Chapter 13
Empirical Evidence of Irrationality

Abstract Whereas the previous two chapters responded to philosophical challenges to our scientific knowledge this chapter explores a more practical threat to scientific knowledge. This challenge comes from research which suggests we are subject to a number of biases and irrational processes when forming our beliefs. Numerous studies have seemingly shown that people are prone to make systematic errors of reasoning in particular kinds of cases. Some take this evidence of human irrationality to undercut our knowledge in general, and hence, our scientific knowledge as well. This chapter argues that this challenge does not pose a significant threat to our scientific knowledge. Although there is evidence for human irrationality, we have ways of keeping this sort of irrationality contained so that it does not "infect" all of our beliefs. So, while we are prone to make systematic errors in certain cases, we are aware of our proclivities, and we can take steps to counteract our natural shortcomings.

In the two previous chapters we explored skeptical challenges to our scientific knowledge that come by way of philosophical arguments threatening a wide swath of our knowledge of the world around us. Although the arguments for these forms of skepticism are interesting, and they pose challenges to our having the knowledge we ordinarily take ourselves to have, we have the requisite tools for overcoming the challenges they present. Considering these sorts of challenges to our scientific knowledge is important because doing so helps illuminate various key features of the nature of knowledge in general, which in turn can help to deepen our understanding of scientific knowledge in particular.

Despite their importance, readers of a more practical disposition are apt to find the traditional skeptical arguments of the previous two chapters unimpressive and perhaps not worth taking all that seriously. In this chapter we will explore a challenge to our scientific knowledge which does not rest upon philosophical speculation. The challenge to our scientific knowledge we will consider in this chapter is more concrete than the previous challenges because it arises from empirical studies of human reasoning. Several studies seem to suggest that we systematically make mistakes in reasoning. On the whole these studies appear to paint a rather dismal picture of our reasoning abilities (Nisbett and Borgida 1975). The mistakes these studies reveal are common enough that they lead some to worry

K. McCain, *The Nature of Scientific Knowledge*, Springer Undergraduate Texts in Philosophy, DOI 10.1007/978-3-319-33405-9_13

we are systematically irrational in our reasoning. Given widespread systematic irrationality, one might worry that we lack the sort of knowledge we take ourselves to have. After all, if we are often wrong when we reason as we do, then it seems that when we get things right as a result of our reasoning it is largely a matter of luck. Consequently, these empirical studies appear to seriously threaten our scientific knowledge and, as will become clear shortly, a large amount of our ordinary (non-scientific) knowledge as well. While the threat posed here may not be as far-reaching as those of the skeptical arguments in the previous chapters, it is still very worrisome—especially, given the fact that this threat is grounded in empirical research not just philosophical speculation.

13.1 The Empirical Evidence

There have been numerous studies on human reasoning and the mistakes we sometimes make when given specific reasoning tasks.[1] Rather than conduct an exhaustive survey of this literature we will confine our examination of the issue to two of the most widely discussed experiments and the evidence they are purported to provide for our systematic irrationality.

13.1.1 Errors in Deductive Reasoning: The Selection Task

One of the most well known experiments concerning reasoning is Peter Wason's (1966) selection task experiment.[2] The purpose of this experiment is to test people's ability to accurately apply principles of deductive reasoning. In this experiment subjects are shown four cards. Each card has a visible side and a covered side. Subjects are told that each card has a letter on one side and a number on the other side. The visible sides are as follows:

Card 1	Card 2	Card 3	Card 4
A	K	4	7

Subjects are then asked which of the four cards they *need* to see the covered side of in order to determine if the following rule is true:

[1]For helpful surveys of this literature see Gilovich (1991), Kahneman et al. (1982), Nisbett and Ross (1980), Plous (1993), and Tweney et al. (1981).

[2]See Wason (1968) for a similar, but more extensive, experiment. See Wason and Johnson-Laird (1972) for further discussion of these sorts of experiments.

> If a card has a vowel on one side, then it has an even number on the other side. (Stein 1996, p. 80)

In order to correctly evaluate this rule subjects should follow this principle:

> *Conditional-Testing Principle*: To test the truth of a conditional, examine cases where the antecedent is true to make sure that the consequent is true and examine cases where the consequent is false to make sure the antecedent is false. (Stein 1996, p. 81)

The reason for this is that the only way a conditional statement, such as the rule subjects are asked to evaluate, can be false is if the antecedent (in this case "a card has a vowel on one side") is true and the consequent (in this case "it has an even number on the other side") is false. So, the only situations in which there is a violation of the rule are situations in which a card has a vowel on one side and an odd number on the other side. Any other combination will be consistent with the rule. Thus, if one is following the Conditional-Testing Principle in her reasoning about this task, she will rightly note that in order to test the rule you need to see the covered sides of card 1 and card 4.

Experimenters found that subjects overwhelmingly give the wrong answers in this sort of experiment. In one instance of the experiment only 5 out of 128 subjects correctly answered that one needs to see the covered sides of cards 1 and 4 in order to test the rule (Stich 1990). Unfortunately, these results are fairly typical. In the many iterations of the experiment it turns out that 33 % of subjects claim that just card 1 needs to be seen in order to test the rule, 46 % claim that card 1 and card 2 need to be uncovered, 17 % maintain that some other, incorrect, combination of cards needs to be examined—less than 5 % of subjects give the correct answer that cards 1 and 4 need to be seen in order to test the rule experimenters presented to them (Stein 1996).

As if these results were not bad enough, matters seem to be even worse. In these sorts of experiments many subjects not only give the wrong answer concerning which cards one needs to see in order to test the rule, but they also continue to give wrong answers in repeated trials (Wason and Johnson-Laird 1972). Even more worrisome is the fact that many subjects will insist on their original (erroneous) answers for testing the rule even after it has been explained to them that the general principles their answers rely upon are mistaken and fail to provide genuine tests of the rule (Wason 1968; Wason and Johnson-Laird 1972). For example, some subjects who said that only the covered side of card 1 needs to be seen in order to evaluate the rule will continue to insist that only card 1 needs to be examined. They will continue to insist this even after it has been explained to them why only examining card 1 fails to implement the correct principle, i.e. it fails to provide a genuine test of the rule in question. These results seem to suggest, at least prima facie, that we exhibit fairly robust and widespread irrationality when it comes to the application of some basic rules of deductive reasoning.

13.1.2 Errors in Inductive Reasoning: The Conjunction Fallacy

Amos Tversky and Daniel Kahneman's (1983) experiment examining our ability to make judgments concerning how comparatively likely various events are is perhaps the most widely discussed experiment suggesting that we make serious errors when reasoning inductively. In this experiment subjects read the following description of a woman:

> Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. (Tversky and Kahneman 1983, p. 297)

After reading the description of Linda subjects are asked to rank a series of statements from the most probable to the least probable based on the information they have been given about her:

1. Linda is a teacher in an elementary school
2. Linda works in a bookstore and takes Yoga classes
3. Linda is active in the feminist movement
4. Linda is a psychiatric social worker
5. Linda is a member of the League of Women Voters
6. Linda is a bank teller
7. Linda is an insurance salesperson
8. Linda is a bank teller and is active in the feminist movement (Tversky and Kahneman 1983, p. 297)

Interestingly, 85 % of subjects ranked (3) as more likely than (8) and (8) as more likely than (6). This is of particular interest because in ranking (8) as more likely than (6) subjects are violating a basic rule of probability theory. When two events, such as being a bank teller (B) and being active in the feminist movement (F), are independent the probability that both events occur is equal to the product of the probability of each event occurring. In other words, the probability of (B&F) = the probability of (B) **x** the probability of (F). Hence, there is no way that the probability (B&F) can be higher than the probability of (B) (or the probability of (F) for that matter). The story is more complicated when the two events are not independent. However, even when the events are not independent there are no cases where the probability of (B&F) will be greater than the probability of (B) (Stich 1990). So, when subjects rank (8) as more likely than (6) they are violating a basic principle of probability theory—they are committing the *conjunction fallacy* (the fallacy of thinking that a conjunction is more probable than one of its component conjuncts).[3]

[3]The reason that committing the conjunction fallacy is irrational is that it makes one susceptible to a "Dutch book". A Dutch book is a series of bets which are such that no matter what the outcome of the series of events betted upon the person accepting the series of bets will lose all of her money. When someone commits the conjunction fallacy she is accepting a distribution of probabilities which could be used to construct a Dutch book against her. For more on how

Again, these results seem to be fairly robust. In order to make sure that subjects did not interpret (6) as meaning "Linda is a bank teller and is not active in the feminist movement" the experiment was run with (6) replaced with "Linda is a bank teller whether or not she is active in the feminist movement". Similar results were found—subjects still tended to rank (8) as more probable than the revised version of (6) (Stich 1990). Further, the subjects of these experiments were not just "statistically naïve undergraduates", but also "presumably *not* statistically naïve psychology graduate students"—suggesting that committing the conjunction fallacy is a fairly widespread error (Stein 1996, p. 94). As with the results of the selection task, the results of these experiments seem to suggest, at least prima facie, that we exhibit a fairly robust and widespread irrationality when it comes to a particular kind of reasoning. In this case the empirical evidence appears to suggest that we make systematic mistakes with respect to inductive reasoning (at least when that reasoning deals with probabilities).

13.2 Responses to the Threat of Irrationality

On the basis of experimental findings like those discussed in the previous section many suppose that we exhibit systematic irrationality. Given the systematic nature of this irrationality and the fact that it seems to be widespread, one might worry that in many of the cases where we do reason correctly it is largely a matter of luck. In other words, one might worry that we cannot sufficiently trust our reasoning abilities in a wide range of cases in order to have knowledge. Thus, it may seem that we know much less than we ordinarily take ourselves to know—this would seriously affect our knowledge in ordinary situations, and possibly our scientific knowledge as well.

Fortunately, as we will see, there are a number of ways to respond to empirical evidence of our irrationality. We will not examine any of these responses in exhaustive detail; nevertheless, we will explore them in sufficient detail to get a good idea of their effectiveness. While it may be that none of these responses is ultimately decisive, they do give us reason to be optimistic about our rationality. More importantly for our purposes, we will see that there is good reason to think that even if we are irrational in many of the ways suggested by these experiments, this does not undermine our scientific knowledge. Let us turn to some of the responses on behalf of our being rational.

committing the conjunction fallacy makes one Dutch bookable see Stein (1996). For an accessible general discussion of Dutch books see Skyrms (1986).

13.2.1 Impossibility Responses

One general category of responses to the threat these empirical studies seem to pose for our rationality is impossibility responses. These responses argue that it is impossible for humans to be systematically irrational. If this is correct, then while the empirical research suggests that we make a variety of mistakes in certain situations it does not (and could not) establish that irrationality is widespread. It will be worth briefly looking at a few attempts at impossibility responses.

The first impossibility response comes from consideration of the principle of charity. The principle of charity was originally developed as a principle of how language translations should be conducted (Quine 1960). According to this principle, when translating the assertions of someone speaking an unknown language we should translate what she says in such a way that we do not understand her to be speaking absurdities. As Edward Stein (1996, p. 112) puts the point, "if I translate someone speaking an unknown language as saying something absurd, I should scrutinize my translation, not the rationality of the person." Of course, this is not to say that it is never appropriate to translate someone as saying something absurd. There are clearly times when the speaker intends to say absurd things. There are also times when a speaker inadvertently says absurd things. Nonetheless, the idea is that in the vast majority of cases when a speaker is making sincere assertions we should interpret her in such a way that what she says is rational.

One way to argue that people must be rational is to argue that the principle of charity applies to how we should understand the principles of reasoning which people employ.[4] Essentially, the idea is that "in order to characterize a person's reasoning competence, you have to assume that the principles she uses are basically rational" (Stein 1996, p. 116). Given this, one might think that it is simply impossible for us to be systematically irrational despite what the empirical research initially seems to suggest.

At first glance this appears to be a promising response to the suggestion that studies like those we explored above show we are irrational. Before this can be settled though, we have to first recognize that there are two version of the principle of charity which might be applied to our reasoning. The *strong version* of this principle is that "people should *never* be interpreted as irrational" (Stein 1996, p. 116). The *weak version* of this principle is that "*unless* there is strong empirical evidence to the contrary, people should be interpreted as rational" (Stein 1996, p. 116). The strong version of the principle of charity is clearly too strong. Surely, there are times when people are irrational. So, the strength of this response rests on the weak version of the charity principle. While the weak version seems to be a plausible principle, it does not seem to show that it is *impossible* for us to be systematically irrational. After all, this principle allows that when there is sufficient empirical evidence it is correct to interpret people as irrational. Thus, it seems that

[4]Davidson (1985) and Dennett (1978, 1987) both argue that the principle of charity applies to our reasoning.

the principle of charity does not provide good grounds for thinking it is impossible that we are systematically irrational.

The second impossibility response also comes by way of a comparison with language. One plausible way of understanding the grammatical rules of a particular language is that they are generalizations, or idealizations, of how competent speakers actually use the language. Hence, when it comes to the English language, for example, the grammatical rules which govern that language are generalizations from how people who speak English use the language. If this is correct, then while competent speakers of a language might at times break the grammatical rules, they cannot be totally off base in their application of the grammatical rules since the rules themselves come from how people actually speak.

Some argue that rules of good reasoning are analogous to grammatical rules for a language (Cohen 1983). The idea is that the rules for good reasoning, like the grammatical rules of a language, are generalizations of how people actually reason. If this is correct, then it simply cannot be that people tend to reason incorrectly—we cannot be systematically irrational.

One might worry about this sort of response for at least two reasons.[5] First of all, one might simply doubt the analogy between reasoning and language use (Feldman 2003). After all, it seems strange to say that people routinely misuse the language that they speak; it does not seem all that strange to think that people routinely make errors of reasoning. The first would make it impossible to understand speakers of a language, but the latter does not appear to yield any analogous results. Second, one might worry that even if we grant the analogy between language and reasoning, it will not be enough to show that people are not systematically irrational. It could be that the rules of good reasoning—what it is rational to infer on the basis of a given body of evidence—are generalizations from when we think very carefully about our evidence and what it supports, i.e. when we are reasoning at our very best. Of course, "our common and unconsidered judgments could be frequently irrational, even if rationality is determined by what we do when we are more careful" (Feldman 2003, p. 162). So, it may be that the claim that we are systematically irrational is consistent with the rules of good reasoning being like those of correct grammar. Thus, it is not clear if this response can provide an adequate defense of human rationality.

Perhaps a different approach to arguing that systematic irrationality is impossible will succeed. One might think that whereas the first two responses seek to show that it is impossible for us to be systematically irrational on the basis of conceptual grounds a third impossibility response which seeks to establish this point on empirical grounds would be more effective. Specifically, this third response involves

[5]Cohen's (1983) defense of this response is quite sophisticated, and it is worthy of careful consideration. However, for our purposes it is enough to grasp the general nature of this sort of response and some reasons that one might be skeptical of its effectiveness. For detailed critical discussion of Cohen's arguments concerning human irrationality see Stein (1996) and Stich (1990).

arguing on the basis of evolution that humans must be rational.[6] This response has two basic steps (Stein 1996). The first step relies on the claim that natural selection will select for mechanisms which yield true beliefs. The second step involves arguing that there is a close connection between our being rational and our having mechanisms which yield true beliefs. The idea is that natural selection "guarantees that all normal cognitive systems will be rational" (Stich 1990, p. 16). Simply put, this response is based on the idea that since we have evolved—as a species we have withstood the battle of survival of the fittest—we must be rational.

While this evolutionary response does have some appeal, it faces some serious objections.[7] One problem with this response is that many of the beliefs and reasoning tasks that the empirical studies we have discussed examine concern fairly abstract matters. It is far from clear that beliefs about abstract and theoretical matters have much connection with survival (Feldman 2003). Another problem is that it seems clear that it is not necessary for beliefs to be true or even rational in order to be conducive to survival. As Stephen Stich (1990, p. 62) explains "a very cautious, risk-aversive inferential strategy—one that leaps to the conclusion that danger is present on very slight evidence—will typically lead to false beliefs more often, and true ones less often, than a less hair-trigger one that waits for more evidence before rendering judgment." Yet, as Stich points out, it is possible that the first sort of cognitive system is selected for and the second is not. The reason for this is that "natural selection does not care about truth; it cares only about reproductive success. And from the point of view of reproductive success, it is often better to be safe (and wrong) than sorry" (Stich 1990, p. 62). Obviously, these considerations do not show that natural selection has failed to lead to our having cognitive processes which guarantee our rationality. They do, however, render the claim that natural selection *must* have so equipped us doubtful.

13.2.2 Questioning the Evidence

Given the challenges facing impossibility responses, one might be tempted to take a different approach to defending our rationality from the sorts of studies discussed above. A plausible way of doing this is to argue that although it is possible we are irrational, there is good reason to doubt that the empirical research has shown that we are. In other words, one might question the evidence of our irrationality purportedly provided by these studies.

Researchers have examined both the selection task and the conjunction fallacy studies by running variations of the studies. The results of these additional studies might be taken to cast doubt on how much evidence the original studies really give

[6] Several philosophers have expressed sympathy for this sort of argument including, among others, Dennett (1987), Goldman (1986), Lycan (1988), Millikan (1984), and Sober (1981).

[7] For detailed discussion and criticism of this argument see Stein (1996) and Stich (1990).

us for thinking that we are systematically irrational. Studies have found that when it comes to selection tasks subjects do considerably better when the selection task is less abstract than the task of the original experiment.[8] In fact when given a selection task involving the evaluation of a rule like "Every time I go to New York, I travel by train" subjects correctly determine which cards need to be uncovered nearly 66 % of the time (Stein 1996, p. 83). This is considerably better than subjects' performance in the original, rather abstract, selection task. In light of results like this, one might think that in the original selection task subjects are simply misinterpreting the task, and so applying a different rule of reasoning than they normally would (Cohen 1983).

Similar results were found for experiments concerning the conjunction fallacy. Studies have shown that when "the problem is phrased in such a way that subjects are being asked to indicate *frequency* rather than *probability*, their responses are in accord with the conjunction principle" (Stein 1996, p. 100).[9] For instance, when subjects are given the description of Linda from above and told that a certain number of people fit it, they answer the question as to how many of those people are (1) bank tellers and how many are (2) bank tellers and feminists in ways that are consistent with the conjunction principle (Gigerenzer 1991). That is to say, when the experiment is formed this way subjects do not tend to commit the conjunction fallacy. Again, this may lead one to think that subjects are merely misunderstanding what is asked of them in the original experiment (Feldman 2003).

The results of these further studies seem to weaken the empirical evidence for our irrationality provided by the original experiments. Nevertheless, it is not clear that this is enough to defend us from the charge of systematic irrationality. One reason this may not be enough is that at least in some cases, such as subjects' misinterpretation in the original selection task experiments and their correct interpretation in the concrete cases, it seems subjects are relying on content-specific rules. That is, they are using rules of reasoning which depend upon the specific subject matter to which the rule is being applied. Their reliance on content-specific rules explains why they use the correct rule when reasoning about the concrete cases, but they get things wrong in the original case with its more abstract content. The reliance on content-specific rules is worrisome because the correct rules of reasoning, such as the Conditional-Testing Principle, are content-neutral (Stein 1996). The correct rules of reasoning apply regardless of the particular subject being discussed. In light of this, one might worry that even if we tend to use content-specific rules which yield the appropriate results in some cases, we are still using the incorrect reasoning rules. Hence, one might worry that we are still exhibiting systematic irrationality—even when we are getting the correct answers to selection tasks!

Another reason that this may not be enough to defend our rationality is that there are a large number of additional studies which purport to show that we are

[8] See Wason and Shapiro (1971) and Johnson-Laird et al. (1972).

[9] See Fielder (1988).

systematically irrational. For example, studies seem to show that belief polarization is quite common. Belief polarization is the phenomenon which occurs when someone begins with a particular view on a topic and then receives mixed evidence on the topic (evidence that is balanced between the two sides of the debate). The person's belief tends to become more polarized; she tends to believe the position that she started with even more strongly than she initially did after receiving the mixed evidence (Lord et al. 1979).[10] It is not clear how subjects could be misinterpreting the task here. So, one might worry that this provides evidence for systematic irrationality. Belief polarization is particularly interesting for our purposes because it may make one worry that while it is widely acknowledged that scientific knowledge is theory-laden to some extent, perhaps these results show that it is problematically so. In addition to this sort of phenomenon, there are numerous studies suggesting that we exhibit a wide range of cognitive biases in our reasoning.[11] One might worry that it will be difficult to show that in all of these cases in which widespread irrationality is suggested by the empirical evidence there is a misinterpretation going on which allows for our being rational despite making mistakes. Consequently, it at least seems to remain an open, empirical question whether we exhibit widespread irrationality (Stein 1996).

13.2.3 A Modest Response

It appears that the impossibility responses and questioning the empirical evidence do not provide clear defenses of our rationality. Given this, one might conclude that the empirical evidence shows we are systematically irrational in various ways—we fall prey to cognitive biases, make mistakes when it comes to deductive and inductive reasoning, and so on. In light of this systematic irrationality one might think that we simply do not, and cannot, have much of the knowledge that we take ourselves to have. Further, one might conclude that we do not have much, if any scientific knowledge. This is a very dismal assessment indeed.

Although one might conclude that we do not have much scientific or ordinary knowledge on the basis of these empirical results, doing so would be a mistake. What is more, this would be a mistake *even if* we grant that the empirical evidence does demonstrate we are systematically irrational (something that we have already noted is not clearly the case). The inferences from our being systematically irrational in various ways to our lacking ordinary knowledge and to our lacking scientific knowledge are both dubious. Here is where a modest response to the empirical evidence can provide us with succor from the skeptical conclusions one might attempt to draw from these studies.

[10] See Gilovich (1991) for discussion.

[11] See Fine (2006), Gilovich (1991), Kahneman (2011), Kahneman et al. (1982), and Nisbett and Ross (1980).

Rather than leading us to skepticism, recognizing that we are systematically irrational (if we really are) and prone to make certain common errors of reasoning should leave us optimistic about our prospects for both ordinary and scientific knowledge. The mere fact that we are in a position to evaluate our reasoning—we can determine that our reasoning in these experimental situations is poor—suggests that we tend to reason correctly a lot of the time. Plausibly, in order to recognize the mistakes that we are making in these sorts of cases *as mistakes* requires reasoning in the correct way. Further, these deviations from correct reasoning would not seem so jarring if we did not accept, and presumably take ourselves to reason in accordance with, the correct ways of reasoning. Consequently, it is reasonable to think that the fact that we recognize the results of these experiments as evidence for thinking we make errors in reasoning suggests that we often do reason correctly. In other words, "the focus on error does not denigrate human intelligence, any more than the attention to diseases in medical texts denies good health. Most of us are healthy most of the time, and most of our judgments and actions are appropriate most of the time" (Kahneman 2011, p. 4).

Another reason that we should be optimistic is that because of the fact that we recognize our susceptibility to various mistakes we can correct for those mistakes. Once we are aware of the effects of belief polarization, for example, we can take steps to temper our assessments of our viewpoints on controversial issues (Kelly 2008). The same is true of other cognitive biases which we are prone toward (Kahneman 2011). In fact, we often take steps to help catch biases and mistakes already. When we recognize that we might be prone to making a particular mistake we will double-check our calculations, ask others for their opinions, and so on. The fact that we are aware that we tend to fall prey to certain mistakes helps us to seek methods for avoiding those mistakes. It is much easier to avoid a pit that you see in the path ahead than it is to avoid one of which you are totally unaware. Recognizing the sorts of errors of reasoning we are likely to commit gives us hope that we can correct for our mistaken natural tendencies.

This sort of correcting or vigilance for avoiding cognitive errors is something that we seem to be particularly good at when we work with others. For instance, when David Moshman and Molly Geil (1998) conducted a version of the selection task experiment in which some participants were given the task individually and others worked in groups they found that groups did dramatically better. Similarly to the original selection task experiment, individuals working on the task only produced the correct answer about 9 % of the time. However, groups had a 70 % success rate with respect to the selection task. Groups that consisted of participants who had already tried the task as individuals did even better—they had an 80 % success rate. In general groups do much better at tasks for which individuals fail to give correct answers.[12] Discussion and interactions with others can help us to avoid errors and greatly increase our performance. Unsurprisingly, peer discussion methods have been found to be very effective for learning (Slavin 1995). Some go so far as to

[12] See Bonner et al. (2002), Laughlin and Ellis (1986), and Moshman and Geil (1998).

claim that various cognitive biases we commit as individuals are not only to be expected, but perhaps are conducive to greater performance on the whole, given the social nature of our reasoning practices (Mercier and Sperber 2011; Sperber and Mercier 2014). Given these findings, it is not surprising that in science we employ particular methods when conducting research—we rely on things like peer review, participate in discussions, and so on.[13] By doing so, we can compensate for, and maybe even take advantage of, various cognitive mistakes that we are prone to make as individuals and attain scientific knowledge in spite of our shortcomings.

13.3 Conclusion

We have seen that there are numerous experiments which seem to suggest that humans are systematically irrational. Although some might take this empirical evidence to show that we lack much of the ordinary and scientific knowledge which we take ourselves to have, this conclusion is premature. While there are difficulties for the responses which seek to establish that widespread irrationality of the sort that might be thought to undermine our knowledge is impossible, it is not obvious that such a response cannot succeed. Additionally, we have seen that there are some grounds for doubting the empirical evidence that is purportedly yielded by experiments on our reasoning. However, we have also seen that there is reason for thinking that it will be difficult to provide sufficient evidence for thinking that we should doubt all of the myriad evidence of our irrationality in the numerous experimental situations that have been tested.

Despite these somewhat dismal prospects for showing that we do not make systematic mistakes of reasoning, we have seen that there is cause for optimism about our prospects for knowledge. By recognizing our tendency to make certain systematic mistakes we can take steps to correct for these effects both as individuals and as groups. Thus, we have seen that while there is empirical evidence that we make various mistakes of reasoning, there is no reason to think that this means we cannot have scientific knowledge or that we lack the knowledge we ordinarily take ourselves to have. At most what we seem to have is further reason for thinking that scientific knowledge should be held tentatively. But, this is something to which we were already committed.

[13] We will discuss these issues further in the final part of the book when we examine some of the social aspects of scientific knowledge.

References

Bonner, B. L., Baumann, M. R., & Dalal, R. S. (2002). The effects of member expertise on group decision making and performance. *Organizational Behavior and Human Decision Processes, 88*, 719–736.

Cohen, L. J. (1983). Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences, 6*, 317–370.

Davidson, D. (1985). Incoherence and irrationality. *Dialectica, 39*, 345–354.

Dennett, D. C. (1978). *Brainstorms*. Cambridge, MA: MIT Press.

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

Feldman, R. (2003). *Epistemology*. Upper Saddle River: Prentice Hall.

Fielder, K. (1988). The dependence of the conjunction fallacy on subtle linguistic factors. *Psychological Research, 50*, 123–129.

Fine, C. (2006). *A mind of its own: How your brain distorts and deceives*. Cambridge, UK: Icon books.

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

Gilovich, T. (1991). *How we know what isn't so: The fallibility of human reason in everyday life*. New York: Free Press.

Goldman, A. I. (1986). *Epistemology and cognition*. Cambridge, MA: Harvard University Press.

Johnson-Laird, P., Legrenzi, P., & Legrenzi, M. (1972). Reasoning and a sense of reality. *British Journal of Pyschology, 63*, 395–400.

Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus, and Giroux.

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

Kelly, T. (2008). Disagreement, dogmatism, and belief polarization. *Journal of Philosophy, 105*, 611–633.

Laughlin, P. R., & Ellis, A. L. (1986). Demonstrability and social combination processes on mathematical intellective tasks. *Journal of Experimental Social Psychology, 22*, 177–189.

Lord, C., Ross, L., & Lepper, M. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology, 37*, 2098–2109.

Lycan, W. G. (1988). *Judgement and justification*. Cambridge, UK: Cambridge University Press.

Mercier, H., & Sperber, D. (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences, 34*, 57–111.

Millikan, R. G. (1984). *Language, thought, and other biological categories*. Cambridge, MA: MIT Press.

Moshman, D., & Geil, M. (1998). Collaborative reasoning: Evidence for collective rationality. *Thinking and Reasoning, 4*, 231–248.

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

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

Plous, S. (1993). *The psychology of judgment and decision making*. New York: McGraw-Hill.

Putnam, H. (1981). *Reason, truth and history*. Cambridge, UK: Cambridge University Press.

Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.

Skyrms, B. (1986). *Choice and chance* (3rd ed.). Belmont: Wadsworth.

Slavin, R. E. (1995). *Cooperative learning: Theory, research, and practice*. London: Allyn and Bacon.

Sober, E. (1981). The evolution of rationality. *Synthese, 46*, 95–120.

Sperber, D., & Mercier, H. (2014). Reasoning as social competence. In H. Landemore & J. Elster (Eds.), *Collective wisdom: Principles and mechanisms* (pp. 368–392). Cambridge, UK: Cambridge University Press.

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

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

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

Wason, P. (1966). Reasoning. In B. Foss (Ed.), *New horizons in psychology* (pp. 135–151). Harmondsworth: Penguin.

Wason, P. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology, 20*, 273–281.

Wason, P., & Johnson-Laird, P. (1972). *Psychology of reasoning: Structure and content*. Cambridge, MA: Harvard University Press.

Wason, P., & Shapiro, D. (1971). Natural and contrived experience in a reasoning problem. *Quarterly Journal of Experimental Psychology, 23*, 63–71.