

The Interpretation of Cognitive Feelings

Dissertation

zur Erlangung des akademischen Grades eines Dr. phil.

der Fakultät für Verhaltens- und Empirische Kulturwissenschaften

der Ruprecht-Karls-Universität zu Heidelberg

vorgelegt von

Christian Unkelbach

aus Bretten

Heidelberg, November 2004

Gutachter:

Prof. Dr. Klaus Fiedler (Betreuer)

Prof. Dr. Herbert Bless

Acknowledgements

The present work was made possible by the help and support of many friends and colleagues, and I wish to express my sincere gratefulness:

To my research assistants, Liesa Büche, Marina Kühne, Viola Schreiber, Martin Stegmüller, Martina Wilke and especially, Myriam Bayer, who provided feedback throughout the whole process, conducted the experiments and spotted the largest holes in my reasoning.

To Peter Freytag, who meticulously read, corrected and improved an earlier draft of this thesis.

To Arie Kruglanski, who sparked the initial idea for the second part of the present thesis.

To Herbert Bless, who was my first contact to “the message within”.

To Henning Plessner, who always beat me in our tennis matches, but in exchange, also gave me most valuable insights.

To all the rest of the Crispies; for many fun lunches, coffee table discussions and shared experiences.

To my advisor, Klaus Fiedler, who for over the last 3 years has been providing guidance by stimulating discussions, by giving me the chance to present my work at many conferences, and by being a grindstone to sharpen hypotheses and theories. Most importantly, he provided the freedom to develop and sometimes to discard my own ideas. Thanks, Klaus!

The research was supported financially by the Deutsche Forschungsgemeinschaft (DFG) grant Fi 294/21, awarded to Klaus Fiedler, and the Deusch-Israelische Projektkooperation (DIP) grant D4.2.

Table of Contents

<u>INTRODUCTION</u>	<u>2</u>
<u>STRUCTURING THE BACKGROUND: ASSUMPTIONS ABOUT COGNITIVE FEELINGS</u>	<u>3</u>
TWO CLASSIC PARADIGMS FOR INVESTIGATING COGNITIVE FEELINGS	5
EXPERIENCING, ATTRIBUTING AND INTERPRETING COGNITIVE FEELINGS	6
EXPERIENCING COGNITIVE FEELINGS	6
ATTRIBUTING COGNITIVE FEELINGS	7
INTERPRETING COGNITIVE FEELINGS	7
<u>PART I – EXPLICITLY PROVIDED INTERPRETATIONS</u>	<u>10</u>
THE GENERAL PARADIGM OF PART I	12
<u>EXPERIMENT 1</u>	<u>14</u>
METHOD	16
RESULTS	18
DISCUSSION	22
<u>EXPERIMENT 2</u>	<u>24</u>
METHOD	25
RESULTS	26
DISCUSSION	29
<u>SUMMARY OF EXPERIMENT 1 AND 2</u>	<u>30</u>
<u>PART II: INTERPRETATIONS LEARNED VIA ECOLOGICAL FEEDBACK</u>	<u>32</u>
THE GENERAL PARADIGM OF PART II	34
<u>EXPERIMENT 3</u>	<u>38</u>
METHOD	38
RESULTS	42
MENTAL ROTATION TASK	42
RECOGNITION TEST – RESPONSE LATENCIES	44
RECOGNITION TEST – SDT-ANALYSIS	45
DISCUSSION	49

EXPERIMENT 4	51
METHOD	52
RESULTS	53
MENTAL ROTATION TASK	53
RECOGNITION TEST - RESPONSE LATENCIES	55
RECOGNITION TEST - SDT-ANALYSIS	56
DISCUSSION	59
GENERAL DISCUSSION	61
A SINGLE DIMENSION?	62
IS THE DISTINCTION OF ATTRIBUTION AND INTERPRETATION NECESSARY?	63
UNCONSCIOUS VS. CONSCIOUS PROCESSES	65
LATENCY AS AN INDEX FOR FLUENCY	67
THE POSSIBLE INHERENT MEANING OF COGNITIVE FEELINGS	69
FORTHCOMING RESEARCH QUESTIONS	70
CONCLUSION	72
REFERENCES	74
SUMMARY	84

The Interpretation of Cognitive Feelings

Introduction

“We are better prepared to study the content of thought than the experience of thinking.”
(Clore, 1992; p. 133)

Imagine you were asked to judge the difficulty of a mathematical problem, for example, the equation $2 + 2 = x$. Your immediate reaction might be that this is a very easy task, probably not even worth the label “problem”. But on what evidence do you base your judgment? A justification why this is considered to be an easy task is not complicated: The mathematical equation consists only of two single digit natural numbers on which the operation of summation is performed to reach the conclusion that $x = 4$. This, however, may not be the reason why you came up with the judgment that this is an easy task. Rather, you might have used the experience that it is easy to solve the equation, that you instantly knew the solution. This experience-based route to judgments, decision making and inferences will be the theme of the present dissertation. It is about the experience of thinking rather than about the content of thought.

The analytic approach to judgments, decisions and inferences (here, two single digits involving only summation) has been the dominant model in cognition research for quite some time, explicitly or implicitly relying on a computer metaphor. One constraint in this tradition has been that the computer metaphor does not account for the role of affect in judgments (e.g., Mellers, Schwartz, & Ritov, 1999). Human beings are not cold information processing units; emotions, moods, and feelings have a profound influence. Yet, the computer metaphor can incorporate such affective states by assuming that they enter judgments as information (e.g., Schwarz & Clore, 1983; Zajonc, 1980), by assuming that they have a systematic influence on what information is considered (Bower, 1981; Forgas, 1995), or by assuming that they have a systematic influence on how information is integrated (Fiedler, 2000; Forgas, 2000).

Besides this obvious constraint, there is another conceptual shortcoming in the computer metaphor. A computer has no monitoring system for its processes. For instance, a computer will work through implemented algorithms without realizing a difference between a simple and a complex operation. It will solve the most complex iterations as well as the simple task of $2 + 2 = x$. A computer cannot use the information how long a calculation takes, how fast something is accessed on its hard drive, or how long it takes to download information from a distant source on the internet. Human beings do. The simple task of judging the difficulty of $2 + 2 = x$ should serve as an illustration; as said, you probably did not think about the quality of the problem analytically. Instead, you used the experience resulting

from your own cognitive processes. Following the terminology of Clore et al. (2001), such experiences will be called *cognitive feelings*. The underlying rationale is that the term feeling is a “generic designation for all kinds of internal signals that provide consciously available feedback from non-conscious affective, bodily, or cognitive processes” (Clore et al., 2001, p. 30).

The impact of such cognitive feelings on judgments is for example demonstrated by Tversky and Kahneman’s well-known availability heuristic (Tversky & Kahneman, 1973). They stated that frequency or probability estimates are based on the ease with which instances of the event come to mind. For example, if asked to judge the frequency of words starting with the letter “r” compared to words with “r” as the third letter, participants largely overestimate the first frequency and underestimate the latter (Tversky & Kahneman 1973, Exp. 3). This is supposedly the case because participants use their subjective experience, their feeling, that it is easier to come up with words starting with “r” compared to words with “r” as the third letter. Words starting with an “r” are more available. But what is the underlying process? The use of the subjective experience of ease in judgments of easiness or difficulty is reasonable (e.g., in judging the difficulty of $2 + 2 = x$), but why does ease influence frequency estimates? There is a silent assumption of an inferential process of the kind that if instances come easy to mind, they must be frequent. The ease is interpreted in terms of frequency. So, within the area of research on the experience of thinking, the interpretation of cognitive feelings will be the main subject of the present thesis.

However, before a model for this interpretation process is presented, a discussion of the underlying assumptions, a distinction of cognitive and affective feelings and a general stake out of the presented research seems is called for.

Structuring the background: Assumptions about cognitive feelings

The influence of feelings on judgments in general has been studied mostly in the area of affective feelings, and there, mostly on a positive – negative valence dimension. The equivalent in the domain of cognitive feelings is an easy – difficult or a fluent – non-fluent dimension. Although there are a number of cognitive feelings like revelation (Watkins & Peynircioglu, 1990), surprise (Reisenzein, 2000; Whittlesea & Williams, 2001), tip-of-the-tongue (Hart, 1965; for a review, see Schwartz, 2002), they are all derivable from this dimension. For example, the revelation effect (Bornstein & Neely, 2001; Bernstein, Whittlesea, & Loftus, 2002) is easy to describe as the experienced difficulty of a cognitive

process followed by sudden unexpected ease, which is termed revelation. The following discussion will therefore concern the ease or fluency of cognitive processes.

Subjective experiences of fluency or ease can result from conceptual processes (e.g. memory retrieval, information integration, memory storage) or perceptual processes, but the distinction is blurred. For instance, on a biological level, the first integration of information takes place on a sub-retinal level. From a definitional point, it is difficult to tell where a perceptual process ends and a conceptual process starts. As an illustration, it is a daily life phenomenon that people are able to react when they hear their own name, even if they listen intensively to something else, for example, when they are engaged in a conversation on a cocktail party (Cherry, 1953). This phenomenon is hard to explain if a sharp distinction of conceptual and perceptual, bottom-up and top-down processes, is assumed. And logically, it is difficult to construe conceptual processes that are devoid of perception and vice versa. Therefore, conceptual and perceptual processes will be treated as interchangeable. On an operational level, fluency and ease will be indexed by the speed of responses. It is hard to test the functional similarity of speed and fluency, but it makes sense intuitively (Reber, Wurtz, & Zimmermann, 2004): Any process that is fluent or easy should be faster than a non-fluent or difficult process, be it conceptual or perceptual.

Furthermore, the assumption is made that judgments are based on available information at the time of judgment (Fiedler, 2000); thus, judgments and decisions based on existing or preformed opinions are not in the focus of the following. To relate again to the example from the beginning: If you already have an opinion available whether $2 + 2 = x$ is easy, you need not use your experience that it is easy to solve the equation; you can just use your existing opinion. A similar assumption is found in research on the influence of affect in judgments and decisions. If opinions and judgments already exist, there is less room for systematic influences of affective feelings (Forgas, 1995).

Yet, if the judgment is constructed, feelings might influence what information is used, how it is used or feelings might enter the judgments as direct input. For example, Schwarz and Clore (1983) showed that people use their mood as input for judgments of how satisfied they are with their lives. This offers the analogy for the impact of cognitive feelings on judgments: If an object evokes positive affect, the evaluation of this object will be positive. In general, evaluative judgments depend on evaluative information, and affective feelings offer a direct source of such evaluative information. The same is true for the presented math problem; the cognitive feeling of ease offers a direct source of information for a judgment about the easiness of a task: Judgments of ease depend on, well, easiness information.

Cognitive feelings, however, influence not only judgments of ease or difficulty. To show this, two classic paradigms will be introduced.

Two classic paradigms for investigating cognitive feelings

So far, it is not surprising that a feeling of ease influences judgments concerning easiness. But cognitive feelings do not only influence ease judgments directly; the experience of thinking *interacts* with the content of thought (Schwarz, 1998). The classic study for this effect is related to judgments of self-assertiveness (Schwarz et al., 1991), and was originally designed to disentangle the amount of recall from the experience of recalling that was present in the original experiments on the described availability heuristic by Tversky and Kahneman (1973). To do this, participants were asked to recall either six or twelve instances of their own behavior. In one condition, they were asked to recall assertive behaviors, in another condition, they were asked to recall non-assertive behaviors. Recalling six assertive instances led to higher judged self-assertiveness than recalling twelve assertive instances, and vice versa for non-assertive instances. The presumed logic is that it is easy to recall six instances, independent of the type of behavior, whereas it is difficult to recall twelve instances. Supposedly, the ease of retrieval was interpreted as frequency or typicality of the behavior.

Following from this example, one major difference from affective feelings is clear: Influences of affective judgments result in main effects of the valence of the affect; that is, affective feelings give immediate evaluative information, which can be termed „How do I feel about it“. However, cognitive feelings can result in interaction effects; here, the content that is recalled interacts with the experienced ease of recall. In the given example, the underlying rationale is that the experience of ease is interpreted as frequency or typicality. If negative information seems to be more frequent or typical, an evaluative judgment is more negative than when negative information seems to be seldom or atypical. The reverse is true for positive information.

The second classic paradigm was introduced in a study by Jacoby, Kelley, Brown, and Jasechko (1989; see also Jacoby, Woloshym, & Kelley, 1989). They had participants judge whether names on a list are famous or not. They manipulated the fluency of names by having participants read non-famous names from a test list of names 24 hours prior to the actual task. Non-famous names that were more fluently processed due to prior presentation had a higher probability to be judged as famous. Presumably, participants interpreted the enhanced fluency of the names as fame.

Thus, whereas affective feelings have an immediate and inherent meaning, cognitive feelings are malleable. In one study, they influence judgments of frequency or typicality, in

another study, they influence judgments of fame. In addition, cognitive feelings influence judgments of clarity (Whittlesea, Jacoby, & Girard, 1990), recognition judgments (Jacoby & Witherspoon, 1982), truth judgments (Reber & Schwarz, 1999), judgments of learning and feelings of knowing (Koriat, 1993, Koriat, Goldsmith, & Pansky, 2000).

This leads back to the question that is at the heart of the presented thesis: How does the same construct (i.e., fluency/ease) influence such a broad variety of judgments? Why is ease or fluency differentially interpreted, for example, as frequency, typicality, or fame? Before turning to this main question of interpretation in more detail, there are yet two more propositions that need discussion: The experience of cognitive feelings and the attribution of cognitive feelings.

Experiencing, attributing and interpreting cognitive feelings

Experiencing cognitive feelings

The first assumption is that people need to experience a cognitive feeling for this feeling to have an impact on judgments. This is not as trivial as it sounds. Cognitive processes take place all the time, we continuously perceive, store, retrieve or integrate information. Which processes, however, give rise to an experience, to a cognitive feeling and which do not? How do we get around the lurking homunculus that watches cognitive processes, serving as a control unit and deciding when to allow an experience? The dilemma is solved by assuming that all cognitive processes result in a subjective experience. Koriat and colleagues (Koriat, 1993; Koriat & Goldsmith, 1996) have used the term “parasitic” to describe the effect that cognitive processes cause subjective experiences as by-products. The key point is whether you notice the subjective experience; a cognitive feeling arises when there is a noticeable difference in the ongoing flow of experience (Whittlesea & Williams, 1998). Processes execute either more easily than expected or more difficultly than expected. For example, the cognitive feeling of familiarity is presumably due to a prior encounter contact with an object. Even when we cannot explicitly remember the object, we often have a vague feeling of familiarity (on the distinction of such “remember” and “know” memory processes, see Kelley & Jacoby, 2000; Gardiner & Richardson-Klavhen, 2000). Yet, why do we experience a feeling of familiarity when we see a face in the crowd on the street, but not when we see a family member at our breakfast table (Whittlesea & Williams, 1998)? It is the discrepancy that makes the experience noticeable. Discrepancies can result from comparisons with expected standards or from comparison with the ongoing flow of experience. The assumption of such a continuous monitoring and evaluation of cognitive processes and the detection of discrepancies is most explicitly made in the SCAPE (Selective Construction And

Preservation of Experience) model by Whittlesea and Leboe (2000; see also Whittlesea & Price, 2001; Whittlesea, 2002). Hansen and Wänke (2004) tested this assumption for ease of retrieval effects. In their study, participants had to generate few (which is easy) or many (which is difficult) arguments supporting a given attitude. Furthermore, they manipulated participants' expectation about the difficulty of the generation task. The standard ease of retrieval effect is that generating few arguments leads to greater change in attitude than generating many arguments (Wänke, Bless, & Biller, 1996). In this study, this was only true when participants' expected ease was discrepant from the experienced ease; that is, when they expected the task to be easy and it was actually difficult or when they expected the task to be difficult and it was actually easy. If the expectation matched the task requirements, no ease of retrieval effects were found. This further supports the notion that it is indeed the discrepancy in the flow of experience that results in cognitive feelings and their influence on judgments and inferences.

Attributing cognitive feelings

The second assumption is about the attribution of the experience; to have an impact on a given judgment or decision, the subjective experience must be attributed to the cognitive process that is involved in the judgment. This again is analogous to the influence of affective feelings on judgments: In the described study about the influence of mood on judgments of personal well-being (Schwarz & Clore, 1983), a reminder of the weather cancelled the effect of mood on the judgment, supposedly because a good or bad mood was attributed to the pleasant or unpleasant weather conditions. The same effect is shown for cognitive feelings in the Schwarz et al. (1991, Exp. 3) study; if the subjective experience of ease is attributed to an external source, in this particular case, the purportedly facilitating or inhibiting effects of music on the recall of autobiographic memories, the ease of retrieval effect vanishes. Applied to the examples of math problems this means: If you cannot solve a problem right away, you might judge it difficult, but not when the difficulty is caused by a noisy environment, which keeps you from concentrating. In other words, if a cognitive feeling is attributed to an irrelevant source, if it is deemed as non-diagnostic, the feeling should not be used as information in judgments.

Interpreting cognitive feelings

Cognitive feelings have an effect on judgments when they are experienced and when they are not attributed to an irrelevant, external source, that is, when they are diagnostic for a given judgment. The analogy in this respect to affective feelings is striking. Provided that a cognitive feeling is indeed experienced and it is deemed diagnostic, how is a cognitive feeling

interpreted? In the case of the Schwarz et al. (1991) study, why do we interpret ease of retrieval as frequency; in the case of the Jacoby et al. (1989) study, why do we interpret fluency as fame?

A first possibility is to follow the analogy of affective and cognitive feelings all the way and assume that the employed manipulations resulted in very distinct experiences and therefore in specific effects on given judgments. This implies that cognitive feelings have an inherent meaning and different manipulations result in distinct experiences like familiarity, feelings of knowing or ease. The meaning of a cognitive feeling could be hard-wired, acquired phylogenetically, much the same way as we acquire the meaning of affective feelings: We do not need to interpret anger, happiness or sadness.

However, even with the exact same manipulation, it is possible to obtain differential effects of cognitive feelings. In a study by Kunst-Wilson and Zajonc (1980), participants were subliminally shown geometric shapes. They were then asked which of two presented shapes they preferred. The pairs always consisted of a new shape and a shape that was subliminally presented before. Overall, they found a preference for the presented shapes, whereas the recognition of the shapes was at chance level. Again, the underlying logic is that prior presentation, even subliminally, facilitates the processing of a given stimulus, which leads in turn to a higher preference for the respective stimulus. Mandler, Nakamura, and van Zandt (1987) replicated the experiment, but in addition to ask for preference, participants in two other conditions were also asked which of the shapes seemed darker or brighter. In all three conditions, participants not only judged the previously subliminally presented shapes as preferable, but also as darker and also as brighter. Therefore, the cognitive feeling of facilitated processing was nonspecific and needed interpretation, which was provided by the question asked.

In another study, Whittlesea, Jacoby and Girard (1990) manipulated the visual clarity of words and participants had to judge whether a word had been presented previously or not. Items with a higher visual clarity had a greater chance to be judged as old, compared to items with lower visual clarity. That is, the enhanced fluency of processing due to greater visual clarity was interpreted as familiarity (i.e., to the fact that one had encountered this item before). The interesting result for the present argument is that prior presentation also influenced judgments of visual clarity. Items that had previously been presented were judged to be higher in visual clarity. This symmetric finding supports the presumed non-specificity of cognitive feelings.

These examples suggest a second possibility: Cognitive feelings deviate from affective feelings in such that they have no inherent meaning; rather, they result from an unspecific

experience that maps onto a easy – difficult dimension. To have an impact on judgments, this unspecific experience needs interpretation. This interpretation in turn leads to judgments that an item is frequent, true, typical, familiar, famous or easy to remember. The concept is not original, but stems from yet another analogy to affective feelings. It was formulated by Schachter and Singer (1962) to explain the multitude of emotions. In short, they proposed that emotions (i.e., affective feelings) result from a cognitive interpretation of an unspecific physiological arousal. The same assumption can explain the multiple effects of cognitive feelings; the continuous monitoring of ongoing cognitive processes maps onto an unspecific easy – difficult dimension which is then interpreted according to the present context. For example, in the study by Schachter and Singer (1962), the context was provided by a confederate who behaved either angrily or happily, leading to an interpretation of the unspecific arousal induced by an adrenaline injection as anger or elation.

As is clear from the previous paragraphs, the latter possibility is advocated here. It is proposed that the impact of cognitive feelings results from the interpretation of unspecific experiences and the interpretation itself is provided by the context. This also incorporates the necessity in the studies by Mandler, Nakamura, and van Zandt (1987) and Kunst-Wilson and Zajonc (1980) that recognition of the presented stimuli is at chance level. If a stimulus is recognized, a different interpretation of the cognitive feeling is provided. The unspecific difference in ease/fluency cannot account for two effects at one time; either the difference results from prior exposure, that is, you recognize them, or it results from your subjective liking, the greater brightness, or the greater darkness of the stimuli. Once an interpretation is made, the experience cannot account for another effect. Similarly, if you have difficulties to come up with instances of your own self-assertiveness, you can conclude that there are not many such instances, or that your autobiographic memory in general is bad, or that the background music inhibits your performance; but you would not interpret the difficulty in all three ways. The conjunction for different interpretations of cognitive feelings is always “or” and not “and”.

The following empirical section, supporting this interpretation model of cognitive feelings, is divided into two parts: Part I will deal with the deliberative interpretation of cognitive feelings; that is, when people use explicit interpretations provided by their environment to infer the meaning of an experiential state. Two experiments will present data from an ease of retrieval paradigm which show that ease of retrieval effects are indeed dependent on explicitly provided interpretations. Part II will deal with the case when interpretations are not explicitly provided, but when interpretations are acquired by a simple learning mechanism, following ecological feedback in a given context; for example, that

instances which come easier to mind are indeed more frequent or more probable. Part II will present two experiments that employ a memory paradigm that lends support to the notion that the interpretation of cognitive feelings can be acquired via such simple ecological feedback. All four experiments in the empirical section aim to show that subjective experiences resulting from cognitive processes have no inherent meaning but depend on the interpretation of these cognitive feelings.

PART I – EXPLICITLY PROVIDED INTERPRETATION

Given that a person experiences a cognitive feeling and this experience is not attributed to some external source, how is the supposedly nonspecific experience interpreted? The answer in most psychology experiments is simple: The interpretation is explicitly provided by the question an experimenter asks. This is most obvious in the study by Mandler, Nakamura and van Zandt (1987). If the question is about brightness, fluent stimuli are judged brighter, if the question is about darkness, fluent stimuli are judged darker. But other classic findings are also open for this construal of providing an explicit interpretation of a cognitive feeling.

Again, the study by Jacoby et al. (1989) serves as an excellent example. To reiterate, participants read a list of names, knowing that none of the names on the list belonged to a famous person, that is, all the names were non-famous. After a delay of 24 hours they were presented with another list, consisting of names of famous and non-famous people. Their task was to decide whether a name was famous or not. The result was that non-famous names presented 24 hours earlier had a higher probability to be classified as famous than new names. In terms of the present discussion, this result can be construed as follows: The presentation of a name resulted in a noticeable difference in the fluency of processing compared to totally new names, even when a name itself could not be recognized. This enhanced fluency was then interpreted according to the task at hand, namely, to judge whether a name is famous or not. The authors construed the result as such that fluency was *misattributed* to the familiarity of the name and when a name seems familiar, it should be famous; yet, the step in between seems unnecessary given that we can simply assume an unspecific experience of fluency which is then interpreted according to the task at hand.

Though also in line with the authors' own hypothesis, the reported response latencies fit also with the interpretation hypothesis: False fame judgments (i.e., classifying a name from 24 hours ago as famous) take longer than true fame judgments (i.e., classifying a famous name as famous). If there was an inherent meaning to a cognitive feeling resulting from prior

experience, one would rather expect that this results in faster responses than in slower responses (Whittlesea, 2002).

It also seems possible to explicitly tell people the meaning of a given experience. For example, Winkielman and Schwarz (2001) had participants remember either 4 (which is supposedly easy) or 12 (which is supposedly difficult) events from their childhood. In addition, participants were provided with a theory stating that happy times fade fast from memory or that unhappy times fade fast from memory. Those who were led to believe the latter judged their childhood happier when retrieval was made difficult (i.e., remembering 12 events) and less happy when retrieval was made easy (i.e., remembering 4 events). This pattern was reversed when participants were led to believe that unhappy times of your life fade fast from memory. Obviously, participants used the explicit ad-hoc theories to interpret the ease or difficulty of remembering in constructing a judgment, in this case, the pleasantness of their childhood.

Furthermore, Skurnik, Schwarz, and Winkielman (2000) report data from Skurnik (1998) that show a reversal of the truth effect (Begg, Anas, & Farinacci, 1992). Usually, greater experienced fluency of processing of a statement leads to greater rated truth of that statement. But when the experimenter provides information that fluency is diagnostic for falseness rather than truth, this effect can be reversed (i.e., participants were told that false statements are more easily processed). The authors place these data in a greater framework of naïve beliefs about cognitive feelings; and these naïve beliefs are open to explicit manipulations. In terms of the present discussion, the interpretation of a given cognitive feeling is open to explicit instructions.

The notion that cognitive feelings are at least to some extent consciously available and open to explicit manipulations is corroborated by the susceptibility of cognitive feelings to attribution manipulations. Attributing a feeling correctly or incorrectly to an external source that is irrelevant for a given task eliminates the influence of the feeling on a subsequent judgment. This has been demonstrated for affective feelings (Schwarz & Clore, 1983; see Schwarz & Clore, 2003, for a review) and for general states of arousal as well (Zanna & Cooper, 1974; Cantor, Zillmann, & Bryant, 1975; Zillmann & Bryant, 1974; see Zanna & Cooper, 2000, for a review). An example for misattribution in the area of cognitive feelings in the Schwarz et al. study (1991, Exp. 3) has already been mentioned; the impact of the cognitive feelings was eliminated by leading participants to believe that background music facilitates or inhibits the recall of autobiographic events.

Another example is given in a study by Wänke, Schwarz, and Bless (1995). They employed a variant of the frequency estimation task by Tversky and Kahneman (1973).

Participants were asked to estimate the frequency of words that begin with the letter “t” relative to words that have “t” as the third letter. First, participants had to retrieve and write down 10 words that have “t” as a third letter and then 10 words that have “t” as a first letter: The latter task is supposed to be easier and therefore should result in higher frequency estimates of words with “t” as a first letter compared to “t” as a third letter. But the latter 10 words were written down on a sheet with pale t’s in the background. Participants were either told that this facilitates or inhibits the retrieval of words that begin with “t”. If the ease of retrieving words that start with “t” was attributed to the background of the sheet, the relative frequency estimates were lower than the estimates of a control group and when inhibition was presupposed, the frequency estimates were even higher than the control group’s estimates. If such simple verbal instructions suffice to reduce or even wipe out the effects of cognitive feelings, they must be at least to some extent explicit and consciously available, which speaks in turn clearly against an inherent or hard-wired meaning of cognitive feelings.

The general paradigm of Part I

The following two experiments first will explore the dependency of ease of retrieval effects on explicitly provided interpretations of the subjective experience. The paradigm is derived from an ease of retrieval design, where the content of retrieved information is plotted against the subjective experience of how easy or how difficult it was to retrieve that information (Schwarz et al., 1991; Wänke, Bless, & Biller, 1996). If it is easy to come up with positive information about a target, the evaluation of this respective target becomes more positive than when it is difficult to come up with positive information. The reverse is true for negative information (Wänke, Bohner, & Jurkowitsch, 1997).

As mentioned, people might have naïve beliefs about what ease of retrieval signifies without explicit instructions provided by the experimenter – the most prominent are frequency, recency, or typicality. To demonstrate that experiential states need interpretation, one must come up with an ease of retrieval context that evokes a feeling that does not lend itself automatically to naïve beliefs and interpretations in terms of frequency or typicality. To do this, an externalized ease of retrieval approach was used.

What is meant by externalized ease of retrieval? Referring again to the definition of feelings, that they provide information about processes that are otherwise inaccessible (Clore et al. 2001), it is not a long shot to extend this notion from internal processes to external processes. For example, imagine an oral examination and the examinee gives an immediate reply to a question. There are at least two interpretations for this quick response: The person has the answer directly available because the material was well rehearsed. But there is also the

negative interpretation that the person did not think about the question thoroughly and has only memorized the material very well without understanding it. On the other hand, a delayed response may indicate that the material is not well rehearsed; but maybe the person was pondering about the question, avoiding a simple and possibly wrong answer. Similarly, imagine you submit a term for an internet search and it takes some time until you get results. This can be easily interpreted that there are not many pages out there that feature that term. But it can also indicate that the search is thorough, whereas a very quick response may indicate a very superficial search. And at last, you could attribute the speed or the slowness to your internet connection.

For the following two experiments, a very similar logic was used: Participants' task was to evaluate four targets and they could ask other people about their opinion concerning the targets. In return, they got either a fast response or a delayed response. This simulates the internal search process for information and ease of retrieval is clearly operationalized as the delay between the request for information and the delivery of the response.

The general idea to externalize internal processes follows Fiedler, Brinkmann, Betsch, and Wild (2000) and has many advantages: First, it gives a clear definition of ease of retrieval, namely a delay of retrieval from an external source. Second, it allows manipulating the ease of information retrieval independently from the amount of retrieved information. Thus, it is possible to keep the size of the information sample constant across participants. And third, by externalizing the retrieval process, it is possible to directly compare participant's judgments across identical targets in terms of the provided information, varying only the ease of information retrieval across targets.

The big disadvantage naturally lies in the question: Would such a manipulation still be a feeling in the sense of experiential information? To show that this manipulation is indeed conceptually similar to the feeling derived from internally retrieved information, Experiment 1 included a simple manipulation that asked participants either to base their judgment on "how they feel about the target" or to base their judgment "rationally, how do you think about the target". A similar manipulation was employed by Verplanken, Hofstee, and Janssen (1998). They showed that it is possible to access the cognitive and affective components of attitudes by simply switching from a question of "how do you think about it" to "how do you feel about it". If the delay manipulation does resemble information derived from feelings, this manipulation should make a difference in the subsequent judgment; that is, a difference that is dependent on the delay manipulation.

In the Schwarz et al. (1991), study, one's own self-assertiveness was the dimension to be judged. Self-related judgments are obviously not ideal for an externalized approach to ease

of retrieval. Therefore, in Experiment 1 and 2, participants' task was to retrieve information about four politicians and judge their quality after the information search. The information that could be searched depicted two of these politicians as good, whereas the other two were depicted as bad. Orthogonal to this variation in valence, the ease of information retrieval was varied: For two politicians the information was quickly available, whereas for the other two the retrieval was delayed. Note that the differential ease or difficulty of retrieval is manipulated; it should make no difference if something is difficult or easy per se, but always in comparison to a standard, either a known standard or a standard available in a given context. The crucial point is the difference in the process.

It is also noteworthy that at the time of judgment the information has already been retrieved; actually, one should expect an influence from a process "how easy/difficult is it to form a judgment from the given information?" rather than from a process "how easy/difficult was it to retrieve the given information?". To explain the influence of experienced ease during data generation, we need to assume a metacognitive summary representation of the experienced feeling at the time of judgment (Koriat & Levy-Sadot, 1999): "I have 12 arguments, but it was very hard to generate these 12 arguments". If such metacognitive summaries are indeed consciously available, further implications are that people should be able to willfully disregard this information. Here, this is equivalent to ignore the second part of the statement: Just use the 12 arguments and forget about the difficulty of generating them.

Experiment 1 aimed to show that the delay in an externalized ease of retrieval approach resembles a cognitive feeling and that this feeling is consciously available and therefore susceptible to attribution manipulations. As such, it is assumed that the same conscious judgment mechanisms that apply to internal information retrieval apply to the retrieval of external information as well: There must be a noticeable discrepancy in the flow of experience (i.e., the information retrieval), this discrepancy needs to be attributed to the process of information retrieval and not to some irrelevant external source and there must be an applicable rule for the interpretation (i.e., what does the delay mean) of the discrepancy.

Experiment 1

The four targets to be judged in Experiment 1 represented within-participants analogs of the between-participants manipulation in the first experiment in the Schwarz et al. (1991) study. Retrieving instances in which you acted self-assertive is analogous to retrieving positive information, whereas remembering instances in which you did not act self-assertive is analogous to retrieving negative information. Retrieving six instances is equivalent to an

immediate retrieval and retrieving 12 instances is equivalent to a delayed retrieval. The combination of these two variables, valence and differential ease, resulted in four target politicians, labeled “good-delay”, “good-no delay”, “bad-delay”, and “bad-no delay”. In addition to these within-participants manipulations, two between-participants variables were manipulated to further test the proposed conceptual similarity of an external with an internal ease of retrieval approach.

The first between manipulation pertained to the attribution of the delay of information retrieval. Remember the comparison to an internet search: Although you might conclude that a topic is not featured on many pages if it takes a long time to process your request, you might not do so if you know that your slow internet connection is responsible for the delay. To show that attribution plays a role in the process, an external source was provided to explain the delayed information retrieval; if the delay is seen as non-diagnostic for the process, that is, attributed to an unrelated cause (e.g., a slow internet connection), effects related to the delay should disappear. So, in one condition, the delay was explained to be due to a technical problem, constituting an “explained” condition, whereas in another condition no explanation was given for the delay, constituting an “unexplained” condition.

The second manipulation pertained to the test whether the delay of external information retrieval conveys anything similar to a generic feeling. As mentioned, Verplanken, Hofstee, and Janssen (1998) showed that people can use different routes in the construction of evaluative judgment. In addition, Strack (1992; p. 251) gives the example of the well-known Müller-Lyer illusion. Your perceptual impression tells you that the line with inward arrows is shorter than the line with the outward arrows. But if you ever heard of the illusion, you should be able to ignore that experience-based information and take a ruler to check for the actual length of the lines; that is, you can construe a feeling-based or a content-based judgment. In the present experiment, this was implemented by asking participants to evaluate the targets based on how they feel about them or how they think about them. This manipulation tests en passant the conscious availability of metacognitive summaries, albeit indirectly. If participants can disregard experiential information willfully, they must be, at least to some degree, consciously available.

The specific hypotheses for the between manipulations are that the interplay of externalized ease of retrieval and retrieved content is contingent on two factors: First, the experienced ease must be attributed to the process and not to an external source, and second, the judgment must rely on feeling-based information and not on content-based information. The hypotheses for the within participants manipulations are that good politicians will be evaluated more positively than bad politicians, but this valence should interact with the delay

of the retrieved information. At this point, it is important to note that the delay in this context is non-specific and does not lend itself to an immediate interpretation in terms of typicality or frequency.

Method

Participants and Design. Fifty-three (41 women, 12 men) psychology students participated as partial fulfillment of a course requirement. They were randomly assigned to one of four experimental conditions, resulting from the orthogonal combination of the two between participants variables of attribution (explanation given vs. no explanation given) and basis for judgment (feeling vs. thinking). The valence of the politicians (good vs. bad) and the ease of information retrieval (immediate vs. delayed) were manipulated within participants.

Materials and Procedure. On arriving, participants were seated in a cubicle with a personal computer and were given a consent form that informed them about the upcoming tasks. If they agreed to participate, they signed the form and the experimenter started the computer.

The experiment was conducted using a computer program written in MS Visual Basic to present instructions and stimuli and to assess the dependent variables. The first screen asked the experimenter to enter a participant identification number, the sex and the age of the participant. The assignment to one of the four between conditions was contingent on the identification number. After this first screen, instructions and the cover story were presented.

The story of the experiment was straightforward: To exclude prior knowledge, participants were told that they had just arrived on an alien planet. On this planet, elections were about to be held. They were informed that their task would be to collect information about four politicians that were candidates in the election and to evaluate the candidates afterwards. If participants had no further questions, they could proceed to the information collection screen. On this screen, each politician was represented by a small picture of a computer-generated alien. The four pictures were identical, only the color of the alien's shirt and the background color varied; thus, the pictures were very similar but nevertheless clearly distinguishable due to the different colors. Participants were told that they could ask other aliens' opinions about each politician. Participants could search this information about the politicians by clicking on an "ask an opinion about this politician" button below this politician. They could ask seven opinions about each politician and they had to ask all seven opinions before proceeding to the evaluation. To avoid problems with the content of statements about politicians, participants were told that every opinion would be translated into a simple scale of ten stars: Ten golden stars signified the highest and best evaluation, whereas

zero stars signified the lowest and worst evaluation. These opinions about the politicians constituted the manipulation of valence. For two of the politicians, the underlying distribution of the seven opinions had a mean of seven stars, whereas the remaining two had a mean of three stars across the seven opinions that could be searched. The distribution was quadratic; so for the bad politicians with a mean of three stars, every value from zero to six stars occurred once, whereas for the good politicians with a mean of seven stars, every value from four to ten occurred once. The four pictures of the aliens were presented together with their respective “ask an opinion” button on the same screen in a four by two alignment grid. The position of each politician in this grid was newly randomized for each participant. If a participant clicked a politician’s respective “ask an opinion” button, a new screen with an alien face would appear with a speech bubble and the statement: “My opinion of this politician is:” and the ten star scale below the statement with an instance randomly drawn from the quadratic distribution. The drawing was realized without replacement, so each participant had a random sequence of the same seven opinions. The order of the information search across the politicians was under participants’ control. After four seconds, the screen with the speech bubble disappeared and the screen with all four politicians returned.

Orthogonal to the quality of a politician, the retrieval of the information (i.e., the opinion) was manipulated. For two of the politicians, the requested opinion was presented immediately. For the remaining two, the presentation of the requested opinion was delayed for three seconds. This constituted the differential ease of information retrieval. Nested within this manipulation was the attribution manipulation: In the “unexplained” condition, the delay was realized by showing a grey screen. In the “explained” condition, a statement was presented that the computer is busy translating the original opinion into the ten star rating system: “Translation in Progress...”.

After participants had completed their twenty-eight information collection trials (i.e. requesting seven opinions for four politicians), they were prompted by the program to proceed with the evaluation of the politicians. But before they could progress to the evaluation, the second between-participants manipulation was realized: In the “think” condition, participants were asked to give a rational and thoughtful evaluation. In the “feel” condition, participants were asked to give an evaluation based on how they feel about each politician.

On the evaluation screen, the politicians were on the same position in the alignment grid as before, but instead of a button, each picture had a vertical scroll bar below it, with the endpoints labeled very good and very bad. Participants were asked to give their evaluation using this scroll bars and to confirm their judgments once they were done by clicking a proceed button. The computer program converted the position of the scroll bar to a value

between 0 and 100, with higher numbers indicating a better evaluation. When they clicked the confirm button, a message box informed them that another evaluation would follow: Using the same scroll bars, participants had to judge the ease of information retrieval for each politician. As such, the endpoints of the scroll bars were labeled “very easy” and “very difficult”. Once they completed this rating, the next screen prompted them to contact the experimenter. The experimenter thanked and fully debriefed participants about the purpose of the experiment.

Results

Due to an error in assignment of participants identification numbers, the frequencies across the four between conditions were not balanced; there were 13 participants in the “no explanation/thinking” condition, 12 participants in the “no explanation/feeling” condition and 14 participants each in the “explanation/thinking” and explanation/feeling” conditions. Thus, all reported F -values are based on Type III sums of squares.

The main dependent variables were the evaluations of the politicians and the rated ease of information retrieval. To reiterate, it is assumed that the evaluation of the targets is based on an interaction of the retrieved information’s content (i.e., the star rating) and the experience of how easy it was to retrieve that information (i.e., was it easy or difficult to retrieve that information). The impact of this experiential information should be contingent upon two factors: First, is the experience attributed to the respective processes and not to an irrelevant external cause (i.e., a slow translation) and second, is the experiential information used for the judgment (i.e., feeling vs. thinking).

Overall Analysis. Table 1 shows the mean evaluation and standard deviations of the four targets across the four between conditions. Using a mixed ANOVA, these evaluations were analyzed as a function of the between participants variables, that is, the basis for the judgment (feeling vs. thinking) and the attribution of the delay (explained vs. unexplained), and the within manipulation of valence (good vs. bad) and information retrieval (immediate vs. delayed). This analysis yielded a strong effect for the valence of politician, $F(1, 49) = 177.69, p < .001$, indicating that participants encoded the information about the targets correctly and judged the good politicians ($M = 82.11, SD = 21.47$) better than the bad politicians ($M = 27.49, SD = 24.25$; higher values indicate a more positive evaluation). More interestingly, the analysis showed an interaction of the manipulation of information retrieval and the basis for the judgment, $F(1, 49) = 5.59, p < .05$, and even most important for the present argument, a marginal significant interaction of all four factors, $F(1, 49) = 3.98, p <$

.06, showing that the evaluations indeed varied as a function of valence, delay, attribution and the basis of the judgment.

Table 1

Mean evaluations of the four targets as a function of information retrieval, attribution, basis of judgments and quality of the target in Experiment 1 (Standard deviations in parentheses).

Explanation	“thinking”		“feeling”	
	3-Stars	7-Stars	3-Stars	7-Stars
	Immediate information retrieval			
not given	20.77 (15.41)	78.69 (24.81)	27.92 (22.83)	88.17 (7.04)
given	21.57 (19.65)	73.00 (28.38)	33.07 (27.02)	86.00 (19.27)
	Delayed information retrieval			
not given	26.46 (25.62)	86.00 (13.69)	15.50 (15.01)	89.00 (9.06)
given	32.43 (30.85)	88.36 (14.94)	40.00 (27.91)	69.57 (32.38)

Note. Higher values represent a more positive evaluation with 0 being the lowest and 100 being the highest evaluation.

To analyze this rather complex pattern further, the same ANOVA was conducted separately for “delay” and “no delay” information retrieval targets. This strategy is advised not only by instructive considerations, but also by the fact that the attribution manipulation is nested within the delay-manipulation; given the immediate information retrieval, an explanation was neither required nor given. Thus, the design was not fully crossed and the dependent variables should be analyzed separately for the “delay” and “no delay” conditions.

“No delay” targets evaluation. A 2 (feeling vs. thinking) x 2 (7 stars vs. 3 stars) mixed ANOVA was employed to analyze the evaluation of the targets without delayed information retrieval. In this analysis it is again apparent that participants clearly distinguished between the 7 star politician ($M = 81.26$, $SD = 21.99$) and the 3 star politician ($M = 25.85$, $SD = 21.69$). This effect was again highly significant, $F(1, 51) = 133.48$, $p < .001$. In addition, there was a main effect for the kind of judgment made. People in the “feeling” condition judged the targets in general more positively than people in the “thinking” condition, $F(1, 51) = 9.31$, $p < .005$. This effect was unexpected and will be addressed in the discussion. No other effect was significant, all $F_s < 1$. As expected, this indicates that the interaction effects from the overall analysis are mainly due to the targets with delayed information retrieval.

“Delay” targets evaluation. To analyze the data from the targets with delayed information retrieval, a 2 (feeling vs. thinking) x 2 (explanation vs. no explanation) x 2 (7 stars vs. 3 stars) mixed ANOVA was employed. This ANOVA showed again that participants judged the 7 star politician better ($M = 82.96, SD = 21.10$) than the 3 star politician ($M = 29.13, SD = 26.67$), $F(1, 49) = 131.72, p < .001$. The analysis also showed an interaction of the valence of the politician and the attribution of the delay, $F(1, 49) = 6.23, p < .05$. This interaction indicates that the good politician was judged better when the delay was not explained than when it was explained, ($M = 87.44, SD = 11.56$ vs. $M = 78.96, SD = 26.53$, respectively). The bad politician on the other hand was judged better when the delay was explained ($M = 36.21, SD = 29.12$) than when it was not explained ($M = 21.20, SD = 21.51$). But this two-way interaction is qualified by the predicted three-way interaction of valence, attribution and basis of the judgment, $F(1, 49) = 4.48, p < .05$.

For an easier interpretation of this interaction, the means are displayed in Figure 1, corrected for the main effects and the lower order interactions (Rosenthal & Rosnow, 1995, 1996). From Figure 1 it is evident that participants evaluated the “good” politician (i.e., 7 stars) more negatively when there was no explanation given for the delay and more positively when there was an explanation given for the delay. This pattern is reversed for the “bad” (i.e., 3 stars) politician. This target was evaluated more positively when no explanation for the delay was given and more negatively when an explanation was given. Notably, this pattern is only evident for participants in the “feeling” condition. For participants who were instructed to form a rational judgment, neither the attribution manipulation nor the delayed information retrieval had an influence on the judgments.

All together, the separate analysis for the “delay” and “no delay” targets suggests that the manipulations were successful in such that the delayed information retrieval influenced judgments only when they were not attributed to a non-diagnostic source and the judgment is not based on rational considerations.

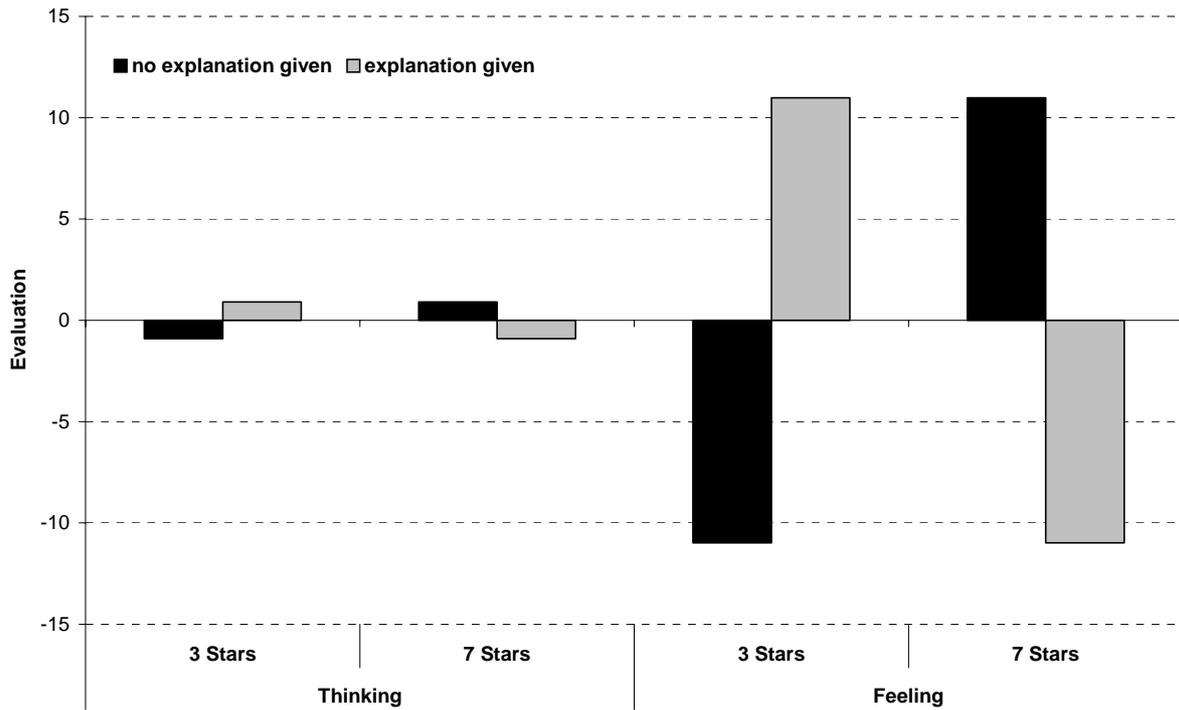


Figure 1. Evaluation of politicians with delayed information retrieval in Experiment 1 as a function of condition, valence and explanation (corrected for main effects).

Rated ease. The rated ease data should shed some light on the claim that the experiential state during information retrieval is consciously available as a metacognitive summary. If this is indeed the case, participants should be able to rate the ease of the information retrieval; these post-hoc ratings of the information retrieval process' ease should covary systematically with the respective evaluations. To test the first proposition, the mean rated ease across the “no delay” targets and the across the “delay” targets was calculated. With a higher rating indicating greater ease, the mean rating for “no delay” targets was $M = 66.78$ ($SD = 24.51$) and the mean rating for the “delay” targets was $M = 52.14$ ($SD = 22.19$). This difference was reliable, $t(52) = 2.83$, $p < .01$. Although this is hardly more than a manipulation check, the correlations between the rated ease and the evaluations in different conditions is not as trivial. The mean correlations across all four between-participants conditions are displayed in Table 2. Although only four correlations are significant, the pattern is nevertheless highly instructive. First, significant correlations only emerged for targets with delayed information retrieval. Second, for the four significant correlations, the pattern from the “thinking” condition is exactly reversed in the “feeling” condition. In the “thinking” condition, the evaluation of the “delay” targets correlates negatively with the rated ease, whereas in the “feeling” condition the exact same correlation is positive, indicating that a higher rated ease leads to more favorable evaluations.

Although an “ease of retrieval” logic would predict that this pattern should be qualified by the valence of the target (i.e., positive correlations for good targets and negative correlations for bad targets in the “feeling” condition and vice versa in the “thinking” condition), the pattern shows that people have an available summary of the information retrieval process which can be used strategically in the formation of judgments.

Table 2

Correlations of rated ease and evaluations for the four targets in the four conditions in Experiment 1

Target	thinking		feeling	
	explained n = 14	unexplained n = 13	explained n = 14	unexplained n = 12
no delay–good	.231	.146	-.130	.269
delay–good	.154	-.306*	.144	.554*
no delay–bad	-.157	.142	-.144	-.136
delay–bad	-.440*	.134	.409*	.138

Note. Positive correlations indicate that higher rated ease covaries with the positivity of the respective evaluation.

* $p < .05$, two tailed.

Discussion

Experiment 1 tested the impact of experiential states during the process of information retrieval on subsequent judgments in an externalized ease of retrieval paradigm. As such, the experiment represents a conceptual replication of the Schwarz et al. (1991) study with an externalized ease of retrieval. The experience during information retrieval was manipulated by a delayed vs. immediate information presentation and this delay was either attributed to the slowness of the translation system or not explained at all. Thus, the design tested the role of attributions about the source of these experiences. In addition, to show that the feeling is available as a conscious meta-summary, participants were asked to give a feeling-based or a content-based evaluation. That is, it was tested if participants can deliberately use their feelings in a judgment or deliberately discard the experiential information.

The strongest effect present in the data was due to the valence of the targets. Such a main effect was not found in the original Schwarz et al. (1991) study, but the explanation is

obvious: Here, valence was manipulated as a within participants variable, whereas the valence in the original experiments (remembering instances of self-assertive vs. non self-assertive behavior) was manipulated as a between-participants variable. Obviously, the within manipulation sensitized participants for the differences in valence.

More interestingly, the simple instruction to give a content-based or a feeling-based judgment was successful; for targets with no delay in information retrieval this instruction resulted in a general more positive evaluation, independent of the actual quality of the target. Apparently, the immediate information retrieval was interpreted as a general positive attribute and not seen in conjunction with the content of the retrieved information. For the targets with delayed information presentation, this general effect was qualified by the actual quality of the target and the attribution of the delay. Moreover, this was only the case when the delay of information retrieval was not attributed to the slowness of the translation system. The latter instruction should render the delay non-diagnostic for the judgment; this misattribution manipulation was successful as well: The resulting three-way interaction implies that the use of feelings is indeed contingent upon deliberate considerations and the attribution of the feeling during the respective process. This finding is corroborated by the correlations of rated ease and respective evaluations, which show a perfect reversal of the significant correlations for targets with delayed information retrieval in the feeling condition compared to the targets in the thinking condition (cf. Table 2).

Most important for the present argument is the finding that when the judgment is feeling-based, and it is not attributed to an external source, the resulting interaction pattern is inconsistent with an ease of retrieval logic. The explanation is straightforward and derives directly from the proposed model: In standard ease of retrieval experiments, the assumption is that people interpret information that is hard to retrieve from memory as either not typical or not frequent. In the present context, no such interpretation was explicitly or implicitly available. If there was an inherent meaning to ease of information retrieval, one would expect the conceptually same pattern independently of the context that elicits the experiential state of ease or difficulty. The systematic pattern suggests that participants might have applied a simple “fast is good” interpretation for the no-delay targets; an interpretation, however, was neither explicitly nor implicitly provided. This strengthens the claim that effects of cognitive feelings depend on interpretation, and this interpretation follows cues provided by the environment. Experiment 2 will show that if cues are provided that allow an interpretation of experienced ease or difficulty terms of frequency or typicality, a classic ease of retrieval pattern can be obtained in the externalized ease of retrieval paradigm.

Experiment 1 showed so far that participants clearly experienced a difference in the speed of information retrieval which resembled a feeling; this feeling is open to attribution manipulations and consciously available as a metacognitive summary which is accessible for strategic use in subsequent judgments.

Experiment 2

As explained, Experiment 1 did not provide an applicable interpretation for the delayed information retrieval. The purpose of Experiment 2 is to show that if the respective interpretation is available for interpreting delayed information retrieval, a classic ease of retrieval pattern can be obtained. To do this, the attribution manipulation was coupled with the interpretation during the retrieval of the respective information. The underlying rationale is the same as when fluency is attributed to a source in the past, this attribution encompasses the interpretation, namely a previous encounter with a stimulus. This logic was used to provide an interpretation for the delayed information retrieval.

Before turning to the implementation of the interpretation, a weakness of Experiment 1 needs discussion and has to be dealt with. The attribution manipulation, which will be used now for interpretation purposes, was nested within the delay manipulation. That is, the attribution manipulation was only taking place for the delay targets. This explains why the interactions were only present for these targets. But if the interpretation is to be coupled with the attribution manipulation, this manipulation must be applicable to all targets. So, instead of having a “no delay” vs. “delay” manipulation, a “short delay” vs. “long delay” manipulation is necessary. This gives the possibility to provide an attribution and interpretation for all four targets.

In Experiment 1, there was either an attribution to slowness of the translation or no information given at all (a grey blank screen). Instead of this blank screen, an attribution to the process of information retrieval was now included, explicitly verbalizing the ease of retrieval logic. When asking for an opinion about a politician, the alien asked appeared immediately but instead of giving the opinion immediately, the statement was made: “I have to think about this for a moment...” For the short delay targets, this screen was presented for 1 second. For the long delay targets, this screen was present for 4 seconds. Thus, the attribution of the delay encompassed an interpretation in terms of frequency or typicality, because non-frequent or non-typical instances are more difficult to generate and take therefore more time. The alternative attribution to the translation process (“Translation in progress...”),

which should render the delay information non-diagnostic, was the same as in Experiment 1, but this misattribution was now presented for all four targets in the respective condition.

Finally, the manipulation of asking participants to make rational vs. spontaneous judgments was dropped, because Experiment 1 showed convincingly that asking for rational judgments eliminates the influence of experiential states. Everything else was parallel to Experiment 1.

Method

Participants and Design. 94 students (41 women, 53 men) from various faculties of the University of Heidelberg participated for payment of € 5. They were randomly assigned to one of the two between-participants conditions, attribution to the translation (“computer attribution”), rendering the delay information non-diagnostic, or to the process of retrieval itself (“person attribution”), encompassing the interpretation in terms of frequency or typicality. The random assignment was again based on participants’ identification number. The opinion about the targets (good vs. bad) and the delay (short vs. long) were manipulated within participants. The orthogonal combination of these two variables resulted once more in four targets: “good-short delay”, “bad-short delay”, “good-long delay”, and “bad-long delay”.

Materials and Procedure. The same computer program as in Experiment 1 was used with some slight modification to implement the described changes. The procedure was also parallel to Experiment 1 with the following changes: The delay for the asked opinions was now 1 sec for the “short delay” condition and 4 sec for the “long delay” condition. This manipulation was orthogonal to the actual quality of the politicians (mean of three vs. seven stars). The interpretation manipulation was realized in combination with the delay manipulation. For half of the participants, a screen appeared with the statement “Translation in progress...”, attributing the delay to the computer and as such, rendering it non-diagnostic for the content of the retrieved information. The other half saw a screen with the face of the respective alien and the statement “I have to think about this for a moment...”, attributing the delay to the person retrieving the information and as such making the delay diagnostic for the content of the retrieved information in terms of frequency or typicality.

Parallel to Experiment 1, participants had to evaluate each politician using a vertical scroll bar beneath each picture with the end points labeled very good and very bad. The program translated the position of the scroll bar to values from 0 to 100, with higher values indicating a more positive evaluation. After completing these ratings, participants were asked to rate the ease of information collection about each target; the response format was the same, but the endpoints of the scroll bars were labeled very easy and very difficult. The position of

the scroll bar was again converted to numerical number from 0 to 100, with higher values indicating higher rated ease. After completing the ratings, the computer program prompted participants to contact the experimenter, who fully debriefed, thanked and paid them.

Results

No one refused to participate and no data were precluded based on the experimenter's observations. Therefore, there were data available from 47 participants in the "person attribution" and 47 participants in the "computer attribution" condition. The dependent variables were the same as in Experiment 1; the judgments about the quality of the four target politicians and the rated ease of information retrieval. In contrast to Experiment 1, the design was fully crossed and thus, an overall analysis of the results is possible. To reiterate, an attribution to the computer should render the experiential information of the delay non-diagnostic, whereas the attribution to the person retrieving the information should trigger the exact same interpretation as a typical ease of retrieval study, namely, that information that is easy to retrieve is more frequent or more typical. Therefore, evaluations should vary as a function of the valence, the delay and the attribution of the retrieved information.

Target evaluations. The mean target evaluations in the two between conditions are given in Table 3. An attribution (person vs. computer) x valence (good vs. bad) x delay (short vs. long) mixed ANOVA was used to analyze these evaluations. As in Experiment 1, the strongest effect is due to the valence of the politicians, $F(1, 92) = 686.42, p < .001$. Good politicians were judged more positively ($M = 83.56, SD = 17.77$) than bad politicians ($M = 25.13, SD = 20.83$). Similar to Experiment 1, this is attributable to the within comparison of good and bad target politicians, which leads to a sensitization for the difference in valence. The only other significant effect was the expected three-way interaction of attribution, valence and delay, $F(1, 92) = 4.79, p < .05$.

Table 3

Mean evaluations of the four targets as a function of information retrieval, attribution and valence of the target in Experiment 2 (Standard deviations in parentheses).

Attribution	short delay		long delay	
	3-Stars	7-Stars	3-Stars	7-Stars
person	18.51 (15.70)	84.11 (15.95)	30.81 (24.54)	83.04 (19.97)
computer	26.36 (20.46)	82.36 (21.31)	24.85 (20.43)	84.72 (13.12)

To interpret this interaction correctly, it is displayed in Figure 2, corrected for the main effects and all lower order interactions (Rosenthal & Rosnow, 1995, 1996). Let us first consider the case when the delay is diagnostic for the retrieved information, that is, when the delay is attributed to the person. Figure 2 shows that given a short delay of information retrieval, the bad politician is rated worse and the good politician is rated better. On the other hand, when the delay of information retrieval is long, the good politician is rated worse and the bad politician is rated better. This pattern is reversed when the delay of information retrieval is attributed to the slowness of the translation computer and therefore not diagnostic for the retrieved information. This result is expected in a typical ease of retrieval study.

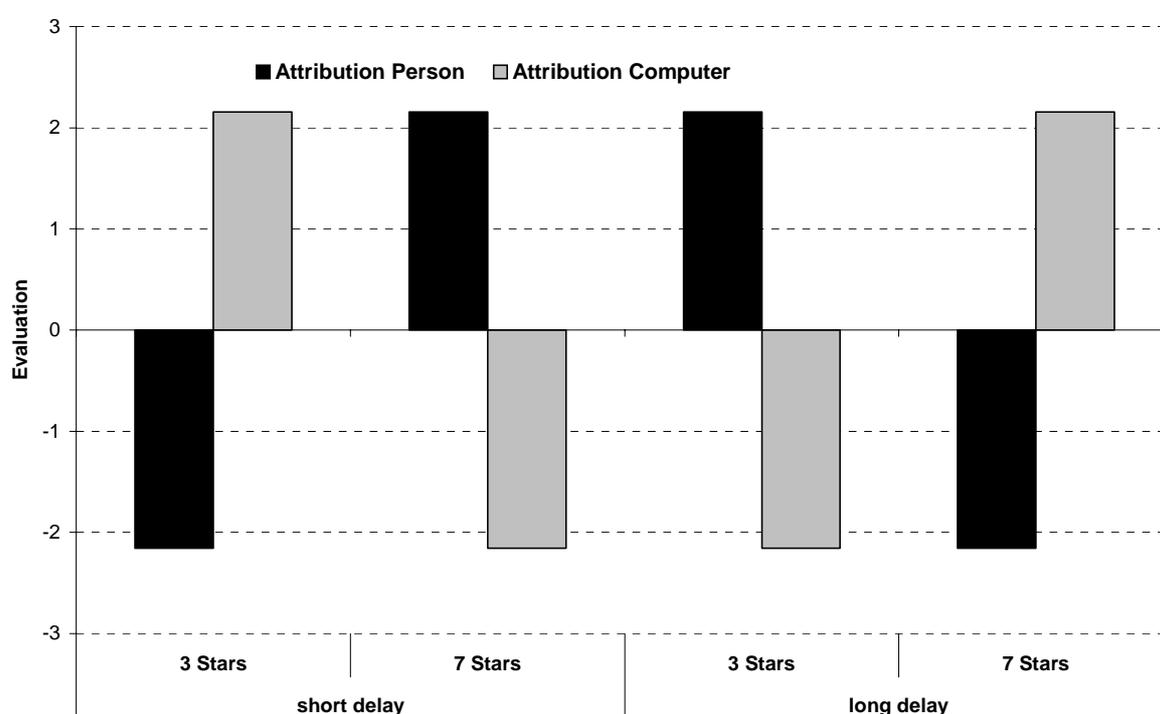


Figure 2. Evaluations of politicians in Experiment 2 as a function of valence, information retrieval delay, and attribution (corrected for main effects).

Rated ease. As in Experiment 1, the mean rated ease was calculated for the short and long delay targets. As expected, the participants rated the information retrieval for short delay targets easier ($M = 70.53$, $SD = 21.95$) than for long delay targets ($M = 63.56$, $SD = 25.31$). This difference was statistically reliable, $t(93) = 2.92$, $p < .005$. Having established this, the more interesting analysis is the correlation of the rated ease and the respective evaluation. Following an ease of retrieval logic, higher rated ease should lead to more positive evaluations for the good politicians, but to more negative evaluations for the bad politicians. In addition, this should only be the case when the ease is diagnostic for the retrieved

information, which is not the case when the delay is attributed to the slowness of the translation computer. To test this, simple correlations were computed between the rated ease and the respective evaluation, separately for the two attribution conditions. The results are displayed in Table 4.

Table 4

Correlations of rated ease and evaluations for the four targets in the two conditions in Experiment 2

Attribution	short delay		long delay	
	3-Stars	7-Stars	3-Stars	7-Stars
person	-.127	.574***	-.287*	.334*
computer	-.162	.044	.000	.255

Note. Positive correlations indicate that higher rated ease covaries with the positivity of the respective evaluation, $n = 47$ in each cell.

* $p < .05$, ** $p < .01$, *** $p < .001$, two-tailed.

As shown above, the retrieval of information about short delay targets was rated significantly easier than for the long delay targets. When this short delay (i.e., the ease of retrieval) is attributed to the person and therefore diagnostic for the retrieved information, the correlations for these short delay targets show the predicted pattern: If ease is rated high for a good politician, the evaluation of this politician is better, as implied by the positive correlation. On the other hand, when ease is rated high for a bad politician, the evaluation of this politician is worse, as implied by the negative correlation. Although the latter correlation is not significant, one has to take into account that the prediction is about differences between good and bad politicians, that is, ease should be used differentially for good and bad targets (good gets better, bad gets worse). Table 4 presents two-sided test of correlations against zero; if one correlation is different from zero and the other has a different sign, the difference between the two correlations is necessarily statistically reliable. The same pattern is true for targets with a long delay of information retrieval. If ease is rated low for a good politician, the evaluation gets worse, resulting again in a positive correlation, whereas the same correlation is negative for a bad politician, indicating that the evaluation got better when the retrieval of information was rated not easy. Both correlations are significantly different from zero, so they are necessarily different from each other.

However, when the delay is attributed to the slowness of the translation computer and therefore independent of the content of the information (i.e., non-diagnostic), there should be no correlation between rated ease and the evaluation of the respective politician. This was confirmed for all four targets; no statistically reliable correlations emerged. But again, the difference in correlations for good and bad politicians is the crucial test. As none of the correlations itself is reliably different from zero, the straightforward logic from above cannot be employed and the difference in correlations between the good and bad politicians must be tested. Using Fisher's difference and z-transformation of correlations (Fisher, 1934), no reliable difference was found, neither for the short delay politicians ($z = -0.973, p < .17$) nor the long delay politicians ($z = -1.22, p < .12$). As hypothesized, there was no evidence for the use of the delay information when the ease/difficulty of information retrieval was attributed to the slowness of the translation computer.

Discussion

The goal of Experiment 2 was to show that when a fitting interpretation is available, a classic ease of retrieval pattern can be obtained in the externalized ease of retrieval paradigm. In addition, given the applicable interpretation, the effects should still depend on the diagnostic value of the experienced ease. To test this, the delay of information retrieval of positive and negative information was manipulated. This delay was attributed either to an external source and therefore rendered not diagnostic or to the target of the retrieved information and therefore rendered diagnostic. The attribution to the target implied an ease of retrieval interpretation for the use of this information in terms of typicality or frequency: If the retrieval of positive (negative) information is retrieved fast (slow), positive (negative) information should be frequent/typical and a better (worse) evaluation should be the result. The reverse is true when the retrieval of positive (negative) information is slow; then a worse (better) evaluation should be the result.

Providing this interpretation, the predicted three-way interaction of valence of information, delay of information retrieval and attribution of the retrieval delay was obtained. If the retrieval delay was attributed to the process (i.e., the person retrieving the information), the good politicians were indeed judged better if the information retrieval was only shortly delayed compared to a long delay, whereas bad politicians were indeed judged worse when the delay was short compared to a long delay.

The effectiveness of the provided interpretation of the retrieval delay in terms of frequency or typicality is further strengthened by the results from the rated ease. Given an attribution to the person, those correlations were systematically positive for good politicians

and negative for the bad politicians. It is noteworthy that these correlations were compared within participants; that is, participants used the retrieval speed differentially for good and bad targets (i.e., the retrieved information about these targets). Participants did not apply a simple heuristic of judging the information retrieval delay as fast for good politician and slow for bad politicians. Rather, they produced highly systematic correlations for good and bad targets given a fast or slow information retrieval. But this is only the case when the speed of retrieval is diagnostic for the retrieved information in terms of typicality or frequency, whereas if the speed is attributed to external reasons, virtually no correlations exist.

The rated ease after the evaluative judgments also suggests that a metacognitive summary of the experienced retrieval speed is available at the time of judgment and that it is used strategically in forming the judgment and open for conscious interpretation (Koriat & Levy-Sadot, 1999).

Summary of Experiment 1 and 2

The general tenet of this thesis is that cognitive feelings need interpretation, much the same way as an unspecific arousal state can be interpreted as different emotions given the appropriate context cues (Schachter & Singer, 1962). Following from this assumption, it can be said that it should not matter whether such experiences are externally, rather than internally generated. Two experiments tested whether it is possible to obtain effects of experiential states on judgments, using an externalized ease of retrieval paradigm. The logic followed the idea of externalizing information search strategies, as done in Fiedler et al. (1999), in combination with the paradigm of Schwarz et al. (1991). But rather than having participants retrieve few (which is easy) and many (which is hard) instances of positive and negative information, the ease of retrieval was manipulated directly in an external information search. This has the huge advantage of experimental control over the retrieved content and the actual ease or difficulty of retrieval, which was manipulated via retrieval speed, that is, the delay of information retrieval and presentation.

Experiment 1 was successful in showing three points: First, effects of the delay manipulation were only visible if participants were asked to make feeling-based judgments. This showed to some extent the resemblance of externally generated experiential states to feelings as the term is defined in the introduction (i.e. giving feedback about otherwise inaccessible processes). Second, the effects also vanished when the delay was attributed to an external cause, and therefore, not diagnostic for the retrieved information. This showed that the ease or difficulty of retrieval in this externalized context is open to the same misattribution

manipulations as in typical ease of retrieval studies. And third, the delay was interpreted differently for positive and negative information, but the resulting pattern did not follow an ease of retrieval logic. The supposed explanation is that the context did not provide the cues for an interpretation of the delay in terms of frequency or typicality.

Experiment 2 provided an interpretation in terms of frequency or typicality embedded in an attribution manipulation. And thus, the expected interaction of delay in information retrieval and the retrieved content was obtained. In addition, the correlations of rated ease of information retrieval and the respective judgments suggested a highly strategic and conscious use of the delay information.

What do these results imply? First, there is no need to assume a privileged access to knowledge about experiential states (Nisbett & Wilson, 1977). People are able to interpret experiential information (here: the speed of information retrieval) that is externally generated as well as internally generated experiential information (in comparison to the content of the information). To illustrate this point, let us refer once more to a candidate in an oral examination: Even if the given answer is correct, the examiner might use the speed of answering, that is, the retrieval of the correct answer from memory as additional information. For example, a quick answer might indicate that the material is well rehearsed and the answer is readily available, leading to a better evaluation. In comparison, a slow response might be interpreted as such that the candidate is not sure about the answer or just that the material is not as well rehearsed, leading to a less positive evaluation.

If the examination were written compared to oral, this information would be lost, much the same way as the ease of argument generation is lost when a yoked group reads the generated arguments and the evaluation is again a function of the number of arguments and not the ease of retrieval (Wänke, Bless, & Biller, 1996).

The reported experiments also provide a possible answer to the question how the experienced ease during information retrieval influences subsequent judgments. As mentioned in the introduction, cognitive feelings should have an impact if it is easy or difficult to integrate the information to a judgment, but not if it is easy or difficult to retrieve the information, because ease is not experienced during the information integration. The offered solution is the idea of a metacognitive summary of how easy or difficult the process of retrieval was. This summary serves then as input during the phase of information integration. Such a conscious and strategic use is supported by the systematic correlations of self-reported ease and the respective evaluations. If participants can give a self-report on the experienced ease, it must be at least partly consciously available. This explanation is also part of the

definition of the term “feeling” that is used here: Feelings provide feedback about otherwise inaccessible processes.

Altogether these results suggest that experiential information, even when it is externalized, can be used in much the same way as internally generated information, that is, information that results from the difficult or easy retrieval from memory. Again, three points in this usage of experiential information are crucial: Is there a notable difference in the flow of experience, is the difference attributed to the process, and is there an applicable interpretation for the experience.

PART II: INTERPRETATIONS LEARNED VIA ECOLOGICAL FEEDBACK

Part I has demonstrated the interpretation of a cognitive feeling in an externalized ease of retrieval context; it was shown that non-specific experiential states are consciously available and that they can be employed strategically in the formation of judgments. But there are at least two points to criticize: First, in the externalized context, such feelings might be non-specific and therefore the necessity of interpretation exists. Cognitive feelings resulting from internal processes might have an inherent meaning nevertheless. Second, people do not go through their daily lives being told what certain experiences mean. The case when an unexplained experiential feeling arises and the experimental situation suggests an explanation only shows that the experience *can* be consciously available and *can* be used vicariously. But how do people know how to interpret cognitive feelings in a real-world setting? It seems unlikely that we use conscious interpretations all the time. And even so, the question remains: How do we come up with these interpretations, conscious or not? Why is ease interpreted as frequency or typicality? Why is fluency mistaken for fame?

Part II will lend support to the hypothesis that the interpretation of cognitive feelings is learned ontogenetically by the route of ecological feedback learning. That is, people learn how a certain experience resulting from conceptual or perceptual processing is to be interpreted; for example, a feeling signifying that you will be able to remember a piece of information that seems inaccessible right now, that you have indeed met a person before whose face seems familiar and that information that is easily accessible is indeed more frequent, recent or typical. The same unspecific experience can account for differential effects in differential contexts. By continuous feedback, people learn interpretations of otherwise non-specific discrepancies in their flow of conceptual or perceptual cognitive processes, which are triggered by the given context. To use the case of an oral examination again as an example: If the examiner often got elaborate and thoughtful responses after a delay or after

hesitation, she might come to interpret a delay positively, whereas if she got frequent wrong or incomplete responses after a delay or hesitation, she might come to interpret a delay negatively.

Such learning might be conscious in the beginning, but can also work in the absence of conscious effort, because the use of a given interpretation (e.g., feelings of familiarity as a basis for recognition judgments; Whittlesea, 1993) has become automatic by frequent application of this interpretation (Bargh & Chartrand, 1999; Bargh & Ferguson, 2000). This sets the present approach aside from the application of lay-theories (Schwarz, in press) and the idea of unconscious attribution (Whittlesea & Williams, 2001a). Attribution in the current thesis is not equal to an interpretation, but the realization that a feeling is diagnostic for a process (e.g., attributing the inability to solve an item in a test to the difficulty of the test and not to the external cause of loud background noise). However, this view does not exclude the prerequisites stated in Part I. The experience must still be attributed to the target at hand and the judgment must allow for the influence of feelings. Conversely, although many judgments can be feeling-based (e.g., the length of the lines in the Müller-Lyer illusion, the assessment of how difficult $2 + 2 = x$), the judgments themselves are still consciously available and open to control. Yet, in the flow of daily life, such conscious control is probably the exception rather than the rule (Bless & Forgas, 2000).

It is important to note that such learning is by no means restricted to cognitive feelings; people learn all the time by simple associative mechanisms the ecological meaning of cues provided in their environment. As already mentioned, one of the standard tasks in everyday life is to interpret perceptual cues to infer the size, distance or depth of objects in our environment (Brunswik, 1941, 1943). But people can learn far more complex patterns, apart from such general functions as visual perceptions. For example, chess players learn the meaning of board configurations, so that they do not need to analyze a whole configuration to see whether black or white will win the game. Rather, they report a feeling whether black or white will win.

In addition, cognitive feelings need not be diagnostic in a given context. This is always the case when the context triggers an interpretation that is not applicable or is not learned in this specific context. In other words, an interpretation is used that has no ecological validity in this context. The principle in experiments using misleading cues follows the analogy of optical illusions. Optical illusions help understand vision because they show what kind of cues the cognitive systems uses to interpret environmental cues, for example, in the perception of depth (Goldstein, 2002). To investigate the impact of those cues on the perception of depth, the cues are manipulated independently from actual depth. Similarly, in

studies on perceptual fluency in a memory context, the fluency of perception is also manipulated independent of prior exposure to investigate the impact of fluency on recognition judgments. If greater fluency leads to a higher probability of “recognizing” a stimulus (again, independent of actual prior exposure), one can conclude that people indeed use this fluency cue as a basis for judgments of recognition. Biased judgments resulting from inferences based on the wrong cue have been demonstrated in the fluency domain (Benjamin, Bjork, & Schwartz, 1998), but are common in other domains as well. For example, the color black is learned to be a cue for aggressiveness and evil in the general context of a western culture, from the evil Black Riders in Tolkien’s “Lord of the Rings” (1954/1955) up to the parody of an evil Black Knight in Monty Python’s “Quest for the Holy Grail” (1975). It has most likely no ecological validity in the context of professional sport, but nevertheless, sport referees sometimes rely on this misleading cue, as in the case when more fouls are called against a team wearing black rather than white uniforms (Frank & Gilovich, 1988).

The General Paradigm of Part II

The empirical part of Part II, consisting of Experiment 3 and 4, will try to show the application of learned interpretations of experiential states, which as cognitive feelings give feedback from otherwise inaccessible processes. Again, the idea is that cognitive feelings (i.e., experiential states during conceptual and/or perceptual processing) follow interpretations that are acquired via feedback learning. As these interpretations are acquired in given contexts, the same feeling can result in different interpretations, which are themselves triggered the given context itself. In other words, the ecological validity of an interpretation in a given context is learned.

Experiment 3 and 4 will manipulate a classic cognitive feeling, the perceptual ease of processing. For reasons of clarity, I will use the term fluency, that is, the fluency of perception. Perceptual fluency may result from various causes; in most cases, it is assumed that having encountered a stimulus before notably facilitates the processing of the stimulus in following presentations (Mandler, 1980; Jacoby & Dallas, 1981; Jacoby, 1983). Nevertheless, the experience of perceptual fluency can be interpreted in different ways, for example, that the stimulus is more likable (Kunst-Wilson & Zajonc, 1980; Zajonc, 1980), that a stimulus is brighter or darker (Mandler, Nakamura, & Van Zandt, 1987) or that a statement has a higher possibility of being true (Reber & Schwarz, 1999). But the general finding in the context of a memory task is that fluency is interpreted as such that one has encountered a stimulus before; that is, participants judge an item with higher probability as “old”, when the item is fluently perceived or processed (Johnston, Hawley, & Elliot, 1991; Johnston, Dark, & Jacoby, 1985;

Mandler, 1980; Whittlesea, Jacoby, & Girard, 1990). These results can be construed according to the present hypothesis as follows: The effect occurs because the difference in perceptual processing is experienced, it is attributed to the presented stimulus (and not, for example, to the dim light) and there is an available interpretation; in the case of a memory test the interpretation is that you have encountered this stimulus before.

Taking experiencing and attributing for granted, the learning of a new interpretation was embedded in a simple recognition task. In a learning phase, participants saw a list of 80 English first names. Then in a following recognition phase, they saw another list of 80 first names; half of them being from the first list, and half of them being new names. Orthogonal to the status of being “old” or “new”, the perceptual fluency of the names was manipulated, in such a way that half of the names could be perceived fluently and half could not be perceived fluently. As argued, the classic prediction is that participants have a higher tendency to classify fluently perceived names as old, because the environment usually provides feedback that we have indeed encountered a stimulus before when it is perceived more fluently.

However, in between the learning and the recognition phase, a training phase took place, in which participants learned a reverse interpretation of cognitive fluency. To do this, it was necessary to associate a cognitive feeling, that is, the experience of fluency, with a reversed response. Again, the classic finding in the context of a memory task is that higher processing fluency leads to a higher probability of the response “yes, this is an old item”, independently of actual prior exposure (Johnston, Hawley, & Elliot, 1991; Whittlesea, 1993). The goal was to produce circumstances that higher processing fluency leads to a higher probability of the reverse response “no, this is a new item”.

This was accomplished using the mental rotation task of Shepard and Metzler (1971). They showed that people are able to determine that two two-dimensional pictures portray objects of the same three-dimensional shape, even though the objects are depicted in very different spatial orientations. They used shapes that were composed of ten concatenated cubes, with three 90 degree angles in the concatenation. They proposed that people perceive the objects indeed as three-dimensional shapes and to check whether two shapes are identical (though presented in different spatial orientations), people mentally rotate the shapes into the same orientation. The main result was that the time to decide whether two pictures depicted the same three-dimensional object or not was a linear function of the angular difference in portrayed spatial orientation of the two shapes. This strongly supported the mental rotation hypotheses, because rotating shapes with a larger difference in angular orientation obviously consumes more time.

This result was used for the following experiments together with the assumption made in the introduction, that fluency is indexed by response latency. Therefore, comparing and rotating shapes with a large difference in angular orientation not only consumes more time, but should also be experienced to be less fluent than comparing and rotating shapes with a small difference in angular presentation. This is analogous to the ease of retrieval experiments in Part I; fast retrieval was assumed to be experienced as easy and slow retrieval was assumed to be experienced as difficult. Here, a short rotation time is assumed to be experienced as fluent and a long rotation time is assumed to be experienced as non-fluent.

The “trick” is now to associate the experience of fluent rotation in the learning phase with the classic default or the reversed non-default response; let us call the classic response “yes, the two shapes are identical” affirmative and the reversed response “no, the two shapes are not identical” negative. Note the underlying analogy of responding affirmatively “yes, the two shapes are identical” and affirmatively “yes, this is an old item”, which means basically that there is a memory match for the presented item, much the same way that there is a match for two identical shapes. To do this, in deviation from the original mental rotation task, there was a perfect contingency of the angular difference and match of the two shapes.

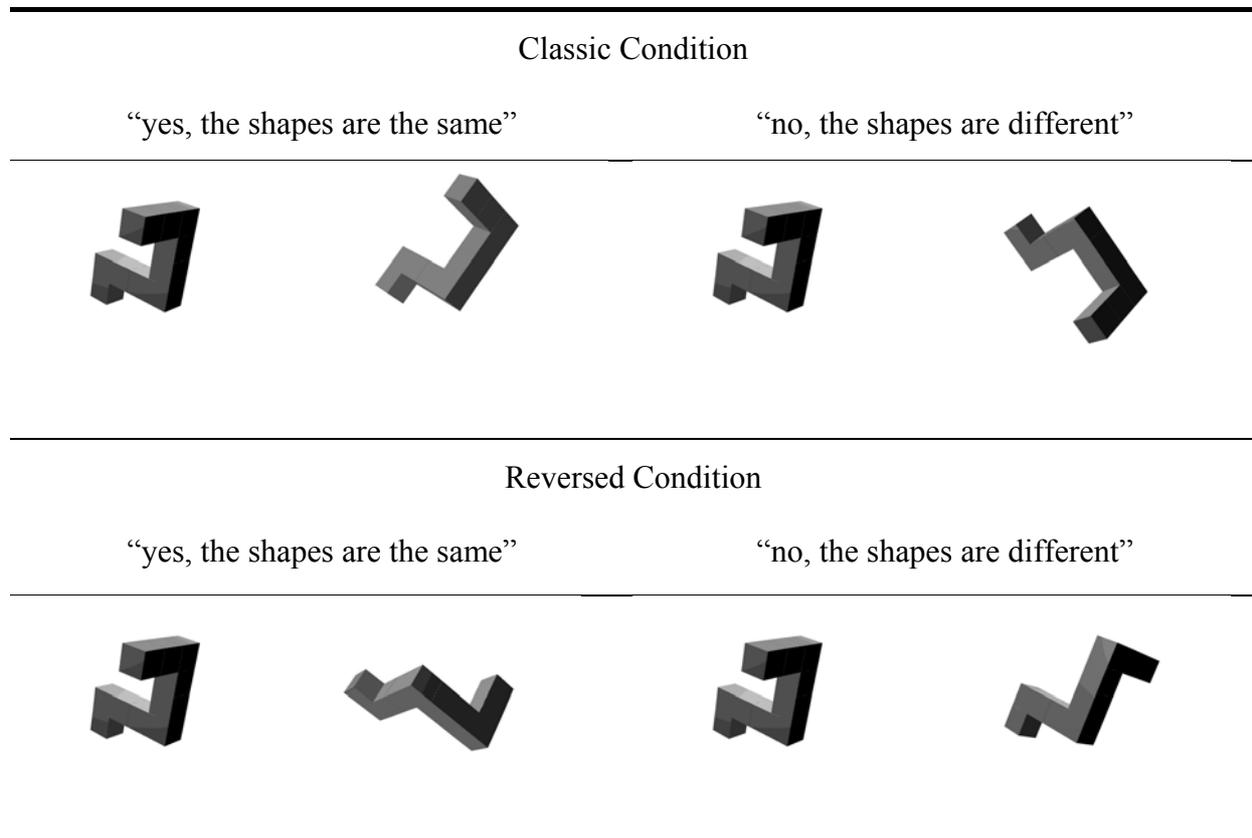


Figure 3. Examples of matching and non-matching three-dimensional shapes in the classic and reversed conditions of Experiment 3 and 4.

In a “classic” condition, two objects with an angular difference of 20 to 40 degrees always depicted the same shape. In contrast, two objects with angular differences of 140 to 160 degrees always depicted two different shapes. The upper panel of Figure 3 presents four exemplars that show the two possibilities of a match and a no-match. After a response (“yes, they are the same”, “no, they are different”), feedback about the correctness of the response was given. Thus, an association of fluency and an affirmative response should result from the training phase. In a “reverse” condition, this pattern is inverted: Two objects with an angular difference of 20 to 40 degrees always depicted two different shapes, and two objects with angular differences of 140 to 160 degrees always depicted the same shape. The lower panel of Figure 3 presents four shapes that show the two possibilities of a match and a no-match in the reversed condition. Through the given feedback, an association of experienced fluency and a negative, non-default, response should result from the training phase.

To facilitate the transfer from the training phase to the recognition phase, perceptual fluency was manipulated in a way that the experience resembles the experience in the training phase. Thus, fluency was manipulated by simply rotating the names as well. High fluency is accomplished by rotating the names in plane clockwise for 20 to 40 and 320 to 340 degrees, which does not make big difference compared to no rotation at all, because the names are only slightly tilted. Low fluency on the other hand is accomplished by rotating the names in plane clockwise for 140 to 160 and 200 to 240 degrees, which turns the names almost upside down. The left part of Figure 4 shows an example of a high fluency name and the right part shows an example of a low fluency name.

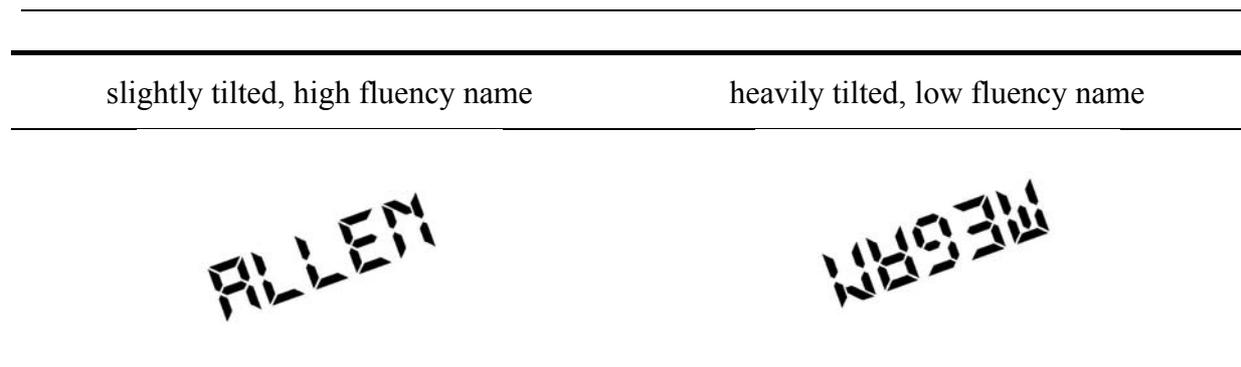


Figure 4. Examples of high and low fluency stimuli used in Experiment 3 and 4.

The general predictions are then straightforward: If people learn in the training phase to associate experienced fluency (non-fluency) with an negative (positive) response, they should have a higher probability to respond negatively (affirmatively) to an item that is fluently (non-fluently) perceived, that is, a name that is only slightly tilted compared to the

almost upside down names. This hypothesis of a reversal of the classic finding in recognition judgments will be tested in Experiment 3.

Experiment 3

To reiterate, the goal of Experiment 3 is to show that the classic interpretation of perceptual fluency (an item being “old”; Johnston, Hawley, & Elliot, 1991) can be reversed by feedback learning. It will be shown that participants judge items that are perceptually fluent with a higher probability as “new”. Following the idea of an interpretation learned via ecological feedback, this should be possible without telling people explicitly how to interpret their cognitive feeling, but by having them learn the interpretation of their cognitive feelings in a training phase.

Method

Participants and Design. 61 (49 female, 12 male) psychology students from the University of Heidelberg participated as a partial fulfillment of a course requirement. They were randomly assigned to one of the two experimental conditions, 31 participants were assigned to the “classic” condition and 30 were assigned to the “reversed” condition. The factor of cognitive fluency (high vs. low) was manipulated within participants.

Material. The material from the Shepard and Metzler (1971) study was conceptually reproduced using the freeware computer tool N3D-PRO by Guliano Cornacchiola. This software allows the creation and manipulation of three-dimensional geometric meshes. With this software, six pairs of concatenated cubical shapes were created. The pairs of shapes only differed in respect to one of the three 90 degree angles in the concatenation of the cubes (cf. Figure 3); thus, the shapes in each pair were similar but easy to distinguish. From each of these twelve shapes, four copies were created by rotating the original shape in depth. Two shapes were rotated 20 and 40 degrees, realizing the short rotation and two shapes were rotated 140 and 160 degrees, realizing the long rotation. Thus, for each original shape there were now four identical shapes created by short rotation (20 and 40 degrees) and long rotation (140 and 160 degrees) and four slightly different shapes, also created by short rotation (20 and 40 degrees) and long rotation (140 and 160 degrees). The total set of shapes consisted of 60 stimuli (cf. Figure 3 for examples). The great advantage of this material is that the difference in angular rotation provides a direct manipulation of cognitive fluency in terms of response latency.

For the memory items, a list of 120 English first names was created. The names were taken from a website that featured the most popular English first names

(<http://www.behindthename.com/top.html>). The following constraints were imposed on the selection of the names; no names with less than 5 letters and more than 6 letters were used and an equal amount of male and female names was required. According to these requirements, the top 1000 names from the year 1989 were searched bottom-up until 120 names were selected for the list. The rationale for selecting English first names in this manner was to generate a pool of stimuli that are very similar in baseline fluency for German students.

For each name a set of eight pictures was created. Four pictures showed a name slightly tilted, by rotating the name by 20, 40, 320, or 340 degrees, making the name easy to read. Four pictures showed the name heavily tilted, almost upside down, by rotating the name 140, 160, 200, or 220 degrees, making the name difficult to read. So, the total set of pictures created from the names (without the non-rotated names) consisted of 960 stimuli.

The experiment was conducted computer-based, using a program written in MS Visual Basic to present instructions, the stimuli and to assess the dependent variables.

Procedure. After arriving in the lab, participants read and signed an informed consent form. They were informed that they would participate in a memory experiment, which required them to learn and to recognize a number of foreign first names. In addition, they were told that the experiment required them to perform a mental rotation task. If participants agreed to take part in the experiment, they signed the form and the experimenter seated them in a cubicle with a personal computer. The experimenter then started the computer program and entered a participant identification number. The assignment to the classic and reversed condition was controlled by this identification number; hence, the experimenter was blind regarding the assignment. The first screen of the computer program presented the general instructions, which were partly redundant to the information given in the consent form. Participants were informed that their task would be to learn a list of names and then to discriminate these learned names later in the experiment from new names. They were also informed that in between the learning and the test phase, they would perform the mental rotation task. The cover story for this task was simple: A certain delay is necessary for a memory experiment and in order to make functional use of this delay, they should provide baseline data for a mental rotation task in another experiment.

After reading these instructions, participants could continue by clicking on a button. On the next screen, they could start the presentation of the first 80 names, which served as the “old” names. The names were presented in the center of the screen. Each name was presented for 1 second with a delay of 1.5 seconds between each name. These 80 names were a new

random selection for each participant from the total list of 120 names described in the materials part.

Following the presentation of the names, the next screen informed participants that they would do an unrelated task to bridge the necessary break for a memory test. Then the mental rotation task was explained. They were told that they had to compare two pictures of geometrical objects and their task would be to decide whether these pictures present the same object or different objects. If they had no further question about this task, they started the mental rotation task. This task implemented now the crucial between participants manipulation.

In the classic condition, a given shape was always compared to the shortly rotated (i.e., a 20 or 40 degree rotated) version of itself. In addition, the direction of comparison was also counterbalanced (i.e., the 20 and 40 degree rotated versions were also compared to the original shape), resulting in four comparisons for each of the six shapes, and therefore in a total of 24 comparisons which required only a short mental rotation to which the correct responses were affirmative: “Yes, it is the same shape”. The given shape was also compared to its yoked shape from the pair, which was similar, but not identical. This non-identical shape was always presented in the 140 and 160 degree rotated version, and the direction of comparison was also counterbalanced. This also resulted in 24 comparisons which required a long mental rotation to which the correct responses were negative: “No, it is a different shape”.

In the reverse condition, following the same logic, from the six basic shapes, 24 comparisons were constructed that required a long mental rotation (i.e., 140 and 160 degrees) to which the correct responses were affirmative: “Yes, it is the same shape”. Conversely, the correct responses to the 24 comparison that required a short mental rotation (i.e., 20 and 40 degrees) were negative: “No, it is a different shape”.

Thus, in both conditions a total of 48 comparisons were required. To ensure the learning effect, these 48 comparisons were repeated, resulting in 96 mental rotation trials. The order of these 96 trials was randomly created for each participant. Participants made their decision by pressing either the “y” or the “-“ key on the keyboard (note that the location of the keys refers to a German keyboard); these keys were clearly marked by a blue and yellow label and the response a key indicated was always visible on the screen: On the left side of the screen a label was visible “Yes, it is the same shape” and on the right side of the screen a label was visible “No, it is a different shape”. After each response, both shapes disappeared and participants received feedback whether their response was correct or incorrect. In the case

of a correct response, a green label appeared with the text “Correct: Those were identical shapes!” for two seconds; in the case of an incorrect response, a red label appeared with the text “Wrong: Those were different shapes!”. After the two seconds, there was a break of one second and then a label appeared with the text: “Next Trial”, followed by the next two shapes in the random order. The responses and the response latencies were recorded during all 96 trials.

When participants finished all 96 trials, another instruction screen appeared, informing them that the actual recognition test was about to start: Their task would be to decide whether a presented name was “old”, (i.e., it was present in the learning phase) or whether it was “new” (i.e., it was not present in the learning phase). In addition, the instructions told participants that the names would be randomly tilted (cf. Figure 4). If participants had no further questions, they could start with the recognition task.

For each participant, 40 names were randomly selected as targets from the 80 names in the learning phase. The remaining 40 names from the original 120 names list served as foils. Half of the targets and foils were assigned the status of high fluency items; they were presented slightly tilted, that is, rotated 20, 40, 320 or 340 degrees. The rotation angle was determined randomly. The remaining 20 foils and 20 targets were assigned the status of low fluency items and they were presented heavily tilted, that is, almost upside down by rotating them 140, 160, 200, or 220 degrees. The sequence resulting from this orthogonal combination of high-low fluency by old-new names was also a new random order for each participant. The tilted versions of a name were taken from the set of 960 rotated names described in the Materials section.

Participants responded to each name with the same keys as in the training phase, that is, the “y” key for old names and the “-“ key for new names. They received no feedback for their decisions and after each keystroke, the next trial started with a delay of one second. The computer recorded the decisions and the response latencies.

After completing the 80 trials, participants were probed for suspicion and they could answer in a free format using the keyboard. Three questions were asked and participants could respond to each of those in a separate textbox: “Besides the given explanations, do you have any ideas about other purposes of the experiment?”, “Do you have any ideas about the purpose of the mental rotation task? Did you notice anything special about this task?”, and “Did you notice anything about the names presented in the recognition task?”.

Once they had completed their answers, the next page gave a written explanation about the purpose of the study; after that, they were prompted to contact the experimenter who thanked and debriefed them further if there were any more questions.

Results

Mental rotation task – Latencies and error rates

The analysis of the data from the mental rotation task has to show that people indeed respond faster to shapes with a small angular difference independent of the given response. In other words, this analysis should show that greater fluency (indicated by faster responses) is associated with different responses in the classic and the reversed condition.

Table 5

Mean percentages of correct responses and mean response latencies (rounded to full ms) in the mental rotation task of Experiment 3 as a function of angular difference and condition (standard deviations in parentheses).

Condition	Percentage Correct		Response Latency	
	Angular Difference		Angular Difference	
	small	large	small	large
Classic	92.20 (0.08)	78.36 (13.72)	2252 (798)	3812 (2049)
Reversed	78.13 (21.72)	70.42 (14.21)	3820 (1953)	4597 (2475)

Note. In the classic condition, a small angular difference requires an affirmative response, whereas the reverse condition requires a negative response for a small angular difference.

The mean response latencies and the percentage of correct responses are displayed in Table 5, separately for small angular difference (20 and 40 degrees) and large angular difference (140 and 160). Remember, in this task there is a perfect contingency of rotation and match of the two shapes. In the classic condition, shapes with a small angular difference always matched, requiring an affirmative response. In the reversed condition, shapes with a large angular difference always matched.

Let us first consider the response latencies. To analyze these data, a 2 (classic vs. reversed) x 2 (small vs. large angular difference) mixed ANOVA was employed. The results from this analysis are clear cut: First, there is a large main effect for angular difference, $F(1, 59) = 58.32, p < .001$. Participants were much faster to decide whether two shapes matched or not when the shapes had only a small angular difference ($M = 3023, SD = 1669$) compared to

a large angular difference ($M = 4198$, $SD = 2284$). Note that this main effect is independent of the condition; that is, whether an affirmative or negative response was required. But there was also a main effect for the between factor condition; $F(1, 59) = 6.37$, $p < .05$. Overall, people in the classic condition responded faster ($M = 3031$, $SD = 1400$) than people in the reversed condition ($M = 4209$, $SD = 2172$). In addition, there was an interaction of the two factors, $F(1, 59) = 6.52$, $p < .05$. As Table 5 shows, this interaction results from the fact that the increase in response latencies for people in the classic condition is much higher than for those in the reversed condition. The results from the response latencies indicate that the classic condition was slightly easier than the reversed condition, resulting from the fact that an initial failure to match to shapes with a small angular difference leads in many cases to further rotation attempts before a negative response is made. But both the main effect for condition and the interaction are rather small in comparison to the strong effect of angular difference. Thus, in the classic condition greater fluency (indicated by response latency) was associated with an affirmative response, whereas in the reversed condition, greater fluency was associated with a negative response.

The same mixed ANOVA was used to analyze the percentage of correct responses for small and large angular difference in the classic and reversed condition. First, there is the same main effect for the angular difference, $F(1, 59) = 24.87$, $p < .001$, indicating that participants made less errors for shapes with a small angular difference ($M = 85.28$, $SD = 17.64$) than for shapes with a large angular difference ($M = 74.45$, $SD = 14.41$). Again, this effect is independent from the correct response required, affirmative or negative. The main effect of condition was stronger for the correct responses than for the response latencies, $F(1, 59) = 11.63$, $p < .005$. Participants in the classic condition had a higher percentage of correct responses ($M = 85.28$, $SD = 9.91$) than participants in the reversed condition ($M = 74.27$, $SD = 14.89$). Parallel to the response latencies, this indicated that the reversed condition was in general more difficult than the classic condition. The interaction was not significant, $F(1, 59) = 2.02$, ns.

All in all, the data from the training phase suggests that fluency as indexed by the time necessary for mental rotation was successfully associated with differential responses in the two conditions, a negative “no” response in the reversed, and an affirmative “yes” response in the classic condition. The fact that the reversed condition seems to be more difficult than the classic condition poses a small problem. However, this is due to the nature of the task: A match and an affirmative response terminates the rotation immediately, which leads to many quick and correct responses in the classic condition for shapes with a small angular

difference, whereas these exact pairs lead to a no-match in the reversed condition; however, this does not automatically lead to a negative response, but can also lead to further rotation attempts.

Recognition test – Response latencies

Johnston, Hawley, and Elliot (1991) reported that in recognition tasks, old items led to faster responses than new items. They took this as evidence that fluency can serve as a basis for recognition judgments, independent of actual recall or recognition. Here, as discussed, response latency will be used as an index for fluency.

The response latencies for the names were classified following three factors: First, status (i.e., whether a name was actually old or new), second, rotation (i.e., whether the name was only slightly tilted or heavily tilted) and third, the kind of response (i.e., whether the response was affirmative or negative). Thus, this classification resulted in eight subgroups of latencies. The data from three participants were excluded because they responded uniformly affirmative or negative. As this uniform response pattern is indicative for no engagement in the task, the data from these participants were excluded from the analysis. The remaining 58 data sets were analyzed using a mixed ANOVA, with status (old vs. new), rotation (slight vs. heavy) and response (yes vs. no) as within factors and condition (classic vs. reversed) as between factor. The results from this seemingly complex analysis are nevertheless straightforward.

First, there was a clear main effect for the rotation of a name, $F(1, 56) = 77.63, p < .001$. As expected, participants responded much faster to slightly titled names ($M = 1263, SD = 420$) than to heavily tilted names ($M = 1751, SD = 693$). Therefore, the rotation indeed manipulated the fluency of the names, as measured by the response latencies.

Second, there was also a main effect for the status of the name, $F(1, 56) = 8.63, p < .01$. Participants responded faster to an old name ($M = 1473, SD = 561$) than to new names ($M = 1542, SD = 678$), independent of the type of response. This effect is of interest for models of memory, but of lesser interest for the present argument.

A third main effect was due to the factor response, $F(1, 56) = 10.53, p < .01$. Participants were faster to give an affirmative response ($M = 1464, SD = 577$) than a negative response ($M = 1550, SD = 662$). This effect matches the idea about two routes to recognition judgments (Mandler, 1980, Kelley & Jacoby, 2000): Recollection and an experience-based feeling of familiarity. The sum of affirmative responses is generally faster, because recollection is fast and only possible for affirmative responses, This explanation is corroborated by a strong interaction of status and response, $F(1, 56) = 25.70, p < .001$. This

interaction is mainly due to fast affirmative responses to old names ($M = 1355$, $SD = 469$), whereas affirmative responses to new names (i.e., errors) were much slower ($M = 1573$, $SD = 653$). Negative responses showed the reverse pattern; they were slightly slower for old names ($M = 1590$, $SD = 621$) than for new names ($M = 1511$, $SD = 703$).

Besides these four strong effects, there were a couple of smaller effects, but none of these effects was reliable on a standard error level (all F s < 2.90 ; p s $< .10$) and because they do not relate to any relevant hypotheses, they are not reported here.

The analysis of the response latencies shows that participants took more time to respond to heavily tilted names, which serves as a manipulation check for the concept of fluency. In addition, the response latencies follow the pattern that is predicted by a two route (remember vs. know; recollection vs. feeling) model of recognition judgments (Mandler, 1980; Johnston, Hawley, & Elliot, 1991, Kelley & Jacoby, 2000).

Now the interesting question is, did participants in the two conditions interpret fluency differentially? That is, did they learn a new interpretation of the cognitive feeling in the reversed condition, based on the feedback provided in the training phase? To answer this question, the pattern of affirmative and negative responses was analyzed using signal detection theory (SDT).

Recognition test – SDT-Analysis

SDT is a very useful tool for analyzing data from all tasks that require dichotomous responses of the form “yes, signal present” and “no, signal not present”, hence signal detection theory. An excellent introduction is given by Swets, Dawes, and Monahan (2000; for computational issues, refer to Stanislaw & Todorov, 1999), but a short summary about the terminology and the meaning of parameters in the theory is called for: A “yes” response when a signal was actually present is considered a “hit”, whereas a “yes” response is considered a “false alarm” when no signal was present. Over a large number of signal and no-signal trials, hits and false-alarms probabilities can be computed from the hit and false-alarm rates. From these probabilities, a SDT analysis delivers two parameter estimates: An estimator for the ability to discriminate between the presence and the absence of a signal (i.e., signal vs. no-signal) and an estimator for the response tendency; that is, favoring one response (“yes, signal present” vs. “no, signal absent”) above the other, independent of actual discrimination ability.

In the case of a memory task, SDT estimates the actual memory capability and the tendency to give an affirmative or a negative answer, independent of the actual memory capability. For the present analysis, the parameter d' was used as an estimator for discrimination ability and the parameter β as an estimator for the response tendency. There

are alternatives to these estimators, depending on the structure of the given data; for the present purpose, however, d' and β are optimally suited (Stanislaw & Todorov, 1999).

For the present approach, the parameter of β is of eminent interest, because it measures the tendency to favor a response given a stimulus independently of the actual memory performance. Consider again the possibility that you can immediately recognize a name; a feeling of fluency or familiarity should not play a role in this decision, because you “remember” the name. Such decisions should mainly influence d' , because they lead automatically to fast and correct responses (see also the analysis of the response latencies). On the other hand, when there is no direct recall, cognitive feelings of fluency play a role, and if the rotation of the names gives rise to such experiential states, as is indicated by the response latencies, this should be detectable in different β 's for the slightly and heavily tilted names. Moreover, because participants in the reversed condition should interpret the experienced fluency differently from those in the classic condition, an interaction of condition (classic vs. reversed) and rotation (slightly vs. heavily) is expected.

The same three participants as in the analysis of the response latencies were excluded from the following analysis, because they responded uniformly to all items. Thus, 29 participants remained in the classic and 29 participants remained in the reversed condition.

Analysis of d' . To reiterate, d' measures the ability to discriminate old names from new names. High values of d' indicate high discrimination ability. Theoretically, d' has a range of zero to infinity; if d' equals zero, the discrimination is at chance level (i.e., the hit and the false alarm rate are the same) and d' reaches infinity if no errors at all are made (i.e., hit rate equals one and the false alarm rate zero). The mean of the estimates of d' for slightly (rotation: 20, 40, 320, and 340 degrees) and heavily tilted (rotation: 140, 160, 200, and 220 degrees) names in the two conditions are displayed in Table 6.

Table 6

Mean estimates of d' in Experiment 3 as a function of fluency and condition (standard deviations in parentheses)

Condition	Fluency	
	high	low
classic	0.78 (0.54)	0.45 (0.55)
reversed	0.99 (0.71)	0.57 (0.51)

Note. Higher d' values indicate a better discrimination ability.

A 2 (condition: classic vs. reversed) x 2 (rotation: slightly vs. heavily tilted) mixed ANOVA was used to analyze these data. This analysis yielded only a strong main effect for rotation, $F(1, 56) = 17.08, p < .001$, indicating that participants discriminated old from new names much better when the names were only slightly tilted ($M = 0.88, SD = 0.63$) compared to heavily tilted names ($M = 0.51, SD = 0.53$). This effect is of lesser interest and only corroborates the fact that heavily tilted name were more difficult to discriminate.

Analysis of β . Again, β measures response bias, independent of the actual discrimination ability; that is, a general tendency to respond affirmatively or negatively over a large number of signal and no-signal trials. Because β is based on a ratio, a value of 1 indicates no bias at all, that is, an equal number of yes and no responses. Values smaller than 1 indicate a tendency to respond affirmatively and values greater than 1 indicate a tendency to respond negatively. The means of the estimators of β in the two conditions are displayed in Table 7, separately for slightly and heavily tilted names.

Table 7

Estimates of β in Experiment 3 as a function of fluency and condition (standard deviations in parentheses)

Condition	Fluency	
	high	low
classic	1.18 (0.44)	1.11 (0.31)
reversed	1.73 (1.16)	1.25 (0.58)

Note. β estimates greater than 1 indicate a tendency to respond negatively to a given item, estimates smaller than 1 indicate a tendency to respond affirmatively.

As this table shows, there was a general tendency to respond negatively. This effect might be due to some feature of the overall task. But for the present argument, differences between the classic and the reverse conditions to fluent and non-fluent stimuli were of interest. To test these, the same mixed ANOVA as for d' was used. This analysis also shows a main effect for rotation, $F(1, 56) = 5.88, p < .05$, indicating that participants had an overall tendency to respond more often negatively to slightly tilted names ($M = 1.45, SD = 0.91$) than to heavily tilted names ($M = 1.18, SD = 0.47$). In addition, there was also a main effect for condition, $F(1, 56) = 5.46, p < .05$. This effect is due to the fact that participants in the

reversed condition had generally a higher tendency to respond negatively ($M = 1.49$, $SD = 0.74$) than participants in the classic condition ($M = 1.15$, $SD = 0.30$). Besides these two main effects, there was also the predicted interaction of condition and rotation. Yet, this interaction effect was not significant on a standard error level, $F(1, 56) = 3.15$, $p < .08$. Nevertheless, for an easier interpretation of this interaction in the presence of two main effects, the means are displayed in Figure 5, corrected for the main effects (Rosenthal & Rosnow, 1991; Rosnow & Rosenthal, 1995).

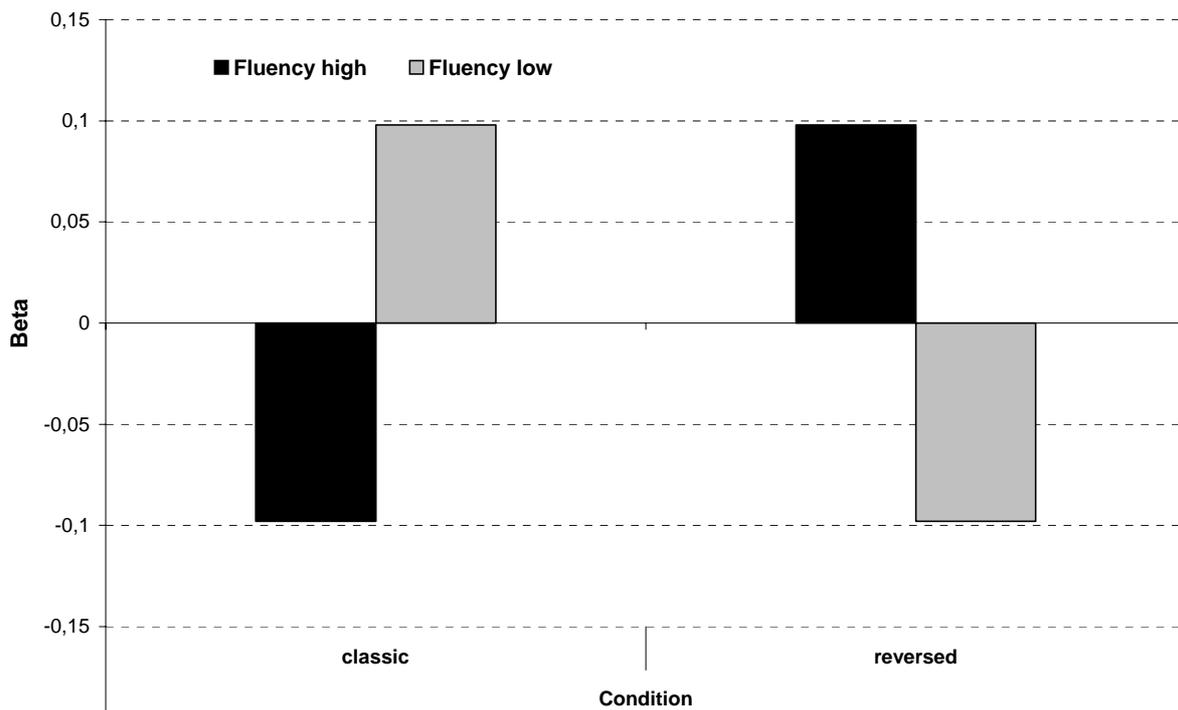


Figure 5. Estimates of β in Experiment 3 as a function of fluency and condition (corrected for main effects). Negative values indicate a tendency to respond affirmatively, positive values indicate a tendency to respond negatively.

The pattern shows the predicted effect; participants who learned in the training phase to associate fluency with an affirmative response had a tendency to respond affirmatively to slightly tilted names, which are fluently perceived (as indicated by the response latencies), and a tendency to respond negatively to the heavily tilted names. On the other hand, participants in the reversed condition, who learned in the training phase to associate fluency with a negative response, had a tendency to respond affirmatively to the heavily tilted names and negatively to the slightly tilted names.

Discussion

The goal of Experiment 3 was to show that the impact of a cognitive feeling is dependent on an applicable interpretation and that such interpretations are learned via ecological feedback in a given context. To support this idea, a well known effect was used: In the context of a memory task, high fluency items have a higher probability to be classified as being old compared to low fluency items. The underlying rationale was that stimuli which are processed more fluently elicit a feeling of familiarity and therefore that stimulus is classified as old. Here, the idea of ecological feedback was introduced. In a real world setting, you might receive the feedback that stimuli (faces, voices, pictures, statements etc.) which elicit fluency have indeed been encountered before. To make this point, a feedback task was created that associated fluency (i.e., short/easy or long/difficult mental rotation) with an affirmative or a negative response. That is, an attempt was made to reverse the strong and robust finding (Johnston, Hawley, & Elliot, 1991; Whittlesea, 1993) that greater processing fluency leads to classifications of a stimulus as old; that is, leading to affirmative responses to stimuli that are processed fluently.

It was assumed that the conceptual similarity of fluency evoked by the mental rotation task in the training phase and the rotation of the names in the recognition test would suffice to bring the reversal about. An easier way would have been to give feedback in the exact same task in a training phase that is used in the following recognition test. But this would have limited the generality of the result; it would have been a much weaker point to show that people can learn the interpretation of a cognitive feeling and apply this interpretation within the same task. The present paradigm demonstrates the importance of the fluency construct as it is the only connection between the mental rotation task and the seemingly unrelated recognition test.

Moreover, the material of the mental rotation task offered an excellent possibility to manipulate fluency. As mental rotation takes time and longer rotation is more difficult, fluency can be manipulated by the angular difference in rotation and can be measured by the response latencies. Similarly, in the actual recognition task, fluency was manipulated by rotating the names, making them fluently or non-fluently perceivable. The response latencies for the stimuli indicate that these manipulations were successful. Especially in the mental rotation task it is important to note that shapes with a small angular difference elicited faster affirmative and in the reversed condition, also faster negative responses. Thus, the mental rotation task in the training phase was successful in associating fluency with an affirmative or a negative response. One minor problem related to the mental rotation task was the greater

difficulty of the task in the reversed condition in general. But the analysis of the SDT measure d' , a measure of actual memory performance, shows no difference between the classic and the reversed condition. Thus, any effects of the differential training are not attributable to greater mental effort.

However, the transfer to the actual recognition test was only partly successful. In the classic condition, it was expected that people who learned to associate fluency with an affirmative response, should also show a higher tendency to respond affirmatively to slightly tilted names. In the reverse condition, people learned to associate fluency with a negative response. Thus, the opposite pattern was expected for the slightly and heavily tilted names. The measure β from signal detection theory was used to test these predictions. The prediction was only partly supported. In the classic condition, participants showed hardly any differential response tendency. But participants in the reversed condition showed a strong tendency to respond negatively to slightly tilted names. This indicates indeed a reversal of the classic interpretation of fluency in recognition judgments; and corrected for the main effects, the pattern shows the predicted interaction effect. However, the effect was not significant on a standard alpha level; yet, the effect size ($d = 0.47$, $r = .23$) indicates that this is rather due to a lack of power of the present study (Cohen, 1990, 1992).

Nevertheless, this also points to a small weakness in the employed design. Let us assume a simplified memory model with two possible ways to come up with a recognition judgment: a recollection-based “know” judgment and a feeling-based “remember” judgment (Gardiner & Richardson-Klavehn, 2000): Cognitive feelings only play a role if there is no recollection of an item. And recollection favors old items, because you can be sure to recognize a name, but you cannot be absolutely sure that an item was not present in the learning phase (Mandler, 1980; Kelley & Jacoby, 2000). Yet, the problem is not severe, because due to the difficulty of the task, there was enough room for errors and uncertainty. Even so, this asymmetry offers an explanation for the somewhat unbalanced results.¹

It is noteworthy that the response latency data also support the idea that there are two routes to recognition judgments, which is, classifying a stimulus as old. The route of direct retrieval or recollection is fast and leads to a state of knowing that an item is old (a “remember” judgment; Kelley & Jacoby, 2000). Hence, cognitive feelings do play a lesser role in such decisions. This is visible in the fast affirmative responses to old items, independent of any fluency manipulation. Second, there is the feeling-based decision, which

¹ Testing the simple effects of the β interaction for high and low fluency names, there is no difference in the classic condition, but a highly significant effect for the reversed condition. Nevertheless, the displayed ANOVA interaction is the most instructive way to interpret the results.

uses the peripheral cue of fluency (a “know” judgment; Kelley & Jacoby, 2000). This leads to the prediction that negative responses should drive the effect, because uncertainty allows for the greater impact of cognitive feelings. The affirmative responses are an amalgam of true recollection, which is less open to feeling-based inferences and feeling-based decision. For negative responses, on the other hand, there is always uncertainty involved. And indeed, the observed interaction is mainly due to participants’ high rate of negative response to only slightly tilted items (i.e., fluently processed items) in the reversed condition. The restricted influence of cognitive feelings to judgments under uncertainty is most pointedly visible in the “tip of the tongue” phenomenon (Hart, 1966). If you do remember a stimulus, there is no room and need for cognitive feelings. But when there is no absolute state of knowing, cognitive feelings come into play.

All together, Experiment 3 gave first evidence that people can learn a reversed interpretation of fluency; they learned to associate a response with an experiential state that was elicited by a cognitive operation, namely, mental rotation. Yet, the results missed a standard level of significance, and one might argue that they are not totally convincing in that respect (but again, refer to Cohen, 1990, 1992 for a discussion of the .05 level in null-hypothesis testing). In addition, the design lacked a control group. It is not clear how participants would have reacted to the rotated names (slightly or heavily tilted) given no particular feedback in the training phase. Therefore, Experiment 4 tried to replicate the results and implemented a control group.

Experiment 4

Besides replicating the results of Experiment 3, the goal of Experiment 4 was to introduce a control group. The question was how participants would respond in the recognition phase if they learned nothing in the training phase. But what constitutes a good control group for the present purpose? For example, one could just omit the training phase and directly go to the recognition test. But this would give a control group a recency advantage; plus, they could wield more free cognitive resources, because they would not perform the rather long 96 trials of the mental rotation task. From such considerations, a suitable control condition was created by changing the mental rotation task in such a way that no contingency between fluency and the correct response existed. That is, new pairs of shapes were created for which half of the shapes with a small angular difference matched and half of these shapes did not match. The same was true for shapes with a large angular difference. Thus, the control condition contained 48 pairs that required only short mental rotation to

determine whether the shapes were identical and 48 pairs that required long mental rotation to determine whether the shapes matched or not. While the classic and reversed conditions had a perfect contingency of angular difference and correct response (i.e., shapes with small/large angular difference always matched and vice versa), the control condition had therefore no contingency at all, but was nevertheless required to perform the same task.

One additional manipulation was added as a between variable in Experiment 4. It was argued that cognitive feelings only play a role when there is uncertainty, measured by response latencies. First, the direct route of remembering is used to determine whether a stimulus is old or new. That is, the first time is devoted to actual memory search. Only when this search yields no result, the secondary route via cognitive feelings is used. For the discussion above follows that only after a certain time the fluency cue is used in making a decision about the status of a name as old or new.

To show that the use of the fluency cue is indeed dependent on time, a time pressure manipulation was included in Experiment 4. In the final recognition test, half of the participants were forced to answer within two seconds, whereas the other half had no time constraints. The idea is that the use of the fluency cue, or cognitive feelings in general, is dependent on available time; here, to terminate an unsuccessful memory search and then to use the present experiential state. The claim is that the proposed process is not as automatic as an affective feelings-as-information approach (Zajonc, 1980), or a direct influence route. It consumes time and resources and therefore, it is unlikely that the interpretation is acquired phylogenetically. This logic is also similar to a design by Neely (1976, 1977), who showed that in semantic priming tasks, new associations can be learned. But they exert an influence only when long SOAs were used. In the case of short SOAs, the effect of newly learned associations vanishes. The same assumption was made here: If there is no time, the influence of the newly learned association, for example, of fluency and a negative response, vanishes.

Method

Participants and Design. 86 students (34 women, 52 men) of the University of Heidelberg participated for a payment of 5 €. They were randomly assigned to one of the six experimental conditions, resulting from the orthogonal combination of the between participants variable time pressure (time pressure present vs. no time pressure present) and learning condition (classic vs. reversed vs. control). The factor of cognitive fluency was again manipulated as a within-participants variable.

Materials and Procedure. The 120 English first names for the learning phase, the 60 three-dimensional shapes for the mental rotation task, and the 960 pictures (four slightly tilted

and four heavily tilted versions of each name) for the recognition phase were identical to Experiment 3. The Visual Basic program was slightly modified to implement the control condition and the time pressure condition

After arriving in the lab, participants read and signed an informed consent form. Then the experimenter seated them in a cubicle and started the computer program. The learning phase was in all conditions identical to Experiment 3. The following mental rotation task was also identical for the classic and the reversed condition, associating a short (long) mental rotation and an affirmative (negative) answer in the classic and vice versa in the reversed condition. However, the control condition showed no contingency between required rotation and the required answer. A given shape was compared twice to a slightly rotated version (i.e., 20 and 40 degrees) of itself and twice to a heavily rotated version (i.e., 140 and 160 degrees) of itself. Similarly, the comparison to the non-identical shape of a pair involved twice the shortly and twice the heavily rotated version of the shape. The direction of comparison was counterbalanced as well and each comparison was presented twice. This resulted in 24 comparisons that required a short mental rotation and an affirmative response, 24 comparisons that required a long mental rotation and a negative response, 24 comparisons that required a short mental rotation and a negative response, and 24 comparisons that required a long mental rotation and an affirmative response.

Subsequent to the mental rotation task, participants continued with the recognition phase. In the time pressure conditions, participants had to make their decision within two seconds; a blue status bar appeared below the presented name indicated the elapsed time and after the two seconds, the status bar and the name disappeared and a label prompted participants to respond immediately. The no time pressure condition was identical to the recognition phase in Experiment 3. The responses and the latencies were recorded as dependent variables. After completing the 80 trials, a final screen debriefed participants about the true purpose of the study and the experimenter thanked and paid them.

Results

The data from one participant in the control condition was deleted prior to analysis because the experimenter noted that she did not follow the instructions properly.

Mental Rotation Task

As in Experiment 3, response latencies served as an index of cognitive fluency and an analysis of the response latencies should show that greater fluency is associated with differential responses in the classic, control and the reversed condition. Due to an error in the

computer program, the responses and their latencies were not recorded in the control condition. This does not pose a serious problem, because the control condition conceptually replicates the original study by Shepard and Metzler (1971) and there is no reason to suspect that the data from the control condition should deviate from the classic finding that small angular differences elicit faster responses than large angular differences. And because there was no contingency of angular difference and a affirmative or negative response, no association of fluency and a response should occur. Thus, the analysis of the mental rotation data is restricted to the classic and reversed condition, without the data from 25 participants in the control condition.

Table 8

Mean percentages of correct responses and mean response latencies (rounded to full ms) in the mental rotation task of Experiment 4 as a function of angular difference and condition (standard deviations in parentheses)

Condition	Percentage Correct		Response Latency	
	Angular Difference		Angular Difference	
	small	large	small	large
classic	93.67 (0.08)	83.41 (15.56)	2432 (1278)	4311 (2346)
reversed	85.58 (16.80)	82.37 (13.98)	3742 (1434)	4348 (1637)

Note. In the classic condition, a small angular difference requires an affirmative response, whereas the reverse condition requires a negative response.

Mean latencies and the percentage of correct responses are displayed in Table 8. This Table shows by large the same pattern that was visible in Experiment 3. Participants were in general faster to respond to shapes with a small angular difference ($M = 3087$, $SD = 1499$) than to shapes with a large angular difference ($M = 4329$, $SD = 2003$). This main effect was again highly significant, $F(1, 50) = 57.52$, $p < .001$. It is important to note that this effect is independent of the required response, that is, affirmative and negative answers were both associated with greater cognitive fluency as indexed by the response latencies. The main effect for the condition was not significant, $F(1, 50) = 2.25$, ns. But there was a strong interaction of condition and angular difference, $F(1, 50) = 15.12$, $p < .001$. This interaction is due to the same fact as in Experiment 3: The increase in response latency for people in the classic condition is much higher than for those in the reversed condition. The strongest effect is nevertheless caused by the angular difference of the rotation task; as such, fluency was

associated with an affirmative response in the classic and a negative response in the reversed condition.

The parallel analysis of the percentage of correct responses replicates Experiment 3 as well. Participants made fewer errors for shapes with a small angular difference ($M = 89.62$, $SD = 13.61$) than for shapes with a large angular difference ($M = 82.89$, $SD = 14.66$), $F(1, 50) = 13.02$, $p < .001$. The main effect for condition was not significant, $F(1, 50) = 1.81$, ns, but the interaction was marginally significant, $F(1, 50) = 3.57$, $p < .07$. This interaction is again caused by the non-symmetrical decrease in correct responses from small to large angular difference in the classic condition. Overall, the analysis of the correct responses indicates that the reversed condition was slightly more difficult, and still more important, responding to a pair of shapes with a large angular difference is more difficult (i.e., less fluent), as indicated by the higher error rate.

Recognition test - Response latencies

The recognition test yielded response latencies as well as correct and incorrect responses to the stimuli from the training phase which were transformed to Hit- and False-Alarm-Rates and ultimately, to SDT parameter estimates.

The response latencies were again classified according to status (i.e., whether a name was actually old or new), fluency (i.e., whether the name was only slightly tilted or heavily tilted) and the kind of given response (i.e., whether the response was affirmative or negative). Four participants had missing values in the subgroups related to affirmative or negative responses; that is, they responded uniformly “yes” or “no” and their data were excluded from the following analysis. The remaining 81 data sets were analyzed using a mixed ANOVA with status, fluency and response as within factors and condition (classic, reversed, control) and time pressure (present vs. not present) as between factors.

This rather complex analysis yielded basically the same results as the analysis from Experiment 3; but naturally, the introduction of time pressure had a profound influence on the response latencies. Any significant main effect was qualified by an interaction with time pressure, as such that under time pressure the existing effects became smaller. Therefore, interactions with time pressure are omitted in the report. The effect of time pressure itself was highly significant, $F(1, 75) = 72.84$, $p < .001$, showing that participants followed the instructions and responded a lot faster in the time pressure condition ($M = 1373$, $SD = 233$) than in the no time pressure condition ($M = 2485$, $SD = 837$). The maximum latency in the time pressure condition was 2114 ms, indicating that basically no one exceeded the given time frame of 2 seconds; the maximum latency in the no time pressure condition was 4545

ms, showing that participants pondered a decision much longer when given the time. This effect shows the success of the time pressure manipulation.

Besides this strong effect of time pressure, the same three main effects as in Experiment 3 emerged: First, participants responded faster to old than to new names ($M = 1913$, $SD = 1115$, vs. $M = 2002$, $SD = 1105$), $F(1, 75) = 6.37$, $p < .02$. Second, participants gave faster affirmative than negative responses ($M = 1874$, $SD = 1018$, vs. $M = 2041$, $SD = 1191$), $F(1, 75) = 20.11$, $p < .001$. And third, participants responded faster to high fluency names than to low fluency names ($M = 1571$, $SD = 699$, vs. $M = 2344$, $SD = 1297$), $F(1, 75) = 157.04$, $p < .001$. This last strong effect is the manipulation check that slightly titled names were processed more fluently than heavily titled names, as indexed by the response latencies.

In addition, the same strong interaction of status (old vs. new) and given response (yes vs. no) as in Experiment was found, $F(1, 75) = 40.14$, $p < .001$. Affirmative responses to old names were faster ($M = 1667$, $SD = 814$) than affirmative answers to new names ($M = 2081$, $SD = 1153$). The reverse is true for negative responses; they were faster for new names ($M = 1924$, $SD = 1056$) than for old names ($M = 2159$, $SD = 1308$). Different from Experiment 3, this interaction was qualified by the fluency of a name, resulting in a significant three-way interaction, $F(1, 75) = 15.71$, $p < .001$. The pattern of means suggests that the described two-way interaction is only true for high fluency items, but not for low fluency items; this is instructive, because if a name is heavily titled, direct recognition of the name is hindered by the necessary mental rotation. All other effects, besides those related to time pressure, were of minor significance, all $F_s < 3$, and do not relate to the present hypotheses.

In summary, the analysis of the response latencies demonstrates that the manipulation of fluency was successful and that the data follow a pattern that is in line with the assumption of two routes to recognition judgments: A fast and efficient recollection route and a slow and malleable feeling route, based on cognitive fluency. Having established this, the interesting question is now the use of this feeling and its interpretation as a function of learning in the mental rotation task and time pressure during the recognition test.

Recognition test - SDT-Analysis

The same four participants as in the analysis of the response latencies were deleted from the SDT analysis. In addition, six participants showed a negative d' value. Their data were also deleted prior to any inference analysis, because a negative d' value is only possible if either the response category were switched or deliberately false responses are given. The loss of data was almost equally distributed across the six cells of the design: In the no time pressure conditions, 13, 11, and 14 participants remained in the classic, control and reversed

condition, respectively. In the time pressure conditions, 13, 11, and 12 participants remained in the classic, control, and reversed condition, respectively. All reported F -values are therefore based on Type III sums of squares.

Analysis of d' . The mean estimates of d' in the six between conditions for high and low fluency items (i.e., names with small or large rotation angle) are given in Table 9. A 2 (time pressure: present vs. not present) x 3 (condition: classic, control, reversed) x 2 (fluency: high vs. low) ANOVA was used to analyze these results.

Table 9

Mean estimates of d' in Experiment 4 as a function of fluency, condition and time pressure (standard deviations in parentheses)

Condition	no time pressure		time pressure	
	high fluency	low fluency	high fluency	low fluency
classic	1.17 (0.49)	0.98 (0.58)	1.05 (0.45)	0.82 (0.40)
control	0.90 (0.35)	1.16 (0.48)	0.68 (0.48)	0.66 (0.51)
reversed	1.06 (0.54)	0.71 (0.44)	0.80 (0.59)	0.77 (0.37)

Note. Higher d' values indicate a better discrimination ability.

This analysis showed only a significant effect for time pressure, $F(1, 68) = 4.60, p < .05$, indicating that without time pressure, participants discriminated better between old and new names. In addition, there was a marginal interaction of condition and fluency, $F(2, 68) = 2.75, p < .08$. Participants in the control group show a greater d' for low fluency items, whereas the classic and reversed condition show a greater d' for the high fluency items. This effect is unexpected and is inconsistent with the findings from Experiment 3; yet, effects on the discriminability are not at the heart of the present study and the result will not be discussed further.

Analysis of β . Of greater interest for the question under investigation is the tendency to respond affirmatively or negatively to a given item independent of actual discrimination ability, as measured by the parameter β . The mean estimates of β in the six between conditions for high and low fluency items are given in Table 10.

The same 2 (time pressure: present vs. not present) x 3 (condition: classic, control, reversed) x 2 (fluency: high vs. low) mixed ANOVA as for d' estimates was used to analyze these data. For the between participants variables, only a main effect for time pressure

emerged, $F(1, 68) = 4.05, p < .05$. As Table 10 shows, participants without time pressure had generally a greater tendency to respond negatively ($\beta = 1.30$) than participants in the time pressure condition ($\beta = 1.11$). The only other significant effect was the expected three-way interaction of time pressure, condition and fluency, $F(2, 68) = 3.01, p < .06$. However, this effect is not significant on a standard alpha level. Yet, performing the same analysis on log transformed β estimations, a standard technique which minimizes the undue influence of outliers in skewed distributions (McClelland, 2000), results in a significant effect on a standard alpha level, $F(2, 68) = 3.28, p < .05$.

Table 10

Mean estimates of β in Experiment 4 as a function of fluency, condition and time pressure (standard deviations in parentheses)

Condition	no time pressure		time pressure	
	high fluency	low fluency	high fluency	low fluency
classic	1.16 (0.42)	1.39 (0.89)	1.37 (0.40)	1.09 (0.27)
control	1.37 (0.50)	1.32 (0.47)	1.00 (0.21)	1.16 (0.23)
reversed	1.36 (0.87)	1.18 (0.40)	1.03 (0.19)	1.01 (0.23)

Note. β estimates greater than 1 indicate a tendency to respond negatively to a given item, estimates smaller than 1 indicate a tendency to respond affirmatively.

The interpretation of this three-way interaction is not as easy as in Experiment 3, but an inspection of the means reveals that in the no time pressure condition, participants in the control condition did not discriminate between high and low fluency items. Participants in the classic condition had a higher tendency to say “no” to low fluency items compared to high fluency items, whereas participants in the reversed condition showed a tendency to say “no” to high fluency items compared to low fluency items. This is the predicted pattern and replicates Experiment 3. Under time pressure, the tendencies are quite different. Almost all effects do vanish; that is, there is only one visible response tendency left: Participants in the classic condition had a tendency to respond negatively to high fluency items. Figure 6 shows the no time pressure condition only, without the control condition, corrected for the main effects. Obviously, the pattern is almost identical to Experiment 3.

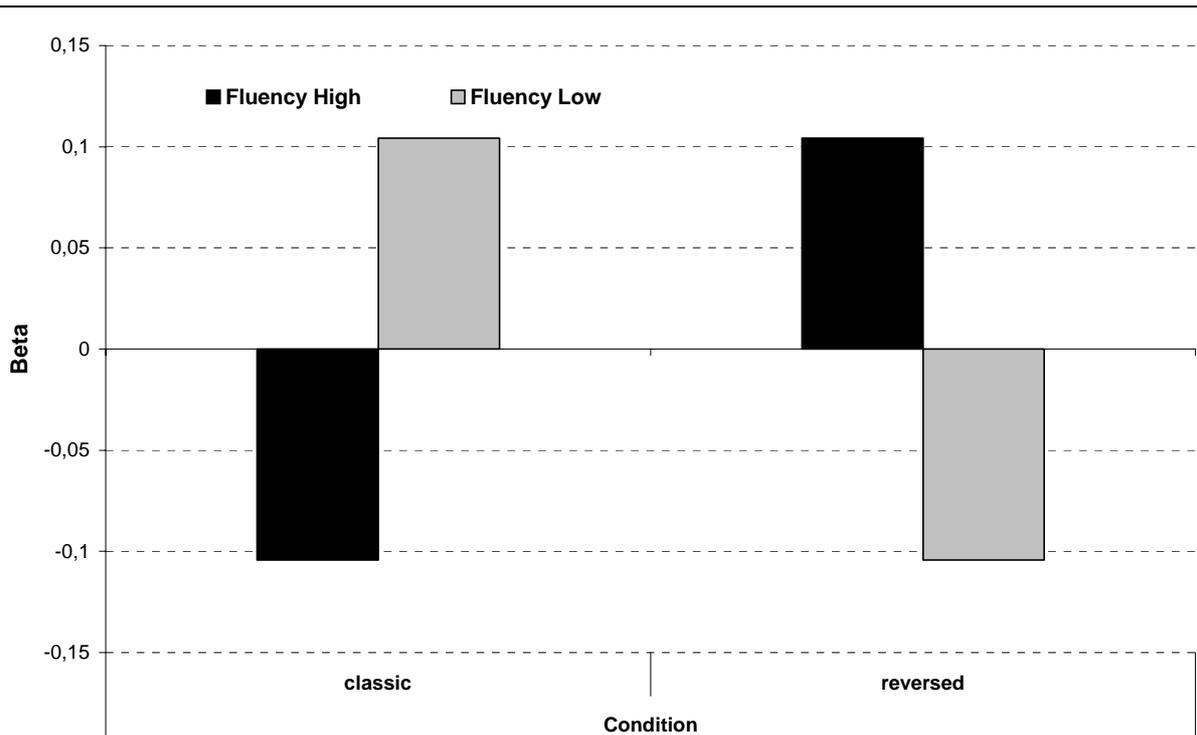


Figure 6. Estimates of β in Experiment 4 without time pressure as a function of fluency and condition (corrected for main effects). Negative values indicate a tendency to respond affirmatively, positive values indicate a tendency to respond negatively.

Discussion

The goal of Experiment 4 was to replicate Experiment 3 with an added control group and a time pressure manipulation. This manipulation rushed half of the participants to respond to a stimulus within a time frame of two seconds. It was included to show that the use of an experiential state is time consuming, that is, the experience needs interpretation, and cannot serve as direct input as affective feelings can be used directly in evaluative judgments. In addition, it was assumed that time pressure interrupts the newly learned associations in the mental rotation task.

For the discussion, let us first consider the response latencies. The mental rotation task was successful to associated affirmative “yes” responses and negative “no” responses with low and high fluency in the reversed condition, and vice versa in the classic condition. Data from the control condition were not available. All participants responded faster to shapes with a small angular difference than to shapes with a large angular difference, although dependent on the condition, affirmative or negative responses were required. The same is true for the crucial recognition task; participants responded generally faster to slightly tilted names (i.e. high fluency names) than to heavily tilted names (i.e. low fluency names). These are the two

important prerequisites for a meaningful interpretation of the results from the signal detection analysis.

The resulting SDT parameters from the recognition test are at the heart of any inferences from Experiment 4. First, it is noteworthy that the only significant influence on d' , a measure for participants' ability to discriminate old (present in the learning phase) from new (not present in the learning phase) names, was the time pressure manipulation. This is not astounding in terms of a speed-accuracy trade-off; yet, it is more important that the actual discrimination ability did not vary as a function of fluency (i.e., response latency), showing that there is no confound present.

Of greater interest are the results for people's response tendencies independent of actual discrimination ability. The main finding was that time pressure indeed interacted with participants' tendency to respond negatively or affirmatively to high or low fluency stimulus. Without time pressure, participants in the classic condition, who learned to associate short/easy mental rotation with an affirmative and long/difficult rotation with a negative response, showed a tendency to respond affirmatively to high fluency names and negatively to low fluency names. On the other hand, people in the reversed condition responded negatively to high fluency names, supposedly because they learned to associate short/easy mental rotation with a negative and long/difficult rotation with an affirmative response. This was the same pattern that was found in Experiment 3. Participants in the control condition did not show a differential bias for high or low fluency names; they showed an equal tendency to respond negatively. Ideally, those participants should have responded without any bias at all, yet they seem to have learned *something* in the mental rotation phase. So what did participants in the control group possibly learn? It is important to remember that participants in the control condition could not rely on the required length of mental rotation; they always had to try all possibilities. The frequent switches between a short or long required mental rotation and an affirmative or negative response might have turned the rotation task more difficult than in the classic and reversed condition. Thus, the task might have resulted in more negative feedback for these participants, thereby inducing a tendency to analyze a given stimulus thoroughly. But the difference between the mental rotation task and the recognition task was that the former always offered a definite solution in terms of "yes (no), the shapes are identical (different)", given enough scrutiny. On the other hand, the latter could not be solved definitely; thus, when in doubt, participants in the control condition might have responded cautiously: "no, this is a new name". This is surely a post-hoc interpretation, and the failure to record responses and latency data in the mental rotation task for the control group prohibits

testing these speculations. Nevertheless, the important point is that participants in the control group did not differentiate in their responses between high and low fluency names, whereas participants in the classic and reversed condition did.

Given the time frame of two seconds, the predicted response tendencies vanish. The only visible effect is the tendency of participants in the classic condition to respond negatively to high fluency names. This is rather unexpected, but the discussion of this effect will be referred to the general discussion. All in all, Experiment 4 successfully replicated Experiment 3, and showed in addition that the application of the learned association in the mental rotation task is dependent on enough available time at the point of judgment. This speaks again for a conscious use of an experiential state, although the interpretation was learned via ecological feedback in a given context, rather than explicitly provided.

General Discussion

The aim of this thesis was to show that experiential states resulting from cognitive processes need interpretation. These experiential states were termed cognitive feelings, because they provide information about otherwise inaccessible processes (Clare et al., 2001). The interpretation is learned via feedback from the environment and has therefore ecological validity in the context in which it is learned. The experiential information can then be used strategically in judgments and inferences.

This claim is similar to Schachter and Singer's (1962) notion that the large variety of affective feelings is only possible by cognitively interpreting an unspecific arousal according to cues available in the environment. A similar assumption is made here: Cognitive feelings are unexplained experiential states which need interpretation. The experiential states themselves are noticeable discrepancies in ongoing perceptual or conceptual cognitive processes and can be termed ease, fluency, surprise, familiarity or revelation. The interpretation is triggered by cues provided in the context of the experience, again, because the interpretation is presumably learned via ecological feedback in a given context.

Experiment 1 and 2 provided evidence that the impact of such experiential states indeed depends on the interpretation and that this interpretation can be explicitly provided by the cues that are present in a given context. Experiment 3 and 4 provided evidence for the proposed learning process, that is, the acquisition of the interpretation: If you are in the frequently mentioned oral examination, and you cannot retrieve the authors of a specific study the examiner asked for, but you have a clear experiential state from the question asked: You can recall the circumstances where you read the paper, you can even recall the content of the

study, but the authors remain inaccessible. You might label this experience a “tip-of-the-tongue” feeling (Hart, 1965; see Schwartz, 2002, for a review). Later, when you leave the examiner’s office, the names pop up from memory. This is the kind of reinforcement that teaches you that next time you have this cognitive feeling (i.e., feedback from your own cognitive processes) when a question is asked, you can label it a feeling of knowing, because you have learned that it signifies that you will be able to retrieve the item later from your memory. Yet, if the context suggests the wrong interpretations or the fluency is misleading as an index (Benjamin, Bjork, & Schwartz, 1998), you might falsely judge that you will be able to retrieve the item later.

Although the analogy to affective feelings is striking, it is important to note the major difference: Affective feelings have an inherent meaning. If you dislike an object, this dislike lends itself as immediate information for a judgment about the object which does not need interpretation (Bless, 2001). But if this object creates a discrepancy in cognitive processes, this discrepancy can be interpreted as increased belief in a statement (Hasher, Goldstein, & Toppino, 1977), increased or decreased length of stimulus presentation (Witherspoon & Allan, 1985), or reduced or increased loudness of background noise (Jacoby, Allan, Collins, & Larwill, 1988). Again, the interpretation itself is learned via ecological feedback in a given context; that is, the interpretation itself depends on the respective context, for example, when the interpretation is explicitly provided by the question an experimenter asks.

These assumptions result in a model that explains why very similar manipulations of fluency can result in very different effects. The experimental evidence presented is a first step in establishing the soundness of the model. However, there are a couple of open questions that warrant discussion.

A single dimension?

The assumption was made that all cognitive feelings (i.e., experiential states resulting from cognitive processes) map onto an easy–difficult or fluent–non-fluent dimension. Is such an assumption feasible? First, it might be sensible to include at least an intensity dimension, which is part of most dimensional models of affective feelings (e.g., Osgood, Suci, & Tannenbaum, 1957). Yet, because ease and difficulty are assumed to be valence-free to begin with, greater experienced ease or fluency is equivalent to an increase on the assumed one dimension. To say something is very easy or very fluent is equivalent to say that you experience fluency or ease intensively.

This assumption of a single dimension for the experience of cognitive processes is not at the heart of the discussion, and it might be difficult to proof empirically. But it is also

supported by the lack of representation in language. Whereas there are hundreds of adjective that describe affective feelings and emotional experiences (e.g., Averill, 1975), there are hardly any descriptive adjectives for cognitive experiences. Statements like “it is difficult to concentrate” or “it is easy to remember” make intuitively sense, but a statement like “a pleasant thought” automatically implies the content of the thought, and not the subjective experience of thinking. All descriptions of thinking as a process circle around the easy-difficult dimension; valence is only implied by the content. And while there are multi-dimensional models for affective experiences (e.g., very early, Wundt, 1910) or categorical models (e.g., Plutchik, 1980; see Ortony & Turner, 1990, for a review), it is difficult to come up with more than one dimension or category for cognitive feelings.

Is the distinction of attribution and interpretation necessary?

Some experimental evidence suggests that it is sufficient to attribute experienced ease to a given stimulus or task; for example, if the ease of generating six instances of self-assertive behavior is attributed to the task itself and not to the music in the background that supposedly facilitates autobiographic memory (Schwarz et al., 1991, Exp. 3). The interpretation seems to be inherent in this attribution, namely, if it is easy to come up with six instances, there must be a great number of such instances and one is a self-assertive person. The functional similarity of attribution and interpretation is explicitly made by Whittlesea and Williams (2001a, 2001b), who state that the feeling of familiarity is the result of attributing an incoherence in the flow of processing (i.e., facilitation/inhibition of processing) to a prior experience. More precise, if the incoherence is attributed to a source in the past, a feeling of familiarity is experienced (Whittlesea & Williams, 1998); that is, if a stimulus is unexpectedly easy to process, this discrepancy is attributed to the fact that one has encountered a stimulus before. The original formulation highlights the functional similarity of these terms even more: “Like the perception of any stimulus quality, the fluency of performance is interpreted within its context and attributed to some source. When this *unconscious attribution* [italics added] is to a source in the past, one experiences a feeling of familiarity, just as, when interpreting convergence of lines toward the horizon, one experiences a sense of depth.” (Whittlesea, 1993; p. 1235). Experiment 2 made explicitly use of the fact that attribution sometimes encompasses the interpretation: When the delay of information retrieval was attributed to the person retrieving the information, an interpretation in terms of frequency or typicality was the result.

But in general, attribution precedes the interpretation, and is therefore a necessary but not always a sufficient step; it is necessary that the cognitive feeling is attributed to a stimulus

or a task, but the way the feeling is interpreted is dependent on an available interpretation in that context (e.g., a memory experiment). In addition, the notion of unconscious attribution seems to be a misnomer; first, attribution is historically (Jones et al., 1987; H. Kelley, 1973) a conscious process; claiming that a conscious process is unconscious is somewhat inelegant. Second, the term attribution encompasses the interpretation of a discrepancy only in the very specific context of memory attributions. It does make sense to label an experienced discrepancy “familiarity” when the discrepancy is attributed to a previous encounter. And it does make sense that the ease of retrieval is attributed to the frequency or typicality of the retrieved information. In these cases, attribution encompasses the interpretation. But what does it mean if it is difficult to retrieve an item from memory? For example, if you are asked to give a judgment about how many good restaurants there are in the vicinity of Heidelberg, you may search your memory and you may have a hard time to come up with more than two. How can you use this difficulty of retrieval in answering the question; that is, forming a judgment? From this example, it is easy to make a distinction for attribution and interpretation. You need to attribute the experienced difficulty to the process of retrieval, not to the noisy surroundings or a bad hang-over from the extensive wine-tasting at exactly one of the good restaurants. Then you can interpret the difficulty that there are not many good restaurants in Heidelberg; that is, interpreting the non-ease as non-frequency or non-typicality, but you can also conclude that you are simply not an expert for restaurants. In this example, the attribution does not automatically encompass the interpretation.

Thus, in the terms of the present thesis, a more precise label would be “automatic interpretation”. By time and repeated application, a conscious and explicit interpretation of phenomenological information can be applied without effort and attention. Nevertheless, the default is a simple associative learning mechanism to begin with, in which no attention is directly involved. This is in line with a more general assumption about cognitive processes, that the acquisition should resemble the execution (Wilson, Lindsay, & Schooler, 2000). Yet, if necessary, the process is still open to conscious control; to use the example from Whittlesea (1993) again, much the same way as converging lines on a paper create an illusion of depth, we can break free from that illusion and avoid the experiential information, because we know that there is no depth involved. And as is easily possible to have people learn new or even re-learn perceptual cues for the experience of depth, which is a necessity for visual-motor coordination and adaptation (Cunningham & Welch, 1994; Field, Shipley, & Cunningham, 1999). Similarly, it is possible to re-learn the usage and interpretation of cognitive feelings, as done in Experiment 3 and 4.

Unconscious vs. Conscious Processes

So far, I have explicitly avoided an argument about explicit, conscious processes vs. implicit, unconscious processes; the topic was only implicitly present in the discussion. Nevertheless, dual-process models (models that claim two distinct classes of processes which can be distinguished along the lines of attention, automaticity, effort, and foremost, consciousness) have been very prominent in social psychology within the last years (see the edited volume by Chaiken & Trope, 1999; for a critique of these models, see Kruglanski & Thompson, 1999; Kruglanski, Thompson, & Siegel, 1999). The question that must be asked in this thesis is how conscious are the postulated processes; although already implicitly answered, the assumptions are made explicit here.

First, as frequently mentioned, cognitive feelings need to be consciously experienced to have an impact on judgments. Without a notable change in the flow of experience, there is no experience of a cognitive feeling and no impact on decisions, judgments and inferences. This is rather an axiom than a point of debate, because the logic actually works backwards: Phenomenal experience is at the heart of consciousness (Marcel, 1983). To be aware and conscious, we need to be conscious and aware of something. It is hard to argue that there is consciousness without phenomenal experience, and thus, we should assume that experience is necessarily conscious because it is one of the defining criteria for the state of consciousness itself (Brentano, 1874). And although much of the last twenty years research, especially in the area of subliminal priming, has shown numerous possible influences on judgments and decision without people's conscious awareness (for reviews, see DeCoster & Claypool, 2004; Klauer & Musch, 2003), the difference is that in most priming studies it is a prerequisite that people should not notice a prime or be consciously aware of a prime. Here, it is by definition the experience that causes the impact.

Second, the conceptual or perceptual processes that cause the experience are not directly accessible. They are not unconscious, because we receive conscious feedback via the experience, but this feedback is indirect. There is, however, no direct sensory organ that informs us about the status of our cognitive processes, whereas we do get direct feedback from many processes in our body (hunger, thirst, pain, etc.). To assume that feelings serve as a monitoring and feedback system for cognitive processes also solves a conceptual problem that is inbuilt in the computer metaphor of cognition (Nelson, 1996): *Quis custodit custodes?* A process cannot be controlled by another same-level process; if the controlling process detects an error, there are always two possibilities: There is an actual error or the control process itself is faulty. Thus, another control process would be necessary that controls the

control process, requiring for itself yet another process, leading to an infinite regress. Cognitive processes require therefore a monitoring and control system on a different level. A feeling-based account offers this kind of different level. The supposition is that cognitive processes are themselves inaccessible, but we get experiential feedback, a “feeling”, of how they execute. These cognitive feelings are not subject to further monitoring or control, because feelings have an immediate and inherent truth value. To construe cognitive feelings as indirect feedback from cognitive processes allows cognitive feelings to circumvent the infinite regress problem that is inherent in metacognition (Nelson, 1996).

Third, the interpretation of the unspecific experience of ease or fluency is assumed to be by default automatic, but is open to controlled and deliberative use. The first part of this assumption derives from the hypothesis that the interpretation is learned via ecological feedback. It would be counterintuitive to assume that a mechanism that is simple and effortless during acquisition should require effort and deliberation during execution (Wilson, Lindsay, & Schooler, 2000). The second part and additional of this assumption derives from the experimental results of misattribution (e.g., Wänke, Schwarz, & Bless, 1995), the strategic use of ease and fluency (Experiment 1 and 2 of the present thesis) and the openness to explicit instructions (Briñol, Petty, & Tormala, 2004). Again, the analogy of perception is helpful: Even in a two-dimensional picture with converging lines you automatically experience depth, without interpreting the given cues deliberately. But with some effort, you can discard the experiential information and see nothing but two lines. Or you could even change the interpretation of given perceptual cues (e.g., discarding or reversing the interpretation); this might be effortful, controlled and conscious in the beginning, but with time the adaptive new interpretation becomes effortless and automatic (Field, Shipley, & Cunningham, 1999). Much the same way, the interpretation of a fluency cue can also adapt given ecological feedback. But it is important to note that in a real world setting misattribution and explicit change of the interpretation is the exception rather than the rule.

Finally, Jacoby and colleagues also discuss the impact of prior presentation on fame judgments in terms of “unconscious influences of the past”. They even provide an elaborate process-dissociation model on how to determine what part of a judgment process is automatic and what part is controlled (Kelley & Jacoby, 1998, 2000): In a simple memory task, there are two instructions, one constituting an exclusion condition and one constituting an inclusion condition. The exclusion instruction is to call an item “old” only when you recognize it without ambiguity (a “remember” judgment). The inclusion instruction is to call an item “old” when you remember it and/or if it feels familiar (a “know” judgment). The difference between

the inclusion and exclusion instruction is supposed to be the automatic and non-deliberative part of the recognition process. In the light of the present argument, the process is neither automatic nor unconscious. The only point is that the influence of prior exposure that leads to a cognitive feeling of familiarity cannot be consciously recognized. The use of the resulting experiential state is strategic and by no means automatic. The design of the suggested process dissociation framework itself shows that participants need to deliberately decide whether they use their cognitive feelings (e.g., resulting from prior exposure) or not (i.e., make only “remember” judgments and no feeling-based “know” judgments).

Another point that is closely related to the question of conscious control is the apparent speed of responses in Experiment 3 and 4, and many other experiments that used recognition tasks to investigate ease and fluency. Many models in psychology imply that there is an absolute threshold in terms of execution time before conscious control is possible (e.g., Zajonc, 1980; Greenwald & Banaji, 1995). If you want to investigate automatic or implicit processes, because of this rationale, time pressure and fast responses are the means of choice. How does this rationale fit with the interpretation hypothesis? First, it is simple to argue that the time frame of about two seconds is way over the assumed threshold for automatic or implicit processes. Second, there is the functional similarity of acquisition and execution: In Experiment 1 and 2, no time pressure was present in forming a judgment, and in Experiment 3 and 4, the acquisition of a new interpretation was effortless and therefore, the execution should also be effortless. But even without a presumed fast and effortless interpretation, there is a simple explanation for the apparent speed of the given responses. Interpretation does not need to take place every time in 80 identical trials, but only in the first couple of trials. Then the same schema is applicable to all following trials. This is corroborated by the configuration of the response latencies across trials in Experiment 3 and the conditions without time pressure in Experiment 4. They show much higher latencies for trials 1 to 8, followed by a drop to a constant response speed level for the rest of the trials. Although this could reflect a mere increase in acquaintance with the task, this familiarization with the task possibly includes the proposed application of a learned interpretation. Thus, it is even possible to assume a highly deliberative interpretation that is executed rapidly in the following.

Latency as an index for fluency

The validity of operationalization and measurement lies at the heart of experimental psychology; experimental psychologists impose complex meaning to very primitive and simple operationalizations and measurements. When we tell people that they failed on a test, when we put them under time pressure or when we hold them accountable for given decision,

we must ensure that we indeed manipulate self-esteem (and not mood), that we indeed manipulate cognitive resources (and not stress) and that we indeed manipulate involvement (and not caution). Whether we use simple ratings on Likert scales to measure our effects, the time an infant looks at a new stimulus or the alignment of paramagnetic molecules in the brain induced by strong magnetic fields, the meaning of these measures is always in the hands of the experimenter. And it is the responsibility of the experimenter to show that the used operationalizations and measures are valid. We must show that a mark on a rating scale indeed measures an attitude, that the longer time a toddler looks at a face than at an abstract painting indeed measures preference for human features, or that a higher concentration of oxygen-saturated blood in a brain area indeed measures brain activity.

In the present approach, latencies were used as an index for ease and fluency. Experiment 1 and 2 used latencies of information retrieval as an independent variable and Experiment 3 and 4 used latencies as a manipulation check for the high or low fluency of processing of given stimuli and exploited the effect that the time necessary for mental rotation is a linear function of the rotation angle.

Experiment 1 tried to show that the manipulation of retrieval latency is equivalent to an experienced feeling by introducing a “feeling” vs. “thinking” distinction in the assessment of the dependent variables. This manipulation was successful, but the validity is not guaranteed. At this point, I do not see a good manipulation check for the methodology used in Experiment 1 and 2.

What was labeled “high fluency” and “low fluency” in Experiment 3 and 4 corresponded to fast and slow response latencies, but the assumption remained untested that latency is an index for fluency. Yet, as already mentioned, it does make intuitively sense to assume that any process that is experienced as fluent or easy should be faster than a process that is non-fluent or difficult. There is a lack of research tools and methodology that allows investigating the experience of thinking (i.e. of ongoing cognitive processes) more precisely. We can rely on the construct of subjective experience to predict and explain experimental results, but we are not able to measure this construct directly. Experiment 1 and 2 used an approach that tried to externalize the proposed construct; but such a strategy is widely open to the critique that there is at best a conceptual similarity of this externalized approach and a true, subjective experience.

The conclusion is that one has to rely on the face validity of latency as an index and concur with Whittlesea and Williams (1998, pp.163) in their conclusion: “We have used response latencies to index the psychological construct of fluency in these studies because

they are probably linearly related, and because we cannot yet imagine a useful index of the qualitative goodness of processing a stimulus. However, we do not want to claim that our measurements of response latency directly correspond to the actual, effective variable to which our subjects responded. We, [...], will have to work on precisising indices of the qualitative aspects of psychological experience.”

The possible inherent meaning of cognitive feelings

The question that motivated the present thesis was whether the meaning of cognitive feelings is ontogenetically or phylogenetically acquired. The proposition is that cognitive feelings are unspecific and have no inherent, phylogenetically acquired meaning. But there is evidence against this claim. First, Reber, Winkielman, & Schwarz (1998) showed in three experiments that more fluent processing of stimuli leads to higher preference for these stimuli. They manipulated fluency by priming the contour of a stimulus (Exp. 1), the contrast of black and white circles (Exp. 2) and the duration of presentation (Exp. 3). More importantly, they tried to show that the effect of fluency is specifically positive; in Experiment 2, they asked not only for ratings of prettiness, but of ugliness, and in Experiment 3, they asked for disliking as well as for liking. In both cases they found a decrease in rated dislike and ugliness with greater fluency. Although the effects were relatively weak, they represent evidence that fluency might be inherently of positive valence.

Besides these effects on self-reported preferences, Winkielman and Cacioppo (2001) used EMG to measure affective responses towards presented stimuli. They manipulated processing fluency of the stimuli by contour priming (Exp. 1) and presentation duration (Exp. 2). In both studies they found a greater activation of the zygomaticus region, which is indicative of experienced positive affect. They explain this effect within a hedonic fluency model (Winkielman, Schwarz, Fazendeiro, & Reber, 2002), that assumes a direct link of processing fluency and positive affective feelings. Furthermore, Reber, Schwarz, and Winkielman (in press) turn the argument around and try to show that aesthetic pleasure derives from greater processing fluency of aesthetically pleasant items compared to aesthetically unpleasant items. These findings clearly contradict the non-specificity assumption of experienced ease or fluency.

But even in the data presented in this thesis are indications of inherent meaning. In Experiment 4, participants in the time pressure condition showed a high tendency to respond negatively to high fluency items. This is hard to explain if one assumes non-specificity. An explanation is that participants' experience surprise about the high fluency and because surprise is unpleasant, they respond negatively (Whittlesea, personal communication).

These results are hard to ignore, especially the EMG data from Winkielman and Cacioppo (2001). But what do they imply for the presented model that the impact of cognitive feelings depends on interpretation and that this interpretation is learned via ecological feedback? First, it is no problem to accept that the proposed single dimension of ease and difficulty corresponds to an affective dimension of pleasantness and unpleasantness. It makes even intuitively sense that processes that are more fluent/easier are also more pleasant, whereas difficulty is unpleasant. This is indeed the pattern of results that was obtained in Experiment 1; immediate information presentation led to a more positive evaluation of the targets in general. But this does by no means preclude the possibility of interpretation. It only adds the point that if no interpretation is available, fluency will be linearly related to any dimension that corresponds to hedonic positivity. In terms of a dimensional model of cognitive feelings, a positive-negative dimension, if at all present, is most likely parallel to the proposed easy-difficult dimension.

On an experimental level, it might be very difficult to distinguish learned interpretations from true, inherent meaning. The debate is similar to the question whether affect precedes cognition or vice versa (Zajonc, 1980, 1984; Lazarus, 1981, 1984). It is only possible to make a common sense argument: The malleability of cognitive feelings has been demonstrated, however, affective feelings are not prone to such interpretation and relearning. This in turn speaks against an inherent meaning of cognitive feelings.

But the main point remains: Given there was ecological feedback for an experiential state, fluency will be interpreted following this feedback. Just imagine a classical conditioning situation where a participant is presented with two buttons and these buttons are marked with high or low contrast figures. Following the logic from above, in the very first trial, the high contrast button is perceived more fluently and therefore, it is evaluated more positively or aesthetically pleasing, and therefore, has a higher possibility to be chosen. But the high contrast button results in a very annoying noise, whereas the low contrast button results in a mild tone. No matter what inherent pleasantness the high contrast button possessed, it will be washed away by actual ecological feedback. It sounds trivial, but this model lies at the bottom of the present model of interpretations learned from ecological feedback.

Forthcoming research questions

Having discussed some of the open points, the question is, what does the proposed model offer for further research? There are many open threads to follow within the presented paradigms. First, when do people actually use their cognitive feelings in judgments and inferences? The presented research showed and discussed some of the boundary conditions.

For example, a cognitive feeling must be seen as diagnostic and the judgment must depend on the available information at the time of judgment. But there are a number of other possible causes that might limit or boost the influence of experiential states in judgments; motivational states like the accountability for decisions, personality traits like “need for closure”, or specifics of the situation like time pressure are prominent candidates.

Second, the algorithm for the learning process needs specification. In the present research, the number of learning trials was rather high compared to what is feasible in real-world settings. The open questions are clear: How sensitive is the mechanism? Is positive or negative feedback more efficient in the learning process? How broad is the generalization of the learned interpretations; are they restricted to a specific context or applicable in all situations? Determining the parameters of the learning process seems to be an important step in further establishing the model.

And third, what is the contribution of the proposed “feeling” component in a decision? Are judgments feeling or content based in an all- or non-fashion? This is an area where a research tool already exists to investigate the contribution in measurable units: The previously mentioned process-dissociation procedure by Kelley and Jacoby (1998, 2000). This tool should be effective to quantifying the impact of cognitive feelings. To reiterate, there is an inclusion (recognition and familiarity lead to “old” responses) and an exclusion (only true recognition leads to “old” responses) condition. If the proposed interpretation model is sound, the re-learning of interpretations should only influence “old” responses which are not dependent on true recognition. That is, any effect of learning should appear in an inclusion condition, and the size of the effect should equal the difference in the amount of “old” responses between the inclusion and exclusion condition.

These are open questions that reside within the present research, but the idea of interpreting experiential feedback from higher mental processes has implications for seemingly unrelated research areas. For example, there has been a debate about intuition in the area of judgments and decision making (Hogarth, 2001). The present model offers an intriguing explanation why intuitive judgments are sometimes better than deliberative judgments. If a person is an expert in a situation, this person has often received feedback from her decisions. Thus, she can rely on what “feels” right, presumably because this feeling results from a cue pattern in that situation that has been encountered before. And actions taken in that situation lead to positive or negative reinforcement, much the same way as the use of cognitive fluency leads to decisions about the memory status of an item. Such feeling-based decisions are possibly better because deliberative considerations necessarily narrow the use of

present information, whereas the feeling is based on the whole cue configuration of the context. In this vein, the present model can make predictions when to rely on your feeling and when not.

Finally, the assumption of learned interpretations offers a contribution to the field of metacognition, by deriving new hypotheses which are directly testable. For example, if judgments of learning (JOL; i.e., the judgment that you will be able to remember an item) are dependent on the proposed feelings resulting from cognitive processes, it should be possible to delineate the cues people use for JOL and teach a more effective application of these judgments. For instance, that you will be able to remember items better when encoding is difficult rather than easy.

All together, the idea of learned interpretations does not contradict existing research, but rather offers a more refined view. Following some of the presented threads in further research will either show that the model is accurate in its assumptions and predictions, or the model might not stand the test of time; the evidence in the present thesis is merely a first step in establishing the validity of the interpretation approach to cognitive feelings.

Conclusion

This thesis tried to show that cognitive feelings (i.e. experienced discrepancies in the ongoing flow of cognitive processes) do not have an inherent, phylogenetically acquired meaning, like affective feelings do. Rather, the hypothesis was presented that cognitive feelings acquire their meaning ontogenetically via ecological feedback in a given context, resulting in the interpretation of an unspecific ease or fluency feeling as a specific and meaningful feeling in this context. Yet, although it has been demonstrated that it is possible to re-learn the interpretation of cognitive feelings (Experiment 3 and 4) and that results of classic paradigms depend on the presence of an applicable interpretation (Experiment 1 and 2), this does not necessarily preclude the possibility of an inherent meaning that is not dependent on interpretation.

But the proposed model of learned interpretations does not only explain and incorporate many experimental findings in the area of subjective experiences. It offers a bridge between the research on ease of retrieval effects which concerned mostly social judgments and the research on fluency effects in the memory domain. Furthermore, it fits into the broader field of social metacognition by offering an explaining of how we draw inferences about and from our own ongoing cognitive processes.

Nevertheless, the presented experiments are only a first step and many research questions remain open; for example, what rules determine the proposed associative learning mechanism? Is awareness (i.e. consciousness) a necessary prerequisite or just a by-product? How far do the learned interpretations from one context generalize and what are the specific cues that trigger an interpretation? These questions all need to be answered in due time.

But even with only this first step, I do believe that the presented ideas and experiments offer a contribution to a social psychology of subjective experiences. They show that the investigation of the experience of thinking is worthwhile, as much as the content of thought is.

References

- Averill, J. R. (1975). A semantic atlas of emotional concepts. *Catalog of Selected Documents in Psychology*, 5, 30.
- Bargh, J. A., & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist*, 54(7), 462-479.
- Bargh, J. A., & Ferguson, M. J. (2000). Beyond behaviorism: On the automaticity of higher mental processes. *Psychological Bulletin*, 126(6), 925-945.
- Begg, I. M., Anas, A., & Farinacci, S. (1992). Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General*, 121(4), 446-458.
- Benjamin, A. S., Bjork, R. A., & Hirshman, E. (1998). Predicting the future and reconstructing the past: A Bayesian characterization of the utility of subjective fluency. *Acta Psychologica*, 98(2), 267-290.
- Benjamin, A. S., Bjork, R. A., & Schwartz, B. L. (1998). The mismeasure of memory: When retrieval fluency is misleading as a metamnemonic index. *Journal of Experimental Psychology: General*, 127(1), 55-68.
- Bernstein, D. M., Whittlesea, B. W. A., & Loftus, E. F. (2002). Increasing confidence in remote autobiographical memory and general knowledge: Extensions of the revelation effect. *Memory & Cognition*, 30(3), 432-438.
- Bless, H. (2001). Mood and the use of general knowledge structures. In L. L. Martin & G. L. Clore (Eds.), *Theories of mood and cognition: A user's guidebook*. (pp. 9-26): Lawrence Erlbaum Associates, Publishers.
- Bless, H., & Forgas, J. P. (2000). The message within: Toward a social psychology of subjective experiences. In H. Bless & J. P. Forgas (Eds.), *Message within: The role of subjective experience in social cognition and behavior*. (pp. 372-392): Psychology Press.
- Bornstein, B. H., & Neely, C. B. (2001). The revelation effect in frequency judgment. *Memory & Cognition*, 29(2), 209-213.
- Bower, G. H. (1981). Mood and memory. *American Psychologist*, 36(2), 129-148.

- Brentano, F. C. (1874). *Psychologie vom empirischen Standpunkt*. Deutschland, Leipzig: Duncker & Humblot.
- Briñol, P., Petty, R., & Tormala, Z. (2004). *The Meaning of Subjective Ease and its Malleability*. Paper presented at the Heidelberg Small-Group Meeting on Social Metacognition, Heidelberg, Germany.
- Brunswik, E. (1941). Perceptual size-constancy in life situations. *Psychological Bulletin*, *38*, 611-612.
- Brunswik, E. (1943). Organismic achievement and environmental probability. *Psychological Review*, *50*, 255-272.
- Cantor, J. R., Zillmann, D., & Bryant, J. (1975). Enhancement of experienced sexual arousal in response to erotic stimuli through misattribution of unrelated residual excitation. *Journal of Personality & Social Psychology*, *32*(1), 69-75.
- Chaiken, S., & Trope, Y. (1999). *Dual-process theories in social psychology*.: Guilford Press.
- Clore, G. L. (1992). Cognitive phenomenology: Feelings and the construction of judgment. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments*. (pp. 133-163). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.
- Clore, G. L., Wyer, R. S. J., Dienes, B., Gasper, K., Gohm, C., & Isbell, L. (2001). Affective feelings as feedback: Some cognitive consequences. In L. L. Martin & G. L. Clore (Eds.), *Theories of mood and cognition: A user's guidebook*. (pp. 27-62): Lawrence Erlbaum Associates, Publishers.
- Cohen, J. (1990). Things I have learned (so far). *American Psychologist*, *45*(12), 1304-1312.
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, *112*(1), 155-159.
- Cunningham, H. A., & Welch, R. B. (1994). Multiple concurrent visual-motor mapping: implications for models of adaptation. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 987.
- DeCoster, J., & Claypool, H. M. (2004). A Meta-Analysis of Priming Effects on Impression Formation Supporting a General Model of Informational Biases. *Personality & Social Psychology Review*, *8*(1), 2-27.
- Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review*, *107*(4), 659-676.

- Fiedler, K. (2000). Toward an integrative account of affect and cognition phenomena using the BIAS computer algorithm. In J. P. Forgas (Ed.), *Feeling and thinking: The role of affect in social cognition*. (pp. 223-252): Cambridge University Press.
- Fiedler, K., Brinkmann, B., Betsch, T., & Wild, B. (2000). A sampling approach to biases in conditional probability judgments: Beyond base rate neglect and statistical format. *Journal of Experimental Psychology: General*, *129*(3), 399-418.
- Field, D. P., Shipley, T. F., & Cunningham, D. W. (1999). Prism adaptation to dynamic events. *Perception & Psychophysics*, *61*(1), 161-176.
- Fisher, R. A. (1934). *Statistical Methods for Research Workers*. 5th ed.: Oliver and Boyd: Edinburgh.
- Forgas, J. P. (1995). Mood and judgment: The affect infusion model (AIM). *Psychological Bulletin*, *117*(1), 39-66.
- Forgas, J. P. (2000). Affect and information processing strategies: An interactive relationship. In J. P. Forgas (Ed.), *Feeling and thinking: The role of affect in social cognition*. (pp. 253-280): Cambridge University Press.
- Frank, M. G., & Gilovich, T. (1988). The dark side of self- and social perception: Black uniforms and aggression in professional sports. *Journal of Personality & Social Psychology*, *54*(1), 74-85.
- Gardiner, J. M. & Richardson-Klavehn, A. (2000). Remembering and knowing. In: E. Tulving & F. I. M. Craik (Eds.), *The Oxford handbook of memory* (pp. 229-244). London: Oxford University Press.
- Goldstein, E. B. (2002). *Sensation and Perception*. Pacific Grove, CA, USA: Wadsworth.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, *102*(1), 4-27.
- Hansen, J. & Wänke, M. (2004). *The role of expectation in the ease of retrieval effect*. Paper presented at the Heidelberg Small-Group Meeting on Social Metacognition, Heidelberg, Germany.
- Hart, J. T. (1965). Memory and the feeling-of-knowing experience. *Journal of Educational Psychology*, *56*(4), 208-216.
- Hasher, L., Goldstein, D., & Toppino, T. (1977). Frequency and the conference of referential validity. *Journal of Verbal Learning & Verbal Behavior*, *16*(1), 107-112.

- Hogarth, R. M. (2001). *Educating intuition.*: University of Chicago Press.
- Jacoby, L. L., & Witherspoon, D. (1982). Remembering without awareness. *Canadian Journal of Psychology*, 36(2), 300-324.
- Jacoby, L. L., Allan, L. G., Collins, J. C., & Larwill, L. K. (1988). Memory influences subjective experience: Noise judgments. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 14(2), 240-247.
- Jacoby, L. L., Kelley, C., Brown, J., & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality & Social Psychology*, 56(3), 326-338.
- Jacoby, L. L., Woloshyn, V., & Kelley, C. (1989). Becoming famous without being recognized: Unconscious influences of memory produced by dividing attention. *Journal of Experimental Psychology: General*, 118(2), 115-125.
- Johnston, W. A., Dark, V. J., & Jacoby, L. L. (1985). Perceptual fluency and recognition judgments. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 11(1), 3-11.
- Johnston, W. A., Hawley, K. J., & Elliott, J. M. (1991). Contribution of perceptual fluency to recognition judgments. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 17(2), 210-223.
- Jones, E. E., Kanouse, D. E., Kelley, H. H., Nisbett, R. E., Valins, S., & Weiner, B. (1987). *Attribution: Perceiving the causes of behavior.*: Lawrence Erlbaum Associates, Inc.
- Kelley, C. M., & Jacoby, L. L. (1998). Subjective reports and process dissociation: Fluency, knowing, and feeling. *Acta Psychologica*, 98(2), 127-140.
- Kelley, C. M., & Jacoby, L. L. (2000). Recollection and familiarity: Process-dissociation. In E. Tulving & F. I. M. Craik (Eds.), *Oxford handbook of memory.* (pp. 215-228): Oxford University Press.
- Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist*, Vol. 28(2), 107-128.
- Kitayama, S., Markus, H. R., & Kurokawa, M. (2000). Culture, emotion, and well-being: Good feelings in Japan and the United States. *Cognition & Emotion*, 14(1), 93-124.

- Klauer, K. C., & Musch, J. (2003). Affective priming: Findings and theories. In J. Musch & K. C. Klauer (Eds.), *Psychology of evaluation: Affective processes in cognition and emotion*. (pp. 7-49): Lawrence Erlbaum Associates, Publishers.
- Koriat, A. (1993). How do we know that we know? The accessibility model of the feeling of knowing. *Psychological Review*, *100*(4), 609-639.
- Koriat, A., & Goldsmith, M. (1996). Monitoring and control processes in the strategic regulation of memory accuracy. *Psychological Review*, *103*(3), 490-517.
- Koriat, A., & Levy-Sadot, R. (1999). Processes underlying metacognitive judgments: Information-based and experience-based monitoring of one's own knowledge. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology*. (pp. 483-502): Guilford Press.
- Koriat, A., Goldsmith, M., & Pansky, A. (2000). Toward a psychology of memory accuracy. *Annual Review of Psychology*, *51*, 481-537.
- Kruglanski, A. W., & Thompson, E. P. (1999). Persuasion by a single route: A view from the unimodel. *Psychological Inquiry*, *10*(2), 83-109.
- Kruglanski, A. W., Thompson, E. P., & Spiegel, S. (1999). Separate or equal? Bimodal notions of persuasion and a single-process "unimodel". In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology*. (pp. 293-313): Guilford Press.
- Kunst-Wilson, W. R., & Zajonc, R. B. (1980). Affective discrimination of stimuli that cannot be recognized. *Science*, *207*(4430), 557-558.
- Lazarus, R. S. (1981). A cognitivist's reply to Zajonc on emotion and cognition. *American Psychologist*, *36*, 222-223.
- Lazarus, R. S. (1984). On the primacy of cognition. *American Psychologist*, *39*, 124-129.
- Mandler, G. (1980). Recognizing: The judgment of previous occurrence. *Psychological Review*, *87*(3), 252-271.
- Mandler, G., Nakamura, Y., & Van Zandt, B. J. (1987). Nonspecific effects of exposure on stimuli that cannot be recognized. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *13*(4), 646-648.
- Marcel, A. J. (1983). Conscious and unconscious perception: An approach to the relations between phenomenal experience and perceptual processes. *Cognitive Psychology*, *15*(2), 238-300.

- Markus, H. R., & Kitayama, S. (1991). Culture and the self: Implications for cognition, emotion, and motivation. *Psychological Review*, 98(2), 224-253.
- McClelland, G. H. (2000). Nasty Data: Unruly, ill-mannered observations can ruin your analysis. In: H. T. Reis & C. M. Judd (Eds.), *Handbook of research methods in social and personality psychology* (pp.393 – 411): Cambridge University Press.
- Mellers, B., Schwartz, A., & Ritov, I. (1999). Emotion-based choice. *Journal of Experimental Psychology: General*, 128(3), 332-345.
- Neely, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory & Cognition*, 4(5), 648-654.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, 106(3), 226-254.
- Nelson, T. O. (1996). Consciousness and metacognition. *American Psychologist*, 51(2), 102-116.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84(3), 231-259.
- Ortony, A., & Turner, T. J. (1990). What's basic about basic emotions? *Psychological Review*, 97(3), 315-331.
- Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The measurement of meaning*. Urbana, IL: University of Illinois Press.
- Plutchik, R. (1980). *Emotion – A psychoevolutionary synthesis*. New York: Harper& Row.
- Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness & Cognition: An International Journal*, 8(3), 338-342.
- Reber, R., Schwarz, N., & Winkielman, P. (in press). Processing Fluency and Aesthetic Pleasure: Is Beauty in the Perceiver's Processing Experience? *Review of Personality and Social Psychology*.
- Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science*, 9(1), 45-48.
- Reber, Wurtz, & Zimmermann (2004). Exploring “fringe” consciousness: The subjective experience of perceptual fluency and its objective bases. *Consciousness and Cognition*, 13, 47-60.

- Reisenzein, R. (2000). The subjective experience of surprise. In H. Bless & J. P. Forgas (Eds.), *The message within: The role of subjective experience in social cognition and behavior*. (pp. 262-279): Psychology Press.
- Rosenthal, R., & Rosnow, R. (1991). *Essentials of behavioral research: Methods and data analysis* (2 ed.). New York: McGraw-Hill.
- Rosnow, R. L., & Rosenthal, R. (1995). "Some things you learn aren't so": Cohen's paradox, Asch's paradigm, and the interpretation of interaction. *Psychological Science*, 6(1), 3-9.
- Rosnow, R. L., & Rosenthal, R. (1996). Contrasts and interactions redux: Five easy pieces. *Psychological Science*, 7(4), 253-257.
- Schachter, S., & Singer, J. (1962). Cognitive, social, and physiological determinants of emotional state. *Psychological Review*, 69(5), 379-399.
- Schwartz, B. L. (2002). *Tip-of-the-tongue states: Phenomenology, mechanism, and lexical retrieval*.: Lawrence Erlbaum Associates, Publishers.
- Schwartz, B. L., Benjamin, A. S., & Bjork, R. A. (1997). The inferential and experiential bases of metamemory. *Current Directions in Psychological Science*, 6(5), 132-137.
- Schwartz, B. L., Travis, D. M., Castro, A. M., & Smith, S. M. (2000). The phenomenology of real and illusory tip-of-the-tongue states. *Memory & Cognition*, 28(1), 18-27.
- Schwarz, N. (1998). Accessible content and accessibility experiences: The interplay of declarative and experiential information in judgment. *Personality & Social Psychology Review*, 2(2), 87-99.
- Schwarz, N. (in press). Meta-cognitive Experiences in Consumer Judgment and Decision Making. *Journal of Consumer Research*.
- Schwarz, N., & Clore, G. L. (1983). Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality & Social Psychology*, 45(3), 513-523.
- Schwarz, N., & Clore, G. L. (2003). Mood as Information: 20 Years Later. *Psychological Inquiry*, 14(3), 296-303.
- Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality & Social Psychology*, 61(2), 195-202.

- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, *Vol. 171*(3972), 701-703.
- Skurnik, I. W. (1998). *Metacognition and the illusion of truth*. Unpublished doctoral dissertation, Princeton University.
- Skurnik, I., Schwarz, N., & Winkielman, P. (2000). Drawing inferences from feelings: The role of naive beliefs. In H. Bless & J. P. Forgas (Eds.), *The Message within: The role of subjective experience in social cognition and behavior*. (pp. 162-175): Psychology Press.
- Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments & Computers*, *31*(1), 137-149.
- Strack, F. (1992). The different routes to social judgments: Experiential versus informational strategies. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments*. (pp. 249-275). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.
- Swets, J. A., Dawes, R. M., & Monahan, J. (2000). Psychological science can improve diagnostic decisions. *Psychological Science in the Public Interest*, *1*(1), 1-26.
- Tolkien, J. R. R. (1954/1955). *The Lord of the Rings*. Great Britain: George Allen & Unwin.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, *Vol. 5*(2), 207-232.
- Verplanken, B., Hofstee, G., & Janssen, H. J. W. (1998). Accessibility of affective versus cognitive components of attitudes. *European Journal of Social Psychology*, *28*(1), 23-35.
- Waenke, M., Bless, H., & Biller, B. (1996). Subjective experience versus content of information in the construction of attitude judgments. *Personality & Social Psychology Bulletin*, *22*(11), 1105-1113.
- Waenke, M., Bohner, G., & Jurkowitsch, A. (1997). There are many reasons to drive a BMW: Does imagined ease of argument generation influence attitudes? *Journal of Consumer Research*, *24*(2), 170-177.
- Waenke, M., Schwarz, N., & Bless, H. (1995). The availability heuristic revisited: Experienced ease of retrieval in mundane frequency estimates. *Acta Psychologica*, *89*(1), 83-90.

- Watkins, M. J., & Peynircioglu, Z. F. (1990). The revelation effect: When disguising test items induces recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 1012.
- Whittlesea, B. W. A. (1993). Illusions of familiarity. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *19*(6), 1235-1253.
- Whittlesea, B. W. A. (2002). False memory and the discrepancy-attribution hypothesis: The prototype-familiarity illusion. *Journal of Experimental Psychology: General*, *131*(1), 96-115.
- Whittlesea, B. W. A., & Leboe, J. P. (2000). The heuristic basis of remembering and classification: Fluency, generation, and resemblance. *Journal of Experimental Psychology: General*, *129*(1), 84-106.
- Whittlesea, B. W. A., & Price, J. R. (2001). Implicit/explicit memory versus analytic/nonanalytic processing: Rethinking the mere exposure effect. *Memory & Cognition*, *29*(2), 234-246.
- Whittlesea, B. W. A., & Williams, L. D. (1998). Why do strangers feel familiar, but friends don't? A discrepancy-attribution account of feelings of familiarity. *Acta Psychologica*, *98*(2-3), 141-165.
- Whittlesea, B. W. A., & Williams, L. D. (2001a). The discrepancy-attribution hypothesis: I. The heuristic basis of feelings and familiarity. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *27*(1), 3-13.
- Whittlesea, B. W. A., & Williams, L. D. (2001b). The discrepancy-attribution hypothesis: II. Expectation, uncertainty, surprise, and feelings of familiarity. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *27*(1), 14-33.
- Whittlesea, B. W., Jacoby, L. L., & Girard, K. (1990). Illusions of immediate memory: Evidence of an attributional basis for feelings of familiarity and perceptual quality. *Journal of Memory & Language*, *29*(6), 716-732.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, *107*(1), 101-126.
- Winkielman, P., & Cacioppo, J. T. (2001). Mind at ease puts a smile on the face: Psychophysiological evidence that processing facilitation elicits positive affect. *Journal of Personality & Social Psychology*, *81*(6), 989-1000.

- Winkielman, P., & Schwarz, N. (2001). How pleasant was your childhood? Beliefs about memory shape inferences from experienced difficulty of recall. *Psychological Science, 12*(2), 176-179.
- Winkielman, P., Schwarz, N., Fazendeiro, T. A., & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. In J. Musch & K. C. Klauer (Eds.), *Psychology of evaluation: Affective processes in cognition and emotion*. (pp. 189-217): Lawrence Erlbaum Associates, Publishers.
- Witherspoon, D., & Allan, L. G. (1985). The effect of a prior presentation on temporal judgments in a perceptual identification task. *Memory & Cognition, 13*(2), 101-111.
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist, 35*(2), 151-175.
- Zajonc, R. B. (1984). On the primacy of affect. *American Psychologist, 39*(2), 117-123.
- Zanna, M. P., & Cooper, J. (1974). Dissonance and the pill: An attribution approach to studying the arousal properties of dissonance. *Journal of Personality & Social Psychology, Vol. 29*(5), 703-709.
- Zanna, M. P., & Cooper, J. (2000). Dissonance and the pill: An attribution approach to studying the arousal properties of dissonance. In E. T. Higgins & A. W. Kruglanski (Eds.), *Motivational science: Social and personality perspectives*. (pp. 306-312): Psychology Press.
- Zillmann, D., & Bryant, J. (1974). Effect of residual excitation on the emotional response to provocation and delayed aggressive behavior. *Journal of Personality & Social Psychology, 30*(6), 782-791.

Summary

The present thesis deals with cognitive feelings, which are derived from an ease or fluency dimension of ongoing cognitive processes. Such experiences influence many judgments and inferences: For example, judgments about the frequency of events, inferences about the truth of statements or decisions whether an object is recognized or not. It is argued that the specific impact of such experiences depends on the interpretation of a given cognitive feeling, which provides meaning for a hitherto unspecific experience. Furthermore, it is argued that such interpretations may either be explicitly provided or implicitly learned. The empirical section supporting this claim is divided into two parts:

Part I shows the impact of such interpretations when they are explicitly provided. This is done within the realm of “ease of retrieval” experiences and their influence on evaluative judgments. Two experiments are reported that lend support to the idea that cognitive feelings, in this case, the experienced ease of retrieval, depend on the available interpretations. Part II shows the impact of such interpretations when they are learned via feedback from the environment. This is done within the realm of fluency experiences and their influence on recognition judgments. Usually, stimuli that are more fluently processed have a higher probability to be classified as “old” than non-fluently processed stimuli. Two experiments show that this robust finding can be reversed when the environment provides feedback that fluency indicates that a stimulus is new rather than old. That is, it is shown that the impact of fluency is dependent on the learned interpretation of fluency.

Building on the presented empirical evidence, the discussion addresses the assumptions underlying the interpretation hypothesis, the implications for existing results from research on cognitive feelings, and new research questions that arise from this model of how we interpret the experience of thinking.

Erklärung

Ich erkläre hiermit, dass ich die vorliegende Dissertation selbstständig angefertigt und keine anderen als die angegebenen Quellen oder Hilfsmittel verwendet habe. Weder die vorliegende Arbeit noch Teile davon sind oder waren Grundlage einer anderen akademischen Prüfung.

Heidelberg, Februar 2005

Christian Unkelbach