in the mind’s movement from ‘seven fishes’ to ‘seven’ is abstraction indeed, but it is more: it is a dramatic decontextualization. In the contexts of our ordinary life we have to deal with quantities of fishes but we never encounter seven” (p. 90). She emphasizes how, in order to master a variety of abstract rule systems (mathematics, logic, etc.), decontextualization must become a comfortable thinking style for a learner: “If the intellectual powers are to develop, the child must gain a measure of control over his own thinking and he cannot control it while he remains unaware of it. The attaining of this control means prising thought out of its primitive unconscious embeddedness in the immediacies of living in the world and interacting with other human beings. It means learning to move beyond the bounds of human sense. It is on this movement that all the higher intellectual skills depend” (Donaldson, 1978, p. 123).

This point is a recurring theme in the literature of cognitive development (e.g., Neimark, 1987; Piaget 1972; Sigel, 1993). Indeed, many developmental theorists, as have influential cognitive psychologists (e.g., Kahneman & Tversky, 1982), emphasize how schooling teaches children to decouple reasoning from the pragmatic inferences of language comprehension. For example, in a paper discussing developmental trends in reasoning, Chapman (1993) draws the specific link between the ability to decouple pragmatic knowledge in the interests of reasoning logically: "children who have mastered the pragmatic rules of language may have to unlearn some of these rules when it comes to formal reasoning. More precisely, they may have to learn that particular contexts exist in which those rules do not apply" (p. 104).

Of course, all of this emphasis on decoupling pragmatic processes of natural language understanding is quite unnatural--unnatural in the sense that it is evolutionarily unprecedented and that it requires overriding many cognitive heuristics that are probably highly automatized (Pollock, 1995; Stanovich, 1999). But it is not just for success in school that we inculcate the decontextualization skills emphasized by Donaldson (1978, 1993). Increasingly, modern society is demanding such skills (Frank & Cook, 1995; Gottfredson, 1997; Hunt, 1995, 1999)--and in some cases it is rendering economically superfluous anyone who does not have them (Bronfenbrenner, McClelland, Wethington, Moen, & Ceci, 1996; Frank & Cook, 1995).

Modern society creates many situations that require radical decontextualization--that require one or more of the fundamental computational biases to be overridden by analytic intelligence. For example, many aspects of the contemporary legal system put a premium on detaching prior belief and world knowledge from the process of evidence evaluation. There has been understandable vexation at odd jury verdicts rendered because of jury theories and narratives concocted during deliberations that had nothing to do with the evidence but instead that were based on background knowledge and personal experience. For example, a Baltimore jury acquitted a murder defendant who had been identified by four witnesses and had confessed to two people because "they had invented their own highly speculative theory of the crime" (Gordon, 1997, p. 258). In this particular case, the perpetrator had wanted to plea bargain for a 40-year sentence, but this was turned down at the request of the victim's family. Similarly, in Lefkowitz' (1997) account of the trial of several teenagers in an affluent New Jersey suburb who brutally exploited and raped a young girl who was intellectually disabled, one juror concocted the extenuating circumstance that one defendant thought he was attending an "initiation rite" even though no evidence for such a "rite" had been presented in months of testimony.

The point is that in a particular cultural situation where detachment and decoupling is required, the people who must carry out these demands for decontextualization are often unable to do so even under legal compulsion. Post-trial reports of juries in a "creative," "narrative," or highly enthymematic mode have incited great debate. If the polls are to be believed, a large proportion of Americans were incensed at the jury's acquittal of O. J. Simpson. Similar numbers were appalled at the jury verdict in the first trial of the officers involved in the Rodney King beating. What both juries failed to do was to decontextualize the evidence in their respective cases--and each earned the wrath of their fellow citizens because it is a cultural (and legal) expectation of citizenship that people should be able to carry out this cognitive operation in certain settings.

The need to decontextualize also characterizes many work settings in contemporary society. Consider the common admonition in the retail service sector of "the customer is always right". This admonition is often interpreted to include even instances where customers unleash unwarranted verbal assaults which are astonishingly vitriolic. The service worker is supposed to remain polite and helpful under this onslaught, despite the fact that such emotional social stimuli are no doubt triggering evolutionarily instantiated modules of self defense and emotional reaction. All of this emotion, all of these personalized attributions--all fundamental computational biases--must be set aside by the service worker and instead an abstract rule that "the customer is always right" must be invoked in this special, socially-constructed domain of the market-based transaction. The worker must realize that he/she is not in an <u>actual</u> social interaction with this person (which if true, might call for socking them in the nose!), but in a special, indeed "unnatural" realm where different rules apply.

Numerous theorists have warned about a possible mismatch between the fundamental computational biases and the processing requirements of many tasks in a technological society containing many abstract and decontextualized symbolic artifacts. Hilton (1995) warns that the default assumption of interactional intelligence may be wrong for many technical settings because "many reasoning heuristics may have evolved because they are adaptive in contexts of social interaction. For example, the expectation that errors of interpretation will be quickly repaired may be correct when we are interacting with a human being but incorrect when managing a complex system such as an aircraft, a nuclear power plant, or an economy. The evolutionary adaptiveness of such an expectation to a conversational setting may explain why people are so bad at dealing with lagged feedback in other settings" (p. 267).

Concerns about the real-world implications of the failure to engage in necessary cognitive abstraction (see Adler, 1984) were what led Luria (1976) to warn against minimizing the importance of decontextualizing thinking styles. In discussing the syllogism, he notes that "a considerable proportion of our intellectual operations involve such verbal and logical systems; they comprise the basic network of codes along which the connections in discursive human thought are channeled" (p. 101). Likewise, regarding the subtle distinctions on many decontextualized language tasks, Olson (1986) has argued that "the distinctions on which such questions are based are extremely important to many forms of intellectual activity in a literate, society. It is easy to show that sensitivity to the subtleties of language are crucial to some undertakings. A person who does not clearly see the difference between an expression of intention and a promise or between a mistake and an accident, or between a falsehood and a lie, should avoid a legal career or, for that matter, a theological one" (p. 341).

Olson's statement reflects a stark fact about modern technological societies--they are providing lucrative employment only for those who can master complexity, make subtle quantitative and verbal distinctions, and reason in decontextualized ways (Bronfenbrenner, McClelland, Wethington, Moen, & Ceci, 1996; Frank & Cook, 1995; Gottfredson, 1997; Hunt, 1995, 1999). Objective measures of the requirements for cognitive abstraction have been increasing across most job categories in technological societies throughout the past several decades (Gottfredson, 1997). This is why measures of the ability to deal with abstraction remains the best employment predictor and the best earnings predictor in postindustrial societies (Brody, 1997; Gottfredson, 1997; Hunt, 1995).

Adler (1991) emphasizes the point that not to make important linguistic, probabilistic, and logical distinctions in a complex social environment has real costs and represents more than just the failure to play an artificial game: "The conversationally induced problem of a lack of shared understanding is a subtle one, not due to any blatant verbal trick. It is reasonable to conjecture that the subtlety results in part from participants limited skill with the rule the experimenter wants to study. To be specific: Our greater pragmatic sophistication alone does not explain the differences in the dominant adult responses to the conjunction effect compared to the Piaget class-inclusion studies. The difference between the number of members of a class and the number of a proper sub-class is so obvious to us that we readily permit the conversationally untoward question--'Are there more dimes or coins?'--at face value. Our greater resistance to the violation of the maxims in the conjunction-effect experiment is partly due, I believe, to a certain lack of either accessibility to or confidence in--though not competence with--the conjunction rule for probabilities. If this is so, then the fact that subjects do not understand the experimenter as he intends his words is itself some evidence of a weakness in subjects understanding of the scope of the conjunction rule in everyday reasoning" (p. 265).

Einhorn and Hogarth (1981) highlight the importance of decontextualized environments in their discussion of the optimistic and pessimistic views of the cognitive biases revealed in laboratory experimentation. Einhorn and Hogarth (1981) note that "the most optimistic asserts that biases are limited to laboratory situations which are unrepresentative of the natural ecology" (p. 82), but they go on to caution that "in a rapidly changing world it is unclear what the relevant natural ecology will be. Thus, although the laboratory may be an unfamiliar environment, lack of ability to perform well in unfamiliar situations takes on added importance" (p. 82).

Critics of the abstract content of most laboratory tasks and standardized tests have been misguided on this very point. The issue is that, ironically, the argument that the laboratory tasks and tests are not like "real life" is becoming less and less true. "Life," in fact, is becoming more like the tests! Try using an international ATM machine with which you are unfamiliar; or try arguing with your HMO about a disallowed medical procedure. In such circumstances, we invariably find out that our personal experience, our emotional responses, our stimulus-triggered intuitions about social justice--all are worthless. All are for naught when talking over the phone to the representative looking at a computer screen displaying a spreadsheet with a hierarchy of branching choices and conditions to be fulfilled. The social context, the idiosyncrasies of individual experience, the personal narrative--all are abstracted away as the representatives of modernist technological-based services attempt to "apply the rules". Consider Toronto writer Todd Mercer (2000) trying to fly across the continent on short notice to be with his 83-year-old father undergoing emergency surgery. Calling Canadian Airlines and finding out that the last-minute scheduled airline fare was $3120, Mercer asked if there was any discount that applied to his situation and was informed that he might be eligible for an "imminent death discount" by, as Mercer puts it "no less a medical/spiritual authority" than the telephone ticket agent. Prodded for the definition of "imminent death" the ticket agent quotes from a document outlining the details of the "bereavement travel program" which clarifies the program's requirements when illness rather than death is the reason for the travel. The ticket agent relates that the person in question must be a patient in intensive care, a patient in the final stages of cancer, or a patient involved in a serious accident. Mercer's father had an aortic aneurysm which made him a "walking time bomb" according to his doctor, but he had not yet gone into surgery and had not yet been put into intensive care. The ruling was that such a situation was in "a gray area" and, as a result, the ticket agent stonewalled by saying "not all operations are life threatening. The imminent death discount is not meant just for operations. It is meant for imminent death"--the latter defined as above, and another round of technical and nuanced argument between Mercer and the ticket agent ensued. This is life in the First World in the early part of the twenty-first century.

The abstract, semantic games encountered by Mercer are nothing compared to what a person faces when deciding on whether to apply for a tax deduction for an infirm relative who lived outside Canada for the year 1994. Canada Customs and Revenue Agency will advise the person that "Your dependent must be: -- your or your spouse's child or grandchild, if that child was born in 1976 or earlier and is physically or mentally infirm; or -- a person living in Canada at any time in the year who meets all of the following conditions. The person must have been: -- your or your spouse's parent, grandparent, brother, sister, aunt, uncle, niece, or nephew; -- born in 1976 or earlier; and -- physically or mentally infirm." Given the ubiquitousness of such abstract directives in our informational and technology-saturated society, it just seems perverse to argue the "unnaturalness" of decontextualized reasoning skills when such skills are absolutely necessary in order to succeed in our society. If one has the post-industrial goal of, say, "going to Princeton," then the only way to fulfill that goal in our current society <u>is</u> to develop such cognitive skills. Situations that require abstract thought and/or the ability to deal with complexity will increase in number as more niches in post-industrial societies require these intellectual styles and skills (Gottfredson, 1997; Hunt, 1995). For intellectuals to use their abstract reasoning skills to argue that the "person in the street" is in no need of such skills of abstraction is like a rich person telling someone in poverty that money is not really all that important.

To the extent that modern society increasingly requires the fundamental computational biases to be overridden, then dissociations between evolutionary and individual rationality will become more common. Cosmides and Tooby (1996) argue that "in the modern world, we are awash in numerically expressed statistical information. But our hominid ancestors did not have access to the modern accumulation which has produced, for the first time in human history, reliable, numerically expressed statistical information about the world beyond individual experience. Reliable numerical statements about single event probabilities were rare or nonexistent in the Pleistocene" (p. 15). "It is easy to forget that our hominid ancestors did not have access to the modern system of socially organized data collection, error checking, and information accumulation....In ancestral environments, the only external database available from which to reason inductively was one's own observations" (Brase, Cosmides, & Tooby, 1998, p. 5).

Precisely. I am living in a technological society where I must: decide which HMO to join based on just such statistics; figure out whether to invest in a Roth IRA; decide what type of mortgage to purchase; figure out what type of deductible to get on my auto insurance; decide whether to trade in a car or sell it myself; decide whether to lease or to buy; think about how to apportion my TIAA/CREF retirement funds; and decide whether I would save money by joining a book club. And I must make all of these decisions based on information represented in a manner for which my brain is not adapted (in none of these cases have I coded individual frequency information from my own personal experience). In order to reason normatively in all of these domains (in order to maximize my own personal utility) I am going to have to deal with probabilistic information represented in nonfrequentistic terms--in representations that the evolutionary psychologists have shown are different from my well-adapted algorithms for dealing with frequency information (Cosmides, & Tooby, 1996; Gigerenzer & Hoffrage, 1995).

Consider the work of Brase et al. (1998), who improved performance on a difficult probability problem (Bar-Hillel & Falk, 1982; Falk, 1992; Granberg, 1995) by presenting the information as frequencies and in terms of whole objects--both alterations designed to better fit the frequency-computation systems of the brain. In response to a query about why the adequate performance observed was not even higher given that our brains contain such well-designed frequency-computation systems, Brase et al. (1998) replied that "in our view it is remarkable that they work on paper-and pencil problems at all. A natural sampling system is designed to operate on actual events" (p. 13). The problem is that in a symbol-oriented postindustrial society, we are presented with paper-and pencil problems all the time, and much of what we know about the world comes not from the perception of actual events but from abstract information preprocessed, prepackaged, and condensed into symbolic codes such as probabilities, percentages, tables, and graphs (the voluminous statistical information routinely presented in USA Today comes to mind).

What we are attempting to combat here is a connotation implicit in some discussions of findings in evolutionary psychology (e.g., Gigerenzer & Todd, 1999) and indeed in the situated cognition literature as well (see Anderson et al., 1996) that there is nothing to be gained from being able to understand a formal rule at an abstract level (the conjunction rule of probability, etc.)--and no advantage in flexibly overriding the fundamental computational biases. To the contrary, modern technological society often puts a premium on the use of such abstract tools. Adler (1984) emphasizes that the efficiency of the cognitive modules underlying intention attribution and social cooperation just as surely extract certain costs. He argues that such costs occur "in situations where the naturalness of our expectations under the cooperative principle leads us to miss subtle, but significant deviations from those expectations" (p. 174). He cites a now-famous experiment by Langer, Blank, and Chanowitz (1978) where a confederate attempts to cut into a line at a copy machine. In one condition a good reason is given ("May I use the Xerox machine, because I'm in a rush?") and in the other a totally redundant nonexplanation ("May I use the Xerox machine, because I have to make copies?"). Despite the fact that the second explanation is much less informative than the first, the compliance rates in the two conditions did not differ. Langer (1989; Langer et al., 1978) terms the compliance in the second case "mindless" and Adler (1984) analyzes it in terms of overly generalizing the Gricean Cooperative Principle. The default assumption that a contribution will be selectively relevant--in this case, that a real reason will follow the request--is false in this condition, yet it triggers exactly the same compliance behavior because it is not overridden ("Yes but all of us are in line to make copies. Why should you go first?").

Langer-type examples of mindlessness abound in many important domains. Consumer Reports (April, 1998) chronicles how some dealers put an item costing $500 and labeled ADM on many automobile price stickers. The dealers are hoping that some people will not ask what ADM means. The dealers are also hoping that even after asking and being told that it means "additional dealer markup" that some consumers will not fully process what that means and will not inquire further about what this additional dealer markup feature is that they are paying for. In short, the dealers are hoping that analytic intelligence will not override Langer-type "mindlessness" and allow the customer to ascertain that ADM is not a feature on the car at all--that it simply represents a request from the dealer to contribute $500 more to the dealership, as if it were a charity. As one dealer put it, "every once in a while somebody pays it, no questions asked" (p. 17). A mindless response here, a failure to override automatic heuristics, and the consumer could simply throw away a good chunk of hard-earned income. The modern consumer world is simply littered with such traps and, often, the more costly the product, the more such traps there are (e.g., automobiles, mutual funds, mortgage closing costs).

Modern mass communication technicians have become quite skilled at implying certain conclusions without actually stating those conclusions (for fear of lawsuits, bad publicity, etc.). Advertisements rely on the fundamental computational biases (particularly its enthymematic processing feature) to fill in the missing information. Of course, such techniques are not limited to election and product advertising. The glossy brochures that our universities send out, full of fresh-faced young people in a wholesome learning environment, have the same logic. Margolis (1987; see Margolis, 1996) warns of the ubiquitousness of this situation in modern society: "We can encounter cases where the issue is both out-of-scale with everyday life experience and contains important novelties, so that habitual responses can be highly inappropriate responses. The opportunity for unrecognized contextual effects akin to the scenario effects...[demonstrated in the laboratory] can be something much more than an odd quirk that shows up in some contrived situation" (p. 168).

As discussed previously, evolutionary psychologists have shown that some problems can be efficiently solved if represented one way (to coincide with how the brain modules represent information) but not if represented in another way (for example, as single-event probabilities rather than frequencies--see Gigerenzer, Hoffrage, & Kleinbolting, 1991). Nevertheless, they often seem to ignore the fact that the world will not always let us deal with representations that are optimally suited to our evolutionarily designed cognitive mechanisms. For example, in a series of elegant experiments, Gigerenzer et al. (1991) have shown how at least part of the overconfidence effect in knowledge calibration studies is due to the unrepresentative stimuli used in such experiments--stimuli that do not match the participants' stored cue validities which are optimally tuned to the environment. But there are many instances in real-life when we are suddenly placed in environments where the cue validities have changed. Metacognitive awareness of such situations and strategies for suppressing incorrect confidence judgments generated by automatic responses to cues will be crucial here. Every high school musician who aspires to a career in music has to recalibrate when they arrive at university and see large numbers of talented musicians for the first time. If they persist in their old confidence judgments they may not change majors when they should. Many real-life situations where accomplishment yields a new environment with even more stringent performance requirements share this logic. Each time we "ratchet up" in the competitive environment of a capitalist economy (Frank & Cook, 1995) we are in a situation just like the overconfidence knowledge calibration experiments with their unrepresentative materials. It is important to have learned strategies that will temper one's overconfidence in such situations (Koriat, Lichtenstein, & Fischhoff, 1980).

<u>Abstraction and the Fundamental Computational Biases in Education</u>

In an article titled "Abstraction is Uncooperative," Adler (1984) points out that what is called for in many problem solving and reasoning tasks (and certainly in many tasks in the heuristics and biases literature) is abstraction--extracting from a given problem representation only the features that fit a general pattern. In the process of abstraction "we are rendering information inessential to the formal structure irrelevant" (p. 165). But Adler (1984) points out that the Gricean Cooperative Principle is directly in conflict with the demands of abstraction. The cooperativeness principle that everything about the experimenter's contribution is relevant (the instructions, the context, every bit of content that is presented) is diametrically opposed to the requirements of abstraction--that we treat as inessential everything that does not fit a certain formal structure.

As Donaldson (1978, 1993) argues, education serves, both explicitly and implicitly, to suppress Gricean tendencies in order to make adopting an abstract stance a more natural processing style. This is why education--at all levels--is often, deliberately and correctly, "not like real-life". Much educational effort is expended demonstrating that so-called recipe knowledge must be supplemented by abstract knowledge in order to enable true understanding. The term recipe knowledge, coined by psychologist Leigh Shaffer (1981), is the knowledge of the way to use an artifact without knowledge of the fundamental principles that govern its functioning. For example, most people know many things about how to use a telephone. But many are completely ignorant of the physical principles upon which the operation of the telephone is based. They do not know why it does what it does; they only know that they can make it work. Our knowledge of many technological products in our society is also recipe knowledge. Of course, this is not an entirely negative thing. Indeed, most technological products have been designed precisely to be used without knowledge of all the principles upon which they are based (the well-known “user-friendly” concept). But it is important to understand the limitations of recipe knowledge. It does not contain the generalizeable, universal principles that allow a full understanding of the physical and social world.

A problem arises when people mistakenly view recipe knowledge as the ultimate goal of university education. Such knowledge is inevitably contextually embedded and nongeneralizeable. By its very nature it is perspective-dependent and even discourages perspective switching--because much recipe knowledge is designed for use only within a particular perspective or context. Thus everything about it works against the principles of decontextualized thinking--the importance of which I have outlined earlier. In fact, in an article on the importance of decentering cognitive styles in education, Floden, Buchmann, and Schwille (1987) stress that "unless students can break with their everyday experience in thought, they cannot see the extraordinary range of options for living and thinking; and unless students give up many commonsense beliefs, they may find it impossible to learn disciplinary concepts that describe the world in reliable, often surprising ways" (p. 485). The deepest concepts across the whole range of human knowledge--from the arts to science to moral philosophy to language--require the ability to cognize across various perspectives and situations. They require a person to see a local situation as part of a more global whole--to go beyond situation- or perspective-specific thinking. Students can only begin to gain access to this knowledge if they are first decentered from the egocentric assumption that their environment, milieu, or perspective is the only one there is--and that the immediately surrounding context is uniformly relevant to solving the problem.

Theorists such as Dennett (1991, 1996) and Bereiter (1997) have emphasized how such decontextualizing cognitive styles are computationally expensive, "unnatural," and therefore rare. As Bereiter (1997) notes, "symbolic processes are themselves exceptional and by no means representative of the great bulk of cognitive activity....They are acquired processes, culturally mediated, that enable the brain to act as if it were a different kind of device from what it had evolved to be. The device we simulate in our conscious thinking is a logic machine" (pp. 292-293). These "unrepresentative" types of cognitive have to be taught because they are instantiated in an evolutionary recent virtual machine (see Dennett, 1991) simulated by the massively parallel brain. This virtual machine is a serial processor (a von Neumann machine), and it is a powerful mechanism for logical, symbolic thought. However, this serial process is not only computationally expensive to simulate--but it is a cultural product. It is part of what Dennett (1991) calls the Gregorian brain--the part of the brain that is capable of exploiting the mental tools discovered by others (see also, Clark, 1997).

Why has schooling increasingly focused on the thinking styles of the capacity-demanding serial processes of the brain? It is conjectured here that the reason is that, increasingly, modernism requires an override of the fundamental computational biases of human cognition by the serial processes of the virtual machine (what Stanovich,1999--extrapolating from various dual-process theorists--calls System 2). In short, the processes fostered by education are like they are because modern society is like it is. Modernism has meant decontextualization. School-like processes of cognitive decontextualization gradually replace personalized interactions that are highly contextualized as modernist forces (market mechanisms, for example) spread further into every corner of the world and every aspect of modern life. Philosopher Robert Nozick (1993) describes the theme of Max Weber's writings as explicating how "economic and monetary calculation, bureacratic rationalization, general rules and procedures came to replace action based on personal ties, and market relations were extended to new arenas" (p. 180).

This shift from premodern to modernist societies was discussed recently by social theorist Francis Fukuyama (1999) who uses the distinction in sociology between gemeinschaft ("community") and gesellschaft ("society"). the former--characteristic of premodern European peasant society--consists of "a dense network of personal relationships based heavily on kinship and on the direct, face-to-face contact that occurs in a small, closed village. Norms were largely unwritten, and individuals were bound to one another in a web of mutual interdependence that touched all aspects of life....Gesellschaft, on the other hand, was the framework of laws and other formal regulations that characterized large, urban, industrial societies. Social relationships were more formalized and impersonal; individuals did not depend on one another for mutual support to nearly the same extent" (pp. 8-9). These two different societal structures no doubt call for different cognitive mechanisms in order to optimally fulfill their cognitive demands. Gemeinschaft is clearly most compatible with the fundamental computational biases of human cognition--it is closer to the EEA in which those computational biases were selected for. Gesellschaft, in contrast, would seem to call for cognitive processes to override those biases on many occasions. School-like cognition--with its contextual overrides--is shaping minds to the requirements of gesellschaft. The requirements of gemeinschaft are likely handled quite efficiently by evolutionary adaptations that are a universal endowment of human cognition.

Consider the tensions between the modern and premodern that still exist in society and how often we resolve them in favor of the former. You are a university faculty member and two students come to you requesting exemptions from a course requirement that neither has taken.

They both want to graduate at the end of this term. Student A relates to you a heartfelt and quite probably true tale of personal woe. Money is low; her single mother is ill; loan possibilities are exhausted; failure to graduate will make job possibilities (and hence the probability of repaying the loans) slim; getting a new lease on an apartment in a tight rental market will be impossible--and the list goes on. All of the details you ascertain to be true, but they are conjoined with one troublesome fact--the student has provided no "reason" why they did not take the required course. This all contrasts with student B--a comfortable student who is a sorority member and who drives a late-model convertible to your class. This student has found a "loophole" in the requirements--a complicated interaction between the changing course requirements of your department over the years and how the university registrar gives credit for courses taken at community colleges in a neighboring state. Which student do you think will be more likely to receive your dispensation? The premodern wars with the modern here. In my experience with universities, the behavior dictated by modernism usually prevails. Student B is likely to be more successful in her request. She has a "reason"; she has shown how the "rules" support her position; she appeals to the regulations that govern academic life as they are officially registered in university documents. Student A provides none of these things. She fulfills none of the requirements or regulations with her heartfelt contextualization of her situation. The department's academic advisor may well sympathize with her plight, but he will argue that "his hands are tied" because she cannot produce a decontextualized reason--an applicable rule in the published regulations--that would support her request. This example reflects the tension of the premodern and the modern and how the latter often triumphs in such conflicts. In a fascinatingly recursive way--demands for radical decontextualization form a large part of the context of modern society! This is in fact part of the stress of modern life--its constant requirement that we override fundamental computational biases that are a natural default mode of the human cognitive apparatus.

Notes

Preparation of this chapter was supported by a grant from the Social Sciences and Humanities Research Council of Canada to Keith E. Stanovich.

1 It cannot be emphasized enough that the term "bias" is used throughout this chapter to denote "a preponderating disposition or propensity" (The Compact Edition of the Oxford Short English Dictionary, p. 211) and not a processing error. That a processing bias does not necessarily imply a cognitive error is a point repeatedly emphasized by the critics of the heuristics and biases literature (Funder, 1987; Gigerenzer, 1996a; Hastie & Rasinski, 1988; Kruglanski & Ajzen, 1983), but in fact it was always the position of the original heuristics and biases researchers themselves (Kahneman, 2000; Kahneman & Tversky, 1973, 1996; Tversky & Kahneman, 1974). Thus, the use of the term bias here is meant to connote "default value" rather than "error". Under the assumption that computational biases result from evolutionary adaptations of the brain (Cosmides & Tooby, 1994), it is likely that they are efficacious in many situations.

2 Epstein, Donovan, and Denes-Raj (1999) note that many investigators have commented on the tendency of participants to spontaneously produce narrative responses to the Linda problem.

References

Adler, J. E. (1984). Abstraction is uncooperative. Journal for the Theory of Social Behaviour, 14, 165-181.

Adler, J. E. (1991). An optimist's pessimism: Conversation and conjunctions. In E. Eells & T. Maruszewski (Eds.), Probability and rationality: Studies on L. Jonathan Cohen's philosophy of science (pp. 251-282). Amsterdam: Editions Rodopi.

Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. Psychological Review, 91, 112-149.

Anderson, J. R. (1990). The adaptive character of thought. Hillsdale, NJ: Erlbaum.

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

Anderson, J. R., Reder, L. M., & Simon, H. A. (1996). Situated learning and education. Educational Researcher, 25(4), 5-11.

Andreassen, P. (1987). On the social psychology of the stock market: Aggregate attributional effects and the regressiveness of prediction. Journal of Personality and Social Psychology, 53, 490–496.

Badcock, C. (2000). Evolutionary psychology: A critical introduction. Cambridge, England: Polity Press.

Baldwin, D. A. (2000). Interpersonal understanding fuels knowledge acquisition. Current Directions in Psychological Science, 9, 40-45.

Banaji, M., & Greenwald, A. G. (1995). Implicit gender stereotyping in judgments of fame. Journal of Personality and Social Psychology, 68, 181-198.

Bar-Hillel, M., & Falk, R. (1982). Some teasers concerning conditional probabilities. Cognition, 11, 109-122.

Barkow, J., Cosmides, L., & Tooby, J. (Eds.) (12992). The adapted mind. New York: Oxford University Press.

Baron, J. (1993). Morality and rational choice. Dordrecht: Kluwer.

Baron, J. (1994). Nonconsequentialist decisions. Behavioral and Brain Sciences, 17, 1-42.

Baron, J. (1995). Myside bias in thinking about abortion. Thinking and Reasoning, 1, 221-235.

Baron, J. (1998). Judgment misguided: Intuition and error in public decision making. New York: Oxford University Press.

Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. Journal of Personality and Social Psychology, 54, 569-579.

Baron-Cohen, S. (1995). Mindblindness: An essay on autism and theory of mind. Cambridge, MA: The MIT Press.

Baron-Cohen, S. (1998). Does the study of autism justify minimalist innate modularity. Learning and Individual Differences, 10, 179-192.

Barton, R. A., & Dunbar, R. (1997). Evolution of the social brain. In A. Whiten & R. W. Byrne (Eds.), Machiavellian intelligence II: Extensions and evaluations (pp. 240-263). Cambridge: Cambridge University Press.

Belsky, G., & Gilovich, T. (1999). Why smart people make big money mistakes--And how to correct them: Lessons from the new science of behavioral economics. New York: Simon & Schuster.

Bereiter, C. (1997). Situated cognition and how to overcome it. In D. Kirshner & J. A. Whitson (Eds.), Situated cognition: Social, semiotic, and psychological perspectives (pp. 3-12). Mahwah, NJ: Erlbaum.

Blackmore, S. (1999). The meme machine. New York: Oxford University Press.

Brase, G. L., Cosmides, L., & Tooby, J. (1998). Individuation, counting, and statistical inference: The role of frequency and whole-object representations in judgment under uncertainty. Journal of Experimental Psychology: General, 127, 3-21.

Brody, N. (1997). Intelligence, schooling, and society. American Psychologist, 52, 1046-1050.

Bronfenbrenner, U., McClelland, P., Wethington, E., Moen, P., & Ceci, S. J. (1996). The state of Americans. New York: Free Press.

Broniarczyk, S., & Alba, J. W. (1994). Theory versus data in prediction and correlation tasks. Organizational Behavior and Human Decision Processes, 57, 117-139.

Brothers, L. (1990). The social brain: A project for integrating primate behaviour and neuropsychology in a new domain. Concepts in Neuroscience, 1, 27-51.

Bruner, J. (1986). Actual minds, possible worlds. Cambridge, MA: Harvard University Press.

Bruner, J. (1990). Acts of meaning. Cambridge, MA: Harvard University

Press.

Bruner, J. (1996). Frames for thinking: Ways of making meaning. In D. R. Olson & N. Torrance (Ed.)^(Eds.), Modes of thought: Explorations in culture and cognition, (pp. 93-105). New York: Cambridge University Press.

Bugental, D. B. (2000). Acquisitions of the algorithms of social life: A domain-based approach. Psychological Bulletin, 126, 187-219.

Buss, D. M. (1999). Evolutionary psychology: The new science of the mind. Boston: Allyn and Bacon.

Buss, D. M. (2000). The evolution of happiness. American Psychologist, 55, 15-23.

Byrne, R. W., & Whiten, A. (Eds.). (1988). Machiavellian intelligence: Social expertise and the evolution of intellect in monkeys, apes, and humans. Oxford: Oxford University Press.

Caporael, L. R. (1997). The evolution of truly social cognition: The core configurations model. Personality and Social Psychology Review, 1, 276-298.

Carrithers, M. (1995). Stories in the social and mental life of people. In E. N. Goody (Ed.), Social intelligence and interaction: Expressions and implications of the social bias in human intelligence (pp. 261-277). Cambridge: Cambridge University Press.

Cartwright, J. (2000). Evolution and human behavior. Cambridge, MA: MIT Press.

Chapman, L., & Chapman, J. (1967). Genesis of popular but erroneous psychodiagnostic observations. Journal of Abnormal Psychology, 72, 193–204.

Chapman, L., & Chapman, J. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. Journal of Abnormal Psychology, 74, 271–280.

Chapman, M. (1993). Everyday reasoning and the revision of belief. In J. M. Puckett & H. W. Reese (Eds.), Mechanisms of everyday cognition (pp. 95-113). Hillsdale, NJ: Erlbaum.

Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. Cognitive Psychology, 17, 391-416.

Clark, A. (1997). Being there: Putting brain, body, and world together again. Cambridge, MA: MIT Press.

Cooper, W. S. (1989). How evolutionary biology challenges the classical theory of rational choice. Biology and Philosophy, 4, 457-481.

Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. Cognition, 31, 187-276.

Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), The adapted mind, (pp. 163-205). New York: Oxford University Press.

Cosmides, L., & Tooby, J. (1994). Beyond intuition and instinct blindness: Toward an evolutionarily rigorous cognitive science. Cognition, 50, 41-77.

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

Cummins, D. D. (1996). Evidence for the innateness of deontic reasoning. Mind & Language, 11, 160-190.

Dawes, R. M. (1988). Rational choice in an uncertain world. San Diego, CA: Harcourt, Brace Jovanovich.

Dawes, R. M. (1989). Statistical criteria for establishing a truly false consensus effect. Journal of Experimental Social Psychology, 25, 1-17.

Dawes, R. M. (1990). The potential nonfalsity of the false consensus effect. In R. M. Hogarth (Ed.), Insights into decision making (pp. 179-199). Chicago: University of Chicago Press.

Dawes, R. M. (1991). Probabilistic versus causal thinking. In D. Cicchetti & W. Grove (Eds.), Thinking clearly about psychology: Essays in honor of Paul E. Meehl (Vol. 1, pp. 235–264). Minneapolis: University of Minnesota Press.

Dawes, R. M. (1994). House of cards: Psychology and psychotherapy built on myth. New York: Free Press.

Dawkins, R. (1976). The selfish gene (New edition, 1989). New York: Oxford University Press.

Dawkins, R. (1982). The extended phenotype. New York: Oxford University Press.

Dawkins, R. (1986). The blind watchmaker. New York: Norton.

Dennett, D. C. (1978). Brainstorms: Philosophical essays on mind and psychology. Cambridge, MA: MIT Press.

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

Dennett, D. C. (1991). Consciousness explained. Boston: Little Brown.

Dennett, D. C. (1995). Darwin's dangerous idea: Evolution and the meanings of life. New York: Simon & Schuster.

Dennett, D. C. (1996). Kinds of minds: Toward an understanding of consciousness. New York: Basic Books.

Denzau, A. T., & North, D. C. (1994). Shared mental models: Ideologies and institutions. Kyklos, 47, 3-31.

Dillard, A. (1999). For the time being. Toronto: Penguin Books Canada.

Doherty, M. E., & Mynatt, C. (1990). Inattention to P(H) and to P(D/~H): A converging operation. Acta Psychologica, 75, 1-11.

Dominowski, R. L. (1995). Content effects in Wason's selection task. In S. E. Newstead & J. S. B. T. Evans (Eds.), Perspectives on thinking and reasoning (pp. 41-65). Hove, England: Erlbaum.

Donaldson, M. (1978). Children's minds. London: Fontana Paperbacks.

Donaldson, M. (1993). Human minds: An exploration. New York: Viking Penguin.

Dulany, D. E., & Hilton, D. J. (1991). Conversational implicature, conscious representation, and the conjunction fallacy. Social Cognition, 9, 85-110.

Dunbar, R. (1998). Theory of mind and the evolution of language. In J. R. Hurford, M. Studdert-Kennedy, & C. Knight (Eds.), Approaches to the evolution of language (pp. 92-110). Cambridge: Cambridge University Press.

The Economist (December 12, 1998). The benevolence of self-interest. p. 80.

Edwards, K., & Smith, E. E. (1996). A disconfirmation bias in the evaluation of arguments. Journal of Personality and Social Psychology, 71, 5-24.

Einhorn, H. J., & Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgment and choice. Annual Review of Psychology, 32, 53-88.

Epstein, S., Donovan, S., & Denes-Raj, V. (1999). The missing link in the paradox of the Linda conjunction problem: Beyond knowing and thinking of the conjunction rule, the intrinsic appeal of heuristic processing. Personality and Social Psychology Bulletin, 25, 204-214.

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

Routledge.
Evans, J. St. B. T. (1989). Bias in human reasoning: Causes and consequences. London: Erlbaum Associates.
Evans, J. St. B. T. (1996). Deciding before you think: Relevance and reasoning in the selection task. British Journal of Psychology, 87, 223-240.
Evans, J. St. B. T., Barston, J., & Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. Memory & Cognition, 11, 295-306.
Evans, J. St. B. T., Newstead, S. E., & Byrne, R. M. J. (1993). Human reasoning: The psychology of deduction. Hove, England: Erlbaum.
Evans, J. St. B. T., & Over, D. E. (1996). Rationality and reasoning. Hove, England: Psychology Press.
Evans, J. St. B. T., & Over, D. E. (1999). Explicit representations in hypothetical thinking. Behavioral and Brain Sciences, 22, 763-764.
Evans, J. St. B. T., Over, D. E., & Manktelow, K. (1993). Reasoning, decision making and rationality. Cognition, 49, 165-187.
Falk, R. (1992). A closer look at the probabilities of the notorious three prisoners. Cognition, 43, 197-223.
Fiedler, K. (1988). The dependence of the conjunction fallacy on subtle linguistic factors. Psychological Research, 50, 123-129.
Flanagan, O. (1996). Self expressions: Mind, morals, and the meaning of life. New York: Oxford University Press.
Floden, R. E., Buchmann, M., & Schwille, J. R. (1987). Breaking with everyday experience. Teacher's College Record, 88, 485-506.
Fodor, J. (1983). Modularity of mind. Cambridge: MIT Press.
Frank, R. H., & Cook, P. J. (1995). The winner-take-all society. New York: Free Press.
Friedrich, J. (1993). Primary error detection and minimization (PEDMIN) strategies in social cognition: A reinterpretation of confirmation bias phenomena. Psychological Review, 100, 298-319.
Fukuyama, F. (1999). The great disruption: Human nature and the reconstitution of social order. New York: Free Press.
Funder, D. C. (1987). Errors and mistakes: Evaluating the accuracy of social judgment. Psychological Bulletin, 101, 75-90.
Garb, H. N., Florio, C. M., & Grove, W. M. (1998). The validity of the

Rorschach and the Minnesota Multiphasic Personality Inventory. Psychological Science, 9, 402-404.

Geary, D. C., & Bjorklund, D. F. (2000). Evolutionary developmental psychology. Child Development, 71, 57-65.

Gebauer, G., & Laming, D. (1997). Rational choices in Wason's selection task. Psychological Research, 60, 284-293.

Gibbard, A. (1990). Wise choices, apt feelings: A theory of normative judgment. Cambridge, MA: Harvard University Press.

Gigerenzer, G. (1991b). How to make cognitive illusions disappear: Beyond "heuristics and biases". European Review of Social Psychology, 2, 83-115.

Gigerenzer, G. (1993). The bounded rationality of probabilistic mental models. In K. Manktelow & D. Over (Eds.), Rationality: Psychological and philosophical perspectives (pp. 284-313). London: Routledge.

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

Gigerenzer, G. (1996b). Rationality: Why social context matters. In P. B. Baltes & U. Staudinger (Eds.), Interactive minds: Life-span perspectives on the social foundation of cognition (pp. 319-346). Cambridge: Cambridge University Press.

Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. Psychological Review, 102, 684-704.

Gigerenzer, G., Hoffrage, U., & Kleinbolting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. Psychological Review, 98, 506-528.

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

Goody, E. N. (Ed.). (1995). Social intelligence and interaction: Expressions and implications of the social bias in human intelligence. Cambridge: Cambridge University Press.

Goody, J. (1977). The domestication of the savage mind. New York: Cambridge University Press.

Gordon, R. A. (1997). Everyday life as an intelligence test: Effects of intelligence and intelligence context. Intelligence, 24, 203-320.

Gottfredson, L. S. (1997). Why g matters: The complexity of everyday life.

Intelligence, 24, 79-132.

Granberg, D. (1995). The Monte Hall dilemma. Personality and Social Psychology Bulletin, 31, 711-723.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. Journal of Personality and Social Psychology, 74, 1464-1480.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), Syntax and semantics: Vol. 3, Speech acts (pp. 41-58). New York: Academic Press.

Griggs, R. A. (1989). To "see" or not to "see": That is the selection task. Quarterly Journal of Experimental Psychology, 41A, 517-529.

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

Hardman, D. (1998). Does reasoning occur on the selection task? A comparison of relevance-based theories. Thinking and Reasoning, 4, 353-376.

Hasher, L., & Zacks, R. T. (1979). Automatic processing of fundamental information: The case of frequency of occurrence. Journal of Experimental Psychology: General, 39, 1372-1388.

Hastie, R., & Rasinski, K. A. (1988). The concept of accuracy in social judgment. In D. Bar-Tal & A. Kruglanski (Eds.), The social psychology of knowledge (pp. 193-208). Cambridge: Cambridge University Press.

Hertwig, R., & Gigerenzer, G. (1999). The conjunction fallacy revisited: How intelligent inferences look like reasoning errors. Journal of Behavioral Decision Making, 12, 275-305.

Hilton, D. J. (1995). The social context of reasoning: Conversational inference and rational judgment. Psychological Bulletin, 118, 248-271.

Hilton, D. J., & Slugoski, B. R. (2000). Judgment and decision making in social context: Discourse processes and rational inference. In T. Connolly, H. R. Arkes, & K. R. Hammond (Eds.), Judgment and decision making: An interdisciplinary reader (Second Edition) (pp. 651-676). Cambridge, MA: Cambridge University Press.

Hoch, S. J. (1987). Perceived consensus and predictive accuracy: The pros and cons of projection. Journal of Personality and Social Psychology, 53, 221-234.

Howson, C., & Urbach, P. (1993). Scientific reasoning: The Bayesian approach (Second Edition). Chicago: Open Court.

Humphrey, N. (1976). The social function of intellect. In P. P. G. Bateson & R. A. Hinde (Eds.), Growing points in ethology (pp. 303-317). London: Faber & Faber.

Humphrey, N. (1986). The inner eye. London: Faber & Faber.

Hunt, E. (1995). Will we be smart enough? A cognitive analysis of the coming workforce. New York: Russell Sage Foundation.

Hunt, E. (1999). Intelligence and human resources: Past, present, and future. In P. Ackerman,P. Kyllonen, & R. Richards (Eds.), Learning and individual differences: Process, trait, and content determinants (pp. 3-28). Washington, DC: American Psychological Association.

Jepson, C., Krantz, D., & Nisbett, R. (1983). Inductive reasoning: Competence or skill? Behavioral and Brain Sciences, 6, 494-501.

Johnson-Laird, P. N. (1999). Deductive reasoning. Annual Review of Psychology, 50, 109-135.

Jolly, A. (1966). Lemur social behaviour and primate intelligence. Science, 153, 501-506.

Kahneman, D. (2000). A psychological point of view: Violations of rational rules as a diagnostic of mental processes. Behavioral and Brain Sciences, 23, 681-683.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.) (1982). Judgment under uncertainty: Heuristics and biases. Cambridge: Cambridge University Press.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. Psychological Review, 80, 237-251.

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

Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. Psychological Review, 103, 582-591.

Kahneman, D., & Tversky, A. (Eds.). (2000). Choices, values, and frames. Cambridge University Press:

Kelley, D. (1990). The art of reasoning. New York: Norton.

Kim, J. (1994, September). Watch out for investing newsletters luring you with outdated returns. Money Magazine, pp. 12–13.

King, R. N., & Koehler, D. J. (2000). Illusory correlations in graphological inference. Journal of Experimental Psychology: Applied, 6, 336-348.

Kirby, K. N. (1994). Probabilities and utilities of fictional outcomes in Wason's four-card selection task. Cognition, 51, 1-28.

Klaczynski, P. A., Gordon, D. H., & Fauth, J. (1997). Goal-oriented critical reasoning and individual differences in critical reasoning biases. Journal of Educational Psychology, 89, 470-485.

Klayman, J., & Ha, Y. (1987). Confirmation, disconfirmation, and information in hypothesis testing. Psychological Review, 94, 211-228.

Koehler, J. J. (1993). The influence of prior beliefs on scientific judgments of evidence quality. Organizational Behavior and Human Decision Processes, 56, 28-55.

Koriat, A., Lichtenstein, S., & Fischhoff, B. (1980). Reasons for confidence. Journal of Experimental Psychology: Human Learning and Memory, 6, 107-118.

Kornblith, H. (1993). Inductive inference and its natural ground. Cambridge, MA: MIT University Press.

Krantz, D. H. (1981). Improvements in human reasoning and an error in L. J. Cohen's. Behavioral and Brain Sciences, 4, 340-341.

Krueger, J., & Zeiger, J. (1993). Social categorization and the truly false consensus effect. Journal of Personality and Social Psychology, 65, 670-680.

Kruglanski, A. W., & Ajzen, I. (1983). Bias and error in human judgment. European Journal of Social Psychology, 13, 1-44.

Kuhn, D. (1991). The skills of argument. Cambridge: Cambridge University Press.

Kuhn, D. (1993). Connecting scientific and informal reasoning. Merrill-Palmer Quarterly, 38, 74-103.

Kuhn, D. (1996). Is good thinking scientific thinking? In D. R. Olson & N. Torrance (Eds.), Modes of thought: Explorations in culture and cognition (pp. 261-281). New York: Cambridge University Press.

Kummer, H., Daston, L., Gigerenzer, G., & Silk, J. B. (1997). The social intelligence hypothesis. In P. Weingart,S. D. Mitchell,P. J. Richerson, & S. Maasen (Eds.), Human by nature: Between biology and the social sciences (pp. 157-179). Mahwah, NJ: Lawrence Erlbaum Associates.

Kunda, Z. (1999). Social cognition: Making sense of people. Cambridge, MA: MIT Press.

Langer, E. J. (1989). Mindfulness. Reading, MA: Addison-Wesley.

Langer, E. J., Blank, A., & Chanowitz, B. (1978). The mindlessness of ostensibly thoughtful action: The role of "placebic" information in interpersonal interaction. Journal of Personality and Social Psychology, 36, 635-642.

Lefkowitz, B. (1997). Our guys: The Glen Ridge rape and the secret life of the perfect suburb. Berkeley, CA: University of California Press.

Levin, I. P., Wasserman, E. A., & Kao, S. F. (1993). Multiple methods of examining biased information use in contingency judgments. Organizational Behavior and Human Decision Processes, 55, 228-250.

Levinson, S. C. (1983). Pragmatics. Cambridge: Cambridge University Press.

Levinson, S. C. (1995). Interactional biases in human thinking. In E. Goody (Eds.), Social intelligence and interaction (pp. 221-260). Cambridge: Cambridge University Press.

Liberman, N., & Klar, Y. (1996). Hypothesis testing in Wason's selection task: Social exchange cheating detection or task understanding. Cognition, 58, 127-156.

Lilienfeld, S. O. (1999). Projective measures of personality and psychopathology: How well do they work? Skeptical Inquirer, 23(5), 32-39.

Luria, A. R. (1976). Cognitive development: Its cultural and social foundations. Cambridge, MA: Harvard University Press.

Macdonald, R. R., & Gilhooly, K. J. (1990). More about Linda or conjunctions in context. European Journal of Cognitive Psychology, 2, 57-70.

Malkiel, B. G. (1999). A random walk down Wall Street. New York: W. W. Norton.

Manktelow, K. I. (1999). Reasoning & Thinking. Psychology Press.

Manktelow, K. I., & Evans, J. S. B. T. (1979). Facilitation of reasoning by realism: Effect or non-effect? British Journal of Psychology, 70, 477-488.

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

Manktelow, K. I., Sutherland, E., & Over, D. E. (1995). Probabilistic factors in deontic reasoning. Thinking and Reasoning, 1, 201-219.

Margolis, H. (1987). Patterns, thinking, and cognition. Chicago: University of Chicago Press.

Margolis, H. (1996). Dealing with risk. Chicago: University of Chicago Press.

Markovits, H., & Nantel, G. (1989). The belief-bias effect in the production and evaluation of logical conclusions. Memory & Cognition, 17, 11-17.

McCain, R. A. (1991). A linguistic conception of rationality. Social Science Information, 30(2), 233-255.

Mercer, T. (2000, September, 16). Navigating the shark-eat-shark world of 'compassionate' airfares. The Globe and Mail (Toronto). p. T5.

Messer, W. S., & Griggs, R. A. (1993). Another look at Linda. Bulletin of the Psychonomic Society, 31, 193-196.

Minsky, M. (1985). The society of mind. New York: Simon & Schuster.

Mitchell, P., Robinson, E. J., Isaacs, J. E., & Nye, R., M. (1996). Contamination in reasoning about false belief: An instance of realist bias in adults but not children. Cognition, 59, 1-21.

Mithen, S. (1996). The prehistory of mind: The cognitive origins of art and science. London: Thames and Hudson.

Neimark, E. (1987). Adventures in thinking. San Diego, CA: Harcourt Brace Jovanovich.

Newstead, S. E., & Evans, J. St. B. T. (Eds.) (1995). Perspectives on thinking and reasoning. Hove, England: Erlbaum.

Nickerson, R. S. (1999). How we know--and sometimes misjudge--what others know: Imputing one's own knowledge to others. Psychological Bulletin, 125, 737-759.

Nickerson, R. S., Baddeley, A., & Freeman, B. (1987). Are people's estimates of what other people know influenced by what they themselves know? Acta Psychologica, 64, 245-259.

Nisbett, R. E., & Ross, L. (1980). Human inference: Strategies and shortcomings of social judgment. Englewood Cliffs, NJ: Prentice-Hall.

Nozick, R. (1993). The nature of rationality. Princeton, NJ: Princeton University Press.

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

Oaksford, M., & Chater, N. (1995). Theories of reasoning and the computational explanation of everyday inference. Thinking and Reasoning, 1, 121-152.

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

Oaksford, M., & Chater, N. (1998). Rationality in an uncertain world. Hove, England: Psychology Press.

Oaksford, M., Chater, N., Grainger, B., & Larkin, J. (1997). Optimal data selection in the reduced array selection task (RAST). Journal of Experimental Psychology: Learning, Memory, and Cognition, 23, 441-458.

Oatley, K. (1992). Best laid schemes: The psychology of emotions. Cambridge: Cambridge University Press.

Oatley, K. (1996). Emotions, rationality, and informal reasoning. In J. V. Oakhill & A. Garnham (Eds.), Mental models in cognitive science: Essays in honor of Phil Johnson-Laird, (pp. 1-22). Hove, England: Psychology Press.

Oatley, K. (1999). Why fiction may be twice as true as fact: Fiction as cognitive and emotional simulation. Review of General Psychology, 3, 101-117.

Olson, D. R. (1977). From utterance to text: The bias of language in speech and writing. Harvard Educational Review, 47, 257-281.

Olson, D. R. (1986). Intelligence and literacy: The relationships between intelligence and the technologies of representation and communication. In R. J. Sternberg & R. K. Wagner (Eds.), Practical intelligence (pp. 338-360). Cambridge: Cambridge University Press.

Olson, D. R. (1994). The world on paper. Cambridge, England: Cambridge University Press.

Perkins, D. N., Farady, M., & Bushey, B. (1991). Everyday reasoning and the roots of intelligence. In J. Voss,D. Perkins, & J. Segal (Eds.), Informal reasoning and education (pp. 83-105). Hillsdale, NJ: Erlbaum.

Piaget, J. (1972). Intellectual evolution from adolescence to adulthood. Human Development, 15, 1-12.

Pinker, S. (1994). The language instinct. New York: William Morrow.

Pinker, S. (1997). How the mind works. New York: Norton.

Plotkin, H. (1998). Evolution in mind: An introduction to evolutionary psychology. Cambridge, MA: Harvard University Press.

Politzer, G., & Noveck, I. A. (1991). Are conjunction rule violations the result of conversational rule violations? Journal of Psycholinguistic

Research, 20, 83-103.

Pollard, P., & Evans, J. St. B. T. (1987). Content and context effects in reasoning. American Journal of Psychology, 100, 41-60.

Pollock, J. L. (1995). Cognitive carpentry: A blueprint for how to build a person. Cambridge, MA: MIT Press.

Reber, A. S. (1992a). An evolutionary context for the cognitive unconscious. Philosophical Psychology, 5, 33-51.

Reber, A. S. (1992b). The cognitive unconscious: An evolutionary perspective. Consciousness and Cognition, 1, 93-133.

Rescher, N. (1988). Rationality: A philosophical inquiry into the nature and rationale of reason. Oxford: Oxford University Press.

Rips, L. J. (1994). The psychology of proof. Cambridge, MA: MIT Press.

Sá, W., & Stanovich, K. E. (2001). The domain specificity and generality of mental contamination: Accuracy and projection in mental content judgments. British Journal of Psychology, 92, 281-302.

Sá, W., West, R. F., & Stanovich, K. E. (1999). The domain specificity and generality of belief bias: Searching for a generalizable critical thinking skill. Journal of Educational Psychology, 91, 497-510.

Schank, R. (1991). Tell me a story. New York: Scribners.

Schustack, M. W., & Sternberg, R. J. (1981). Evaluation of evidence in causal inference. Journal of Experimental Psychology: General, 110, 101-120.

Schwarz, N. (1996). Cognition and communication: Judgmental biases, research methods, and the logic of conversation. Mahweh, NJ: Lawrence Erlbaum Associates.

Shaffer, L. (1981). The growth and limits of recipe knowledge. Journal of Mind and Behavior, 2, 71–83.

Shefrin, H., & Statman, M. (1986). How not to make money in the stock market. Psychology Today, 20, 53–57.

Shiller, R. (1987). The volatility of stock market prices. Science, 235, 33–37.

Shontz, F., & Green, P. (1992). Trends in research on the Rorschach: Review and recommendations. Applied and Preventive Psychology, 1, 149–156.

Sigel, I. E. (1993). The centrality of a distancing model for the development of representational competence. In R. Cocking & K. Renninger (Eds.), The development and meaning of psychological distance (pp. 141-158).

Hillsdale, NJ: Erlbaum.

Skyrms, B. (1996). The evolution of the social contract. Cambridge: Cambridge University Press.

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

Slugoski, B. R., & Wilson, A. E. (1998). Contribution of conversation skills to the production of judgmental errors. European Journal of Social Psychology, 28, 575-601.

Sperber, D., Cara, F., & Girotto, V. (1995). Relevance theory explains the selection task. Cognition, 57, 31-95.

Sperber, D., & Wilson, D. (1986). Relevance: Communication and cognition. Cambridge, MA: Harvard University Press.

Stanovich, K. E. (1999). Who is rational? Studies of individual differences in reasoning. Mahweh, NJ: Erlbaum.

Stanovich, K. E., & West, R. F. (1997). Reasoning independently of prior belief and individual differences in actively open-minded thinking. Journal of Educational Psychology, 89, 342-357.

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

Stanovich, K. E., & West, R. F. (1998b). Individual differences in framing and conjunction effects. Thinking and Reasoning, 4, 289-317.

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

Stanovich, K. E., & West, R. F. (1998d). Who uses base rates and P(D/~H)? An analysis of individual differences. Memory & Cognition, 28, 161-179.

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

Sternberg, R. J. (1997). The concept of intelligence and its role in lifelong learning and success. American Psychologist, 52, 1030-1037.

Sternberg, R. J., & Wagner, R. K. (Eds.) (1986). Practical intelligence. Cambridge: Cambridge University Press.

Stevenson, R. J. (1993). Rationality and reality. In K. Manktelow & D. Over (Eds.), Rationality: Psychological and philosophical perspectives (pp. 61-82). London: Routledge.

Stevenson, R. J., & Over, D. E. (1995). Deduction from uncertain premises. Quarterly Journal of Experimental Psychology, 48A, 613-643.

Sutherland, S. (1992). Irrationality: The enemy within. London: Constable.

Thagard, P. (1992). Conceptual revolutions. Princeton, NJ: Princeton University Press.

Thagard, P., & Nisbett, R. E. (1983). Rationality and charity. Philosophy of Science, 50, 250-267.

Thaler, R. H. (1992). The winner's curse: Paradoxes and anomalies of economic life. New York: Free Press.

Thomas, M., & Karmiloff-Smith, A. (1998). Quo vadis modularity in the 1990s? Learning and Individual Differences, 10, 245-250.

Tomasello, M. (1998). Social cognition and the evolution of culture. In J. Langer & M. Killen (Eds.), Piaget, evolution, and development (pp. 221-245). Mahwah, NJ: Lawrence Erlbaum Associates.

Tomasello, M. (1999). The cultural origins of human cognition. Cambridge, MA: Harvard University Press.

Tooby, J., & Cosmides, L. (1992). The psychological foundations of culture. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), The adapted mind, (pp. 19-136). New York: Oxford University Press.

Tsimpli, I., & Smith, N. (1998). Modules and quasi-modules: Language and theory of mind in a polyglot savant. Learning and Individual Differences, 10, 193-216.

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

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

Updegrave, W. L. (1995, August). Why funds don't do better. Money Magazine, 58–65.

Voss, J., Perkins, D. & Segal J. (Eds.). (1991). Informal reasoning and education. Hillsdale, NJ: Erlbaum.

Wagenaar, W. A. (1988). Paradoxes of gambling behavior. Hove, England: Erlbaum.

Wason, P. C. (1966). Reasoning. In B. Foss (Eds.), New horizons in psychology (pp. 135-151). Harmonsworth, England: Penguin.

Whiten, A., & Byrne, R. W. (Eds.) (1997). Machiavellian intelligence II:

Extensions and evaluations. Cambridge: Cambridge University Press.

Widiger, T., & Schilling, M. (1980). Towards a construct validation of the Rorschach. Journal of Personality Assessment, 44, 450–459.

Wolford, G., Miller, M. B., & Gazzaniga, M. S. (2000). The left hemisphere's role in hypothesis formation. Journal of Neuroscience, 20 (RC64), 1-4

Wood, J. M., Nezworski, T., & Stejskal, W. J. (1996). The comprehensive system for the Rorschach: A critical examination. Psychological Science, 7, 3–10.