Cognitive Illusions

A Handbook on Fallacies and Biases in Thinking, Judgement and Memory

Edited by
Rüdiger F. Pohl

Psychology Press
Taylor & Francis Group
HOVE AND NEW YORK

First published 2004 by Psychology Press
27 Church Road, Hove, East Sussex BN3 2FA

Simultaneously published in the USA and Canada
by Psychology Press
270 Madison Avenue, New York NY 10016

Psychology Press is a part of the Taylor & Francis Group

Typeset in Sabon by RefineCatch Ltd, Bungay, Suffolk
Printed and bound in Great Britain by MPG Books Ltd, Bodmin
Cover design by Sandra Heath

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Cataloging in Publication Data
Cognitive Illusions : a handbook on fallacies and biases in thinking, judgement and memory / [edited by] Rüdiger Pohl.—1st ed.
p. cm.
Includes bibliographical references and index.
ISBN 1-84169-351-0 (hardcover)
1. Errors. 2. Prejudices. I. Pohl, Rüdiger, 1954–

BF323.E7C64 2004
153.7′4—dc22 2004009834

ISBN 1-84169-351-0

Contents

4 Confirmation bias

Margit E. Oswald and
Stefan Grosjean

Creating and testing hypotheses represents a crucial feature not only of progress in science, but also in our daily lives in which we set up assumptions about reality and try to test them. However, the lay scientist stands accused of processing his or her hypotheses in such a way that he or she is biased to confirm them. "Confirmation bias" means that information is searched for, interpreted, and remembered in such a way that it systematically impedes the possibility that the hypothesis could be rejected – that is, it fosters the immunity of the hypothesis. Here, the issue is not the use of deceptive strategies to fake data, but forms of information processing that take place more or less unintentionally. In this chapter we are going to study how far the accusation of a confirmation bias in hypothesis testing is justified. But let us first try to solve two problems:

First, assume somebody presents you with the following task:

> I made up a rule for the construction of sequences of numbers. For instance, the three numbers "2–4–6" satisfy this rule. To find out what the rule is, you may construct other sets of three numbers to test your assumption about the rule I have in mind. I gave you one set of three already, and for every three numbers you come up with, I will give you feedback as to whether it satisfies my rule or not. If you are sure you have the solution, you may stop testing, and tell me what you believe the rule to be.

How would you proceed? Which sets of three numbers would you form in order to test your guesses? Please stop your reading at this point to answer this question. Thereafter, continue reading to find out whether you proceeded in manner similar to the participants in a study by Wason (1960).

Wason wanted to demonstrate in his study that most people do not proceed optimally in testing hypotheses. As to what is optimal, Wason followed the lead provided by the philosopher of science Popper (1959). According to Popper, the general mistake consists in trying to *confirm* a hypothesis rather than trying to falsify it. Participants in Wason's experiment typically proceeded in the following manner: Given the sequence of

three numbers "2–4–6", they first formed a hypothesis about the rule (e.g., "a sequence of even numbers"). Then they tried to test this rule by proposing more sets of three numbers satisfying this rule (e.g., "4–8–10", "6–8–12", "20–22–24"). The feedback to these examples was always positive ("This set corresponds to my rule"). After several rounds of such testing, many participants felt secure about their hypothesis, and stopped searching since they thought they had found the correct rule. However, when they had stated the rule they assumed to be correct, they were told that they were wrong. Actually, the rule was simply "increasing numbers". Now, since the set of numbers satisfying the rule as hypothesized by the participants represents a subset of all the possible sets of three numbers satisfying the correct rule, the testing strategy of the participants led to a spurious confirmation of their hypothesis. Their testing questions were all answered in the affirmative although the rule assumed was wrong.

What was the participants' mistake? Wason argues that their error consisted in failing to test sets of three numbers that did *not* correspond to what was assumed to be the rule. Thus, a sequence like "4–5–6" would have been an appropriate test. This is because it does *not* correspond to the rule assumed by the participant at this stage, and yet it prompts a positive feedback (since it does correspond to the correct rule). Thus, participants' assumptions about the rule would have been falsified. Therefore it is not surprising that those (very few) participants who generated sequences that did not correspond to their hypothesis (a negative test strategy) needed, on average, fewer rounds of testing than others to find the correct rule.

For the second problem, imagine that you are presenting two scientific studies on the effectiveness of the death penalty to people opposed to it and to people in support of it, as Lord, Ross, and Lepper (1979) did. One of the studies you present supports the conclusion that the death penalty has a deterrent effect, and thus lowers the crime rate. The other study contradicts the effectiveness of the death penalty. Additionally, Lord et al. gave their participants clues hinting at weak points in both studies (e.g., shortcomings of the cross-sectional or longitudinal surveys used). In your opinion, how will research findings that either confirm or deny the death penalty's deterrent effects be judged, and what impact will they have on the supporters and opponents of death penalty? Would you assume that the ambiguity of the findings will lead to greater relativity in both kinds of attitudes?

As the experiment demonstrated, participants gave higher ratings to the study that supported their own opinion, while pointing to shortcomings in the research that questioned their point of view. This pattern was observed even when both studies had supposedly been carried out using the same method. This kind of *confirmation bias* led to the remarkable outcome that participants were even more convinced of their original opinion after reading both studies than before.

From the starting point provided by the studies of Wason (1960) and Lord et al. (1979), we want to discuss in this chapter two questions that

represent a kind of thread weaving through the references on hypothesis testing. The first question refers to the procedure people typically use to *search* for information, as in the experiment by Wason (1960). Does this represent a confirmation bias at all, or should we concede, after careful investigation, that this procedure may not be optimal but is still a rather effective testing strategy? The second question refers to the conditions under which a confirmation bias occurs. Is the process of hypothesis confirmation caused by a strong desire for confirmation, or does it also happen in a "cold", that is, non-motivational, fashion? The assumption that individuals more or less constantly seek to confirm their hypotheses is widely shared and not only by lay people (see Bördlein, 2000). This tendency exists not merely because the possibility of rejecting the hypothesis is linked to anxiety or other negative emotions, as in the case of the death penalty (see above), but supposedly also because "cognitive processes are structured in such a way that they inevitably lead to confirmation of hypotheses" (Kunda, 1990, p. 494). However, if this were true, it would entail a serious challenge to the position that attributes to humans the ability to adapt effectively to changing environments by virtue of their evolution (Cosmides, 1989; Gigerenzer & Hug, 1992). We can thus see why different views on the question of hypothesis confirmation entail more profound epistemological problems. Until the 1960s this controversy was mainly the province of philosophers, including philosophers of science. The matter was not studied in psychology until it was introduced by Wason who, while still a student, had become fascinated by the philosophy of science taught by Karl Popper at the London School of Economics from 1946 to 1969.

CONFIRMATION BIAS OR POSITIVE TEST STRATEGY?

To be blunt, Wason (1960) was of the opinion that humans do not try at all to test their hypotheses critically but rather to *confirm* them. This position did not remain unchallenged. Klayman and Ha (1987), for example, presented an approach that disputed the view of humans as "hypotheses confirmers". They showed that the behaviour of the participant in the Wason experiment described above may be interpreted as a "positive test strategy" (see Text box 4.1). The positive test strategy is something different from the attempt to confirm hypotheses, or even to "immunize" them against rejection.

The positive test strategy (PTS) restricts the exploration space: A fundamental problem in testing hypotheses consists in the fact that exploration spaces can become very large if all the cases are considered that might be relevant to the hypothesis. A systematic search through the "whole universe" for events that could falsify the hypothesis can, from a pragmatic point of view, scarcely be accomplished. In this case, the PTS provides a heuristic that aids in restricting the exploration space to those cases that

Text box 4.1 A definition of the positive test strategy (PTS)

We propose that many phenomena of human hypothesis testing can be understood in terms of a general *positive test strategy*. According to this strategy, you test a hypothesis by examining instances in which the property or event is expected to occur (to see if it does occur), or by examining instances in which it is known to have occurred (to see if the hypothesized conditions prevail).

(Klayman & Ha, 1987, p. 212)

have some probability of being the relevant ones. Let us take, as an example, the hypothesis "John always becomes aggressive when he is provoked". To test this hypothesis, we could proceed according to the PTS. In this case, we would look for occasions on which John was provoked, in order to see whether he actually reacted aggressively. However, we could also look for occasions on which John did not become aggressive, in order to determine whether he had previously been provoked. Such a search strategy is called a *negative test strategy*, since the *non-occurrence* of the critical event (i.e., John has not been provoked) would confirm the hypothesis.

In this example, the choice of a PTS is very rational. It restricts the exploration space to relevant events without making us the victim of a confirmation strategy: It is always possible to identify occasions on which John was provoked without reacting aggressively. Without any doubt, *falsification* of the hypothesis would be possible. The choice of a *negative test strategy* could become very cumbersome, in this case. The number of cases where John does *not* react aggressively could be very large – moreover, it would not be very conclusive if it only revealed that he was not provoked on these occasions. Only occasions on which John did not react aggressively although he was provoked would count. However, if we hoped to find such events by considering all possible instances in which John did not react aggressively, we would be performing the equivalent of a search for a needle in a haystack. A PTS would enable us to find the "needle" much more easily since we would not have to search the whole "haystack". Here we search explicitly for occasions on which he was provoked – and, as said before, it could result in discovering that he did not become aggressive at all. Following the argument advanced by Klayman and Ha (1987), the PTS as such is not a biasing one. Depending on the structure of the task at hand, the PTS may well result in a falsification of the hypothesis. However, if our task is one in which all events satisfying the current hypothesis (rule) represent a genuine subset of the set of those events that satisfy the correct hypothesis (rule) searched for – as in the 2–4–6 experiment by Wason (1960) – then the PTS does lead to a confirmation bias.

Klayman and Ha (1987) are not the only authors to have critically

discussed the proposition that people only try to confirm their hypotheses. For example, Trope and Bassok (1982) note that we are justified to call a PTS a confirmation-seeking strategy only if the diagnosticity of the test is deficient. What does that mean? The assumption that a person is, for example, characterized by a certain trait such as honesty or introversion may be tested by asking questions referring to that trait. (For simplicity, we will restrict ourselves in the following to questions that can be answered by "Yes" or "No".) Now, questions that correspond to a PTS are worded in such a way that a "Yes" would confirm the hypothesis. Is it not thinkable that persons select such questions that lead to "Yes" answers independently of whether or not the hypothesis is true? Let us assume a woman who wants to test the hypothesis that her new partner is introverted. To do so, she asks the question "Do you read books occasionally?" and he replies "Yes", confirming her hypothesis. It becomes clear that this is a spurious confirmation when we admit that extraverts do also read books occasionally. Following Trope and Bassok (1982), we could call such a test strategy a "non-diagnostic strategy" since a "Yes" answer is likely both if the hypothesis is true and if its alternative is true. In contrast to this strategy, Trope and Bassok defined a "diagnostic strategy" (see Text box 4.2).

Text box 4.2 A definition of the diagnostic strategy

In this diagnostic strategy, the lay interviewer searches his or her stored representations of personality traits for behavioral features that are distinctively associated *either* with the hypothesized trait category *or* with its alternative(s).
(Trope & Bassok, 1982, p. 561)

Trope and Bassok (1982) ran experiments to solve the problem of whether persons prefer a positive over a negative test strategy even if the questions belonging to the PTS are non-diagnostic. Participants had to select questions in order to test a hypothesis about a person who was previously unknown to them. The result was that diagnostic questions were preferred in all cases, even if they represented instances of a negative test strategy. Devine, Hirt, and Gehrke (1990) replicated these results. Taken together, the studies referred to above suggest that persons prefer diagnostic questions but, whenever possible, they like to phrase their questions in the form of a PTS.

The experiment proposed below illustrates a third condition that may contribute to situations where the application of a PTS leads to a confirmation bias. In these cases, the person searching for information applies predominantly a PTS, and this biased way of interrogation influences the behaviour of the interaction partner. The experiment can be run in a classroom or lecture. It was inspired by the studies of Snyder and Swann (1978) and of Zuckermann, Knee, Hodgins, and Miyake (1995).

However, before we come to the proposal of a classroom experiment, let us summarize those cases in which the application of a PTS may lead to a confirmation bias (see Text box 4.3).

Text box 4.3 Cases in which the positive test strategy (PTS) may lead to a confirmation bias

1. If the correct hypothesis (rule) is *more general* than the one assumed by the person. In this case, necessarily all events corresponding to the hypothesis (rule) of the person also correspond to the correct hypothesis (rule) – as in Wason's (1960) 2–4–6 experiment.
2. If, in addition, the questions asked are *non-diagnostic* and very likely to be answered in the affirmative, independently of the truth of the hypothesis (Devine et al., 1990; Trope & Bassok, 1982).
3. If the interrogation behaviour influences the interaction partner in a way that he or she responds affirmatively (Snyder, Tanke, & Berscheid, 1977; Zuckerman et al., 1995).

AN EXPERIMENT ON THE INFLUENCE OF INTERROGATION BEHAVIOUR

Two groups of participants are supposed to test a social hypothesis (one group tests a target person, the interviewee, as to whether he or she is introverted, and the second group whether he or she is extraverted). Text box 4.4 provides the necessary procedural details to run an appropriate classroom study. The following results are expected:

1 The interviewers test the hypothesis assigned to them by means of a positive test strategy.
2 The interviewee behaves affirmatively (i.e., replies predominantly "Yes").
3 The interviewers in each group give significantly different ratings to the personality of the interviewee with respect to the introversion–extraversion dimension and in the direction of their respective initial hypotheses.

If the results come out as expected, they can be interpreted as follows: The participants applied a PTS in testing the social hypothesis. As we have seen above, this is not yet confirmation bias as such (since the interviewee could very well answer "No", thus rejecting or questioning the hypothesis). However, it is more probable that the interviewee responds affirmatively. According to Zuckerman et al. (1995), this may be due to a selective memory process (see below), or additionally to a social norm according to which "Yes" answers appear more friendly. Now, the confirmation bias occurs

Text box 4.4 Details of the classroom demonstration

Procedure

A person (e.g., an acquaintance of the experimenter) plays the role of the "interviewee". This person should be completely unknown to the participants. It is the task of the interviewee to answer the question put to him or her with a "Yes" or a "No". The interviewee must not hear the instructions given to the other participants.

Before giving the instructions, the class is split into two groups. One stays in the room while the other is sent to another room (not the same room as the interviewee). Both groups are given different instructions: One is asked to test the hypothesis that the interviewee is *extraverted*. To test this, all participants in the group are asked to write a common list of questions they are going to ask the interviewee, and these must be questions that can clearly be answered "Yes" or "No". The other group is asked to test the hypothesis that the interviewee is *introverted*. They, too, are asked to write down the questions to be posed to the interviewee later on, and again these must be clearly answerable with "Yes" or "No".

The participants are allowed about 10 minutes to construct and write down their questions. Thereafter, the interviewee is called into the classroom where only one of the groups is present. The interviewee is ignorant with respect to the instructions given to the two groups. He or she is told to answer questions exclusively with a "Yes" or a "No", and make no further explanations or statements – if necessary, the experimenter has to remind him or her of this instruction. Now the first group is to put questions they have devised and written down in advance. This task is limited to about 10 minutes. During this time, the experimenter records:

1. How many questions are asked, and how many of these are instances of a positive test strategy (all questions that are phrased in such a way that a "Yes" reply would confirm the hypothesis count as instances of a positive test strategy).
2. How often the interviewee replies "Yes".

After the time is up, the experimenter interrupts the interview (whether or not all the questions have been put), and asks the participants to rate the interviewee on a 7-point rating scale from –3 to +3, marked at the ends as "introverted" and "extraverted". The scale may be projected as a transparency by an overhead projector, and the participants are asked to write down the number they would assign.

Following this, the other group is called back to the classroom. Now it is their turn to ask the interviewee the questions they have prepared. The first group may remain present but is urged not to talk during this procedure. The process is exactly the same as before. Finally, the second group is also asked to rate the interviewee on the introversion–extraversion scale.

Thereafter, the following three dependent variables are calculated: (1) the proportion of questions that correspond to a positive test strategy in each

group; (2) the proportion of "Yes" replies by the interviewee in each group; (3) the mean ratings of the interviewers on the introversion–extraversion scale in each group. To analyze the data of the classroom experiment proposed above, chi-square tests with $df = 1$ for 2 × 2 contingency tables could be applied, for Variable 1 with two columns for the groups and two rows for PTS questions and non-PTS questions, and for Variable 2 with two columns for the groups and two rows for "Yes" and for "No" replies. For Variable 3, the difference between the mean ratings of the two groups could be tested by means of a *t*-test for independent samples, with $df = (n_1 + n_2 - 2)$.

Results

Based on results of similar studies (e.g., Snyder et al., 1977; Zuckerman et al., 1995), the main findings should be as follows:

1. The questions asked correspond predominantly to a positive test strategy (here: Significantly more than half of the questions are phrased in such a way that an affirmative reply would confirm the respective hypotheses).
2. The interviewee responds in the affirmative (here: More than half of the replies are "Yes").
3. The interviewers treat the responses of the interviewee as a confirmation of their hypothesis (here: The group testing the extraversion hypothesis rates the interviewee as significantly more extraverted than the other group testing the introversion hypothesis).

because the interviewer is not aware of this acquiescence tendency – and thus derives unjustified conclusions from the replies of the interviewee. This way a "self-fulfilling prophecy" is generated.

However, we should also consider possible methodological problems: Could it be that questions about extraverted forms of behaviour might differentiate better between introversion and extraversion than do questions about introverted behaviour (see Devine et al., 1990)? Does the result depend on how introverted or extraverted the interviewee is in reality? And does it make sense to test the target person's introversion/extraversion in advance?

EVIDENCE FOR A "TRUE" CONFIRMATION STRATEGY

Since Klayman and Ha (1987), several other authors have postulated that PTS as such is not a confirmation bias. In spite of this, an astonishing confusion is still to be found in the literature. There are two possible reasons for this: (1) Wason (1960, 1968) publicized the concept of "confirmation bias" in his original studies, and therefore this phrase is often employed although what is actually intended is a PTS (see Devine et al., 1990). In those cases, though,

some authors additionally note that besides the PTS there also exists a "true confirmation bias" (see Poletiek, 2001). (2) Some authors seem to have real problems differentiating between a confirmation bias and a PTS. According to the PTS, persons have the tendency to ask questions in such a way that their hypothesis *would* be confirmed *if* the answer was affirmative. Those authors seem to miss the conditional clause here and transform it instead into propositions like: "Persons have the tendency to seek only for confirmatory evidence" (Doherty, Mynatt, Tweney, & Schiavo, 1979, p. 113). Although the two statements differ only in a few words, they mean totally different things.

The discussion so far has shown that it is not at all easy to demonstrate a confirmation strategy because we always have to prove that the respective procedure *systematically impedes a possible rejection* of the hypothesis. For this reason, it may be justified to ask if it is possible at all to decide in an unambiguous way whether a certain strategy represents a confirmation strategy. We intend now to elaborate on this question, and will include strategies in addition to those concerned with information search (see Text box 4.5).

Text box 4.5 Three ways of introducing confirmation bias

In addition to strategies applied in the *search* for new information (information-gathering strategies), there are other possibilities not yet discussed here for immunizing hypotheses. In the process of remembering (information-recollection strategies) people may selectively *recall* mainly such information that would confirm their hypotheses. This contributes to a considerable degree to a resistance to change. Also, in the process of interpretation (strategies of information interpretation), persons may systematically *re-interpret* existing information contradicting their hypothesis, or attribute less importance to it than to confirming information, in spite of the objectively equal value of each kind of information.

Based on the discussion presented so far, we may talk of a true confirmation strategy if we can show that, in the search for information, the test (the question asked) will very likely confirm the hypothesis (i.e., be answered affirmatively), and does so independent of the truth of the respective hypotheses. Studies previously undertaken have demonstrated, however, that persons – even children (see Samuels & McDonald, 2002) – are seldom so foolish as to apply such a worthless search strategy.

A true confirmation strategy is also involved if another person's opinion is already known and if it can be assumed that questions will be answered according to one's expectation. Thus, in daily life people often ask

like-minded acquaintances for their opinion. Whether, and under what conditions, people systematically search for such sources of information already knowing that a confirmation of their initial assumption is likely, has been little studied to date. The exceptions are those studies inspired by *dissonance theory* in which the phenomenon of "post-decisional regret" has been explored (Frey, 1981; Gagné & Lydon, 2001). According to this theory, people are inclined to avoid sources of information that could question the quality and correctness of a decision that, once made, they are hardly likely to revoke, for example, the purchase of an automobile. However, in such cases it is clear that individuals are no longer neutral with respect to the outcome of their search for information, that is, their starting point is not a "cold" hypothesis. But even in emotionally less involving situations people may search for confirming information, as Betsch, Haberstroh, Glöckner, Haar, and Fiedler (2000) demonstrated in a recent study. If participants have established a *decision routine* in one task, and are expected to test this procedure in another context, they are inclined to maintain this routine. They search mainly for information that can a priori be assumed to favour the routine applied so far.

With respect to *selective remembering*, it has to be said that in recall there seems to be no general advantage for hypothesis-confirming information over hypothesis-contradicting information. Many approaches, inspired by schema theory (Neisser, 1976), assume that schema-consistent information, that is, information consistent with expectations, is not only encoded more easily but also recalled more easily than inconsistent information (Taylor & Crocker, 1981). Other approaches, such as Woll and Graesser's (1982) "schema pointer plus tag model", emphasize that hypothesis-contradicting information will be particularly salient, and thus may be processed more extensively or in a more detailed manner. This kind of processing might make expectancy-disconfirming information particularly likely to be encoded and remembered – perhaps even better than expectancy-congruent material. For both approaches there exists extensive empirical evidence, as is clearly demonstrated in the meta-analysis published by Stangor and McMillan (1992).

Whether the results point to a general advantage in recalling consistent information (congruency effect) or inconsistent information (incongruency effect) seems to depend essentially on the method applied to measure recall. Moreover, an advantage for the recall of consistent information seems to depend on certain additional conditions: (a) Individuals have already established hypotheses, (b) these hypotheses refer to social groups (stereotypes) and not to individuals, and (c) there exists a temporal delay between the processing of stimulus information and the recall or judgement. How far the advantage in recall of congruent data is related to the fact that the hypotheses are motivationally supported was not revealed by this particular meta-analysis.

However, several studies demonstrate that selective, directed recall occurs

when people consider a particular personality trait as especially desirable. Sanitioso, Kunda, and Fong (1990) convinced their students in one of the conditions in their experiment that extraverted persons are particularly successful in their academic and professional careers, and in the other condition that this was instead true for introverted persons. In addition, participants were asked to think of possible reasons for this relation between personality and success. Later on, in a seemingly unrelated study, they were asked to list autobiographical memories reflecting their own standing on the introversion–extraversion dimension. The introvert-success induced participants (a) were more likely to list introverted memories first, (b) generated introverted memories faster, and (c) tended to list overall more introverted memories than did the extravert-success induced participants. The "search tends to be biased, so that memories consistent with the desired trait are more likely to be accessed than memories that are inconsistent with it" (Sanitioso et al., 1990, p. 239). The authors could not find any comparable effect when the two personality traits were instead activated by semantic priming. The enhanced accessibility seems to be due to motivational factors rather than to priming of memory.

The other possible means to confirm hypotheses identified above refers to the case where a hypothesis tester gathers information (strategy of information interpretation) and biases the interpretation of that information so that the hypothesis appears to be true. Systematic preservation of the original hypothesis could occur by assessing the expectancy-congruent information systematically to be *more important* than the incongruent information, or by *increasing confidence* in the hypothesis on the basis of congruent information more than such confidence is decreased by incongruent information (cf. Bacon, 1620/1990, p. 109). Of course, here we must assume that the two kinds of information do not differ in their diagnostic relevance.

A difference in the weighting of congruent and incongruent information was clearly demonstrated in the already cited study on one's attitude towards the death penalty (Lord et al., 1979). Participants evaluated information incongruent with their attitude much more critically than congruent information even when both kinds were acquired by the same method. Of course, we can assume in this study that attitudes towards the death penalty are not emotionally neutral, so that the confirmation bias may have been caused motivationally. But it can also be shown with fairly "neutral" hypotheses that data consistent with the hypothesis are accorded a higher weight than inconsistent data (Gadenne & Oswald, 1986; Slowiaczek, Klayman, Sherman, & Skov, 1992). In a study by Gadenne and Oswald (1986), for example, participants were told a crime story in which a theft is committed under circumstances to be clarified. Thereafter, the participants were asked to rate the importance of statements (of medium diagnosticity) pointing to or arguing against the possibility that a certain person, A, committed the theft. Information implicating this suspect was significantly more strongly weighted if participants had already adopted the hypothesis that

A was the offender (such adoption being induced in the experiment) than if this was not the case. Given information exonerating the suspect, the induction of a hypothesis about the identity of the offender had no effect on the weighting of the evidence.

It is not yet clear what causes the bias in weighting of data congruent with the hypothesis. Some evidence suggests that participants confronted with an unexpected event may well ask whether this might be "the exception to the rule", while they simply accept expected events without further questioning as to whether the event might also be compatible with the alternative hypothesis (Kruglanski & Mayseless, 1988). However, as long as the alternative hypothesis itself, or the possibility that the event could be explained by an alternative hypothesis, is not considered, an overestimation of the diagnostic relevance of events congruent with the hypothesis occurs very rapidly.

CAN WE ALWAYS ASSUME A CONFIRMATION BIAS IN THE CASE OF MOTIVATIONALLY SUPPORTED HYPOTHESES?

The analyses reported so far made clear that persons do not principally proceed in a confirmatory fashion when testing a hypothesis. Does this also hold for motivationally supported hypotheses, that is, those hypotheses with respect to which there are positive or negative emotions depending on the outcome of the test? Can we assume that confirmatory strategies are in general applied to testing hypotheses, given the assumption that people are generally motivated to seek positive emotions, and avoid negative ones? Although this will often be true – many studies (e.g., Trope & Liberman, 1996, p. 258) and our daily life provide corroboration here – we will see nonetheless that the answer is not so simple. Obviously, it is not always possible to simply believe what we want to believe.

In this context, Scott-Kakures (2001) drew attention to those many situations in daily life where just the opposite phenomenon occurs, namely a tendency to confirm *undesired* or unwelcome assumptions. Thus, people are not infrequently inclined towards intensive testing of assumptions like "I forgot to turn off the tap in the kitchen when I left home", "The strange red spots on my back might be an indication of cancer", or "My daughter Sabine overestimates her capabilities and will endanger herself mountain climbing". In such cases, it may well be that those events are systematically remembered, or that information is searched for which would increase the probability of the undesired hypothesis rather than decrease it. This tendency to seek confirmation of the negative, however, is incompatible with the general statement that hypothesis confirmation occurs because the hypothesis is desired.

If the confirmation of a hypothesis is desired or is associated with positive emotions, persons will be motivated to use a confirmation strategy. However, they will not do so at the risk of a *spurious confirmation* if this error