You don't have to believe everything you read:
Background knowledge permits fast and efficient validation of information

Tobias Richter, Sascha Schroeder, and Britta Wöhrmann

University of Cologne

accepted for publication in *Journal of Personality and Social Psychology* (2008)

Author note

The research reported in this article was supported by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG; grant RI 1100/2-2 given to Tobias Richter and Norbert Groeben). We would like to thank Inga Hoever and Anna Sieben for their help in preparing stimulus materials and collecting data.

Corresponding author:

Tobias Richter

University of Cologne, Department of Psychology

Bernhard-Feilchenfeld-Str. 11

50969 Köln, Germany

E-mail: tobias.richter@uni-koeln.de

Telephone: +49-221-470 2773

Fax: +49-221-470 5002

Abstract

In social cognition, knowledge-based validation of information is usually regarded as relying on strategic and resource-demanding processes. Research on language comprehension, in contrast, suggests that validation processes are involved in the construction of a referential representation of the communicated information. This view implies that individuals can use their knowledge to validate incoming information in a routine and efficient manner. Consistent with this idea, Experiments 1 and 2 demonstrated that individuals are able to reject false assertions efficiently when they have validity-relevant beliefs. Validation processes were carried out routinely even when individuals were put under additional cognitive load during comprehension. Experiment 3 demonstrated that the rejection of false information occurs automatically and interferes with affirmative responses in a non-semantic task (epistemic Stroop effect). Experiment 4 also revealed complementary interference effects of true information with negative responses in a non-semantic task. These results suggest the existence of fast and efficient validation processes that protect mental representations from being contaminated by false and inaccurate information.
Key words: beliefs, comprehension, situation model, truth value, validation

You don't have to believe everything you read: Background knowledge permits fast and efficient validation of information

When scientists review publications in their field of expertise, when judges listen to testimonies in the courtroom, or when educated laypeople read newspaper articles about a topic which is relevant to them, they usually not only comprehend but also validate the presented information by judging its truthfulness or plausibility. Despite the close connection of the two types of activities in everyday life, there is a widespread consensus in the social and cognitive psychology literature that comprehension and validation of information form subsequent stages of processing with very different characteristics. Theories of language comprehension are primarily concerned with stimulus-driven comprehension processes that proceed in an automatic fashion (e.g., Kintsch, 1988; van Dijk & Kintsch, 1983). Theories of attitude change, in contrast, largely portray knowledge-based validation as relying on thoughtful and slow processes that occur only under specific conditions (e.g., Chen & Chaiken, 1999; Petty & Wegener, 1999). In a similar vein, Gilbert (1991) has proposed a dual-stage model that makes the explicit assumption of an initial stage of effortless comprehension followed by an optional stage of effortful validation.

In contrast to these theoretical propositions, ideas and findings from social cognition and language comprehension research suggest a more regular role of validation in the processing of incoming information. In the psychology of language, it is now a commonplace assumption that comprehension entails the construction of a referential representation of the communicated information: The content of this representation is not the message itself but rather the state of affairs that is described in the message (Johnson-Laird, 1983; van Dijk & Kintsch, 1983; Zwaan & Radvansky, 1998). It is difficult to explain how comprehenders would be able to construct adequate referential representations without relying on routine and efficient validation processes

(Wyer & Radvansky, 1999). Likewise, there is evidence from a variety of social and cognitive psychological paradigms demonstrating that individuals spontaneously activate their knowledge to validate incoming information (e.g., Schul & Burnstein, 2004; Trope & Gaunt, 2000). Vice versa, once a referential representation has been constructed it may be used to validate new information, leading to constructive judgment biases and perseverance of false information (e.g., Fiedler, Armbruster, Nickel, Walter, & Asbeck, 1996; Ross, Lepper, Strack, & Steinmetz, 1977; Schroeder, Richter, & Hoever, in press).

Against this background, we will argue that validation is a routine companion of comprehension. Starting from a discussion of the dual-stage model proposed by Gilbert (1991), we will develop the alternative view that individuals rely on fast and efficient validation processes in order to construct an adequate referential representation of the communicated information. We will review research on language comprehension and social cognition that highlights conditions and consequences of these processes. One controversial implication of the view advocated here is that individuals are able to reject false information fast and efficiently when they have accessible, certain, and relevant background knowledge. A second implication is that they do so routinely, i.e. without following specific processing goals. We will report results from four experiments that tested these predictions.

Comprehension vs. Validation of Information: Fast and Automatic vs. Slow and Strategic?

Gilbert (1991) has proposed a dual-stage model of comprehension and validation that is based on ideas of the Dutch philosopher Baruch Spinoza (1677/1997). The core assumption of this model is that the comprehension of information invariably entails its initial acceptance. Only at a later stage, resource-demanding validation processes may take effect, allowing individuals to "unbelieve" false information that was initially accepted as true and to re-represent this information as false. Most of the evidence in favor of this model comes from experiments in which participants learned a number of fictitious facts (e.g., about an imaginary animal called “glark”; Gilbert, Krull, & Malone, 1990) and later had to verify these assertions. In one experiment, an affirmation bias occurred for false assertions that had been speed-read in one of the trials that preceded the verification trial. In another experiment, the word *true* or *false* appeared on a computer screen to indicate the validity of the assertions. In some of the trials, additional cognitive load was induced by a secondary task such as pressing a key in response to a tone that sounded shortly after the presentation of the validity information. Again, an affirmation bias occurred: Participants tended to misjudge false assertions learned in these trials as being true in the verification task but not vice versa. Gilbert et al. (1990) interpreted these results as evidence for the assumption that there is an initial acceptance of assertions that can only be reversed by resource-demanding validation processes at a later stage. Gilbert, Tafarodi and Malone (1993) found a similar affirmation bias for false information presented as part of fictitious crime reports in which some of the sentences were marked as conveying true or false information (identified by different ink colors) and participants were put under time pressure or additional cognitive load by a secondary task while they were reading the crime reports. Those participants who were put under time pressure or had to perform the secondary task, mistakenly judged information marked as false as being true in a later verification task and gave prison-term recommendations that were biased in the direction of the information marked as false.

Over the last decade, Gilbert’s (1991) ideas of an initial acceptance of information vs. an optional and effortful validation have inspired research on a variety of persuasive effects and other phenomena, ranging from age differences in the rejection of false information (Chen & Blanchard-Fields, 2000), belief change through fictional narratives (Appel & Richter, 2007; Green & Brock, 2002; Prentice, Gerrig, & Bailis, 1997), unintended effects of warnings about false consumer claims (Skurnik, Yoon, Park, & Schwarz, 2005) to acquiescence in questionnaire responses (Knowles & Condon, 1999). However, the broad impact of the model notwithstanding, the experimental tests of the core assumptions of Gilbert’s (1991) model were concerned with a rather special case of the relation between comprehension and validation. The studies conducted so far were based on pseudo facts such as word definitions in a non-existent language, assertions about an imaginary animal (Gilbert et al., 1990) or a fictitious crime suspect (Gilbert et al., 1993) as stimuli. As a consequence, the information that participants had to study in the course of the experiments and that they later had to verify was not related to any knowledge or beliefs that they could hold. This restriction seems critical because in most situations where individuals validate information, they compare incoming information with their own beliefs about the particular subject matter. In the absence of any validity-relevant background information, the information that an assertion is false is logically as well as psychologically equivalent to the negation of an abstract assertion (cf. Carpenter & Just, 1975). Thus, the results of Gilbert et al. (1990, 1993) leave open the possibility that people might be able to validate and reject false information early in information processing, provided that they possess relevant background knowledge.

The Cartesian View: Comprehension Without Validation

What are the theoretical competitors to a Spinozan dual-stage model that links comprehension to the initial acceptance of the message content? According to Gilbert (1991), the main alternative to such a model is put forward by the French philosopher René Descartes in his *Principles of Philosophy* (1644/2006). Descartes assumed that individuals initially withhold a judgment concerning the validity of a proposition during comprehension (holding it *in aequilibrio*) and determine its truth status at a later point by means of a thoughtful analysis (Descartes, 1644/2006). With these assumptions, the Cartesian model is also a dual-stage model with a comprehension phase and a subsequent resource-dependent validation stage, but unlike the Spinozan model, it implies that assertions are usually comprehended without being assigned a truth value. However, this implication seems problematic from several perspectives: Most of the current philosophical theories of language and theoretical approaches in linguistic semantics, for example, tie the meaning of linguistic expressions either to truth conditions (e.g., Davidson, 2001; Montague, 1973; Partee, 1975) or to assertion conditions under which a speaker is justified to assert a proposition (e.g., Dummett, 1981). In light of these theories, it seems implausible that individuals should be able to comprehend a proposition without any reference to the conditions that would render the proposition true or false. Suggesting a similar conclusion, modern cognitive theories of language comprehension assume that comprehenders draw on their prior knowledge and experience as well as on the communicated information to construct a situation model (Johnson-Laird, 1983; van Dijk & Kintsch, 1983; Zwaan & Radvansky, 1998). Situation models are referential representations of the state of affairs or event described in a message. As such, they radically differ from the verbal input and abstract propositional representations that may be derived from this input. Rather, their structure corresponds to the structure of states of affairs in the world, with entities (objects and persons) and relations between these entities represented within a spatio-temporal framework. For example, the propositional representation of the assertion *Mountains are high* consists of a predicate HIGH which takes the concept MOUNTAINS as an argument (HIGH(MOUNTAINS)). The same assertion can be represented referentially as a visual scene featuring a line of mountains at the horizon that rise high to the sky. The example illustrates two important properties of referential representations. They bear analogical relationships to objects and events in the world and usually include properties derived from prior knowledge and previous experiences. In the present context, it is important to note that there is no psychologically plausible way how individuals could construct such a representation without performing some kind of validation on the incoming information (for a similar argument, see Wyer & Radvansky, 1999)

An Epistemic View: Validation Routinely Accompanies Comprehension

The Cartesian model is not the most powerful competitor to a Spinozan model. A more plausible alternative approach would start from the assumption that validation routinely accompanies the comprehension of incoming information. There are strong theoretical reasons in favor of this view. Most importantly, the idea that the construction of a situation model is an integral part of comprehension strongly suggests that individuals engage in validation processes on a regular basis. In order to construct an adequate representation of the situation described in a communicated message, individuals need to determine whether its contents are plausible or not (Schroeder et al., in press).

How do individuals manage to achieve adequate representations of the state of affairs communicated in a linguistic message? We suggest that they use epistemic validation processes that check whether incoming information is internally consistent and consistent with comprehenders' knowledge. Epistemic validation is a major determinant of whether a particular piece of information becomes part of the current situation model and, ultimately, part of a comprehender's world view. Strategic, slow, and resource-dependent elaborative processes would be unsuited for this purpose. Rather, if validation is indeed a routine companion of comprehension, it must rely on a fast and efficient type of epistemic validation processes that we call *epistemic monitoring*. The nature of these processes and their relationships to existing research in social cognition and the cognitive psychology of language comprehension will be discussed next.

Epistemic Monitoring: Routine and Effortless Validation of Information

Epistemic monitoring processes check for inconsistencies between incoming information on the one hand and elements of the current situation model or knowledge and beliefs retrieved from long-term memory on the other hand. We expect that these processes are carried out routinely and require little cognitive effort, although resolving the inconsistency may be based on more resource-demanding processes. By themselves, epistemic monitoring processes require little cognitive resources because they refer to information that is already part of working memory, such as elements of the currently active situation model, or to elements of long-term memory that can easily be made available by passive memory-based processes (e.g., McKoon & Ratcliff, 1995; Myers & O'Brien, 1998; O'Brien & Albrecht, 1992). As a consequence, even when processing of information is not governed by specific processing goals and cognitive resources are depleted, inconsistencies of incoming information with active contents of working memory or with easily accessible, contextually relevant knowledge in long-term memory should be noticed (cf. Wyer & Radvansky, 1999).

Evidence for Routine Validation Processes

The existence of routine validation processes that use little cognitive resources and enable individuals to reject false or belief-inconsistent information in a fast and efficient manner is a point of contention in social cognition. Some theories such as the dual-state model proposed by Gilbert (1991) seem to suggest that such processes do not exist or are irrelevant in ordinary comprehension situations. Other theoretical approaches are more or less mute about the issue. In persuasion research, for example, rejection of information contained in persuasive messages has often been conceptualized in terms of elaborative processes such as active counter-arguing (e.g., inoculation theory, McGuire, 1964; McGuire & Papageorgis, 1961), or it has been attributed to peripheral cues such as negative affect (Zuwerink & Devine, 1996), mood, and sources that cause individuals to disregard or to derogate the message content (e.g., Bless, Bohner, Schwarz, & Strack, 1990). In comparison to these lines of thinking, the ideas that individuals regularly and efficiently use their own knowledge and beliefs to monitor the validity of incoming information have received less attention. However, despite the lack of a comprehensive theoretical treatment, the idea of routine validation processes is not entirely new to social cognitive research. It coheres well with several prominent approaches in the area, and the empirical evidence reviewed in the next section indirectly suggests a significant role of epistemic monitoring in social information processing. In addition, psycholinguistic research indicates that epistemic monitoring processes are closely tied to the comprehension of verbally communicated information.

Evidence for Routine and Efficient Validation Processes in Social Cognition

Recent studies on resistance against persuasion have shown beliefs to remain unchanged by persuasive attempts even under conditions that involve little thoughtful elaboration, such as low personal involvement and additional cognitive load while the message is processed (Wegener, Petty, Smoak, & Fabrigar, 2004). With regard to these results, Wegener et al. (2004) speculated that there is a non-thoughtful route to resistance that is based on using one's own attitude as an acceptance or rejection cue. To the extent that cognitive mechanisms contribute to the non-thoughtful route to resistance, the concept of epistemic monitoring seems to be well suited to describe these mechanisms. Stronger support for the idea of routine and effortless validation processes comes from research on encoding under distrust (Dodd & Bradshaw, 1980; Schul, Mayo, & Burnstein, 2004). For example, Schul et al. (2004) demonstrated that distrust causes individuals to spontaneously activate knowledge that is incongruent with the content of a message. In one of their experiments, participants provided speeded judgments of whether a target word was an adjective or a noun. The target words were primed by words that were either congruent (e.g., *transient-temporary*) or incongruent (e.g., *hollow-full*) with the target words. In each trial, a mindset of trust or distrust was created by first presenting a round-eyed or a narrow-eyed face that had been associated with true or false message contents in a prior induction phase. Under distrust, incongruent primes caused greater facilitation effects than congruent primes. Given the short stimulus-onset asynchrony (100 ms) between prime and target and the indirect judgment task used in the experiments by Schul et al. (2004), their results cannot be explained by conscious validation processes. Rather, consistent with the assumption of epistemic monitoring processes, these results point to the automatic and spontaneous activation of message-incongruent information.

With respect to the conditions that govern the fast and efficient activation of validity-relevant information, research by Trope and Gaunt (2000) on the integration model of dispositional inference is particularly instructive. In research on person perception, one common finding is that perceivers view behaviors of a target person as indicative of personal dispositions, even if strong situational demands are in effect that prescribe the observed behaviors (e.g., Gilbert, 1997; Gilbert & Malone, 1995; E.E. Jones, 1979). From the perspective of the integration model, both dispositional and situational information are integrated in attributional inferences from the start. However, dispositional factors are frequently more easily available than situational factors, with the result that situational information discounting dispositional attributions tend to be disregarded (e.g., Trope, 1986; Trope, Cohen, & Maoz, 1988). Thus, integration models imply that the likelihood of considering situational information varies with the availability of this information. Consistent with this assumption, Trope and Gaunt (2000) demonstrated that perceivers considered situational information in dispositional inferences even under additional cognitive load if the situational information was made perceptually salient, accessible in memory, or specific to the behavior in question. For the theoretical view advocated here, these findings are important because the informational properties highlighted in the experiments by Trope and Gaunt (2000) correspond to properties of knowledge that should moderate the occurrence of epistemic monitoring processes. Just as salient, easily accessible, and specific situational information is used to discount premature dispositional inferences, we propose that information that is relevant and either active in working memory or easily accessible in long-term memory can be used to validate incoming information in a fast and efficient manner.

Evidence for Routine and Efficient Validation Processes in Language Comprehension

A large body of evidence from the psychology of language completes this view by showing that the validation of incoming information is closely tied to the construction of a referential representation. Research by Singer and colleagues (e.g., Singer, Halldorson, Lear, & Andrusiak, 1992) on bridging inferences in comprehension suggests that knowledge that is inconsistent with implicit background assumptions (enthymemes) of causally related sentences is routinely used in epistemic monitoring. For example, after reading inconsistent causal sequences such as *Dorothy poured the bucket of water on the bonfire - The fire grew hotter*, responses to questions such as *Does water extinguish fire?* were facilitated relative to temporal sequences that were used as controls (Singer, 1993). One interpretation of this finding is that the validation of implicit premises is an integral part of constructing situation models of causal relationships. Along the same lines, experiments on logical-propositional inferences have shown that when a text signals logical relationships between statements (e.g., by sentence particles such as *if*, *or*, or *not*) and premises stated earlier are activated in working memory, comprehenders routinely detect logical inconsistencies and devote more processing time to the sentence where the inconsistency becomes apparent (Lea, 1995; Lea, Mulligan, & Walton, 2005). Similarly, inconsistencies of other types of situational information such as those between protagonists' actions and their previously described goals in a narrative (Albrecht & Myers, 1995) or inconsistencies in information about a protagonist's location (O'Brien & Albrecht, 1992) are detected when the antecedent information is still active in working memory. Finally, Singer (2006) demonstrated with a reading time paradigm that comprehenders tacitly verify information presented in everyday stories even when they do not follow an intentional validation strategy. Reading times for affirmative target sentences in the stories were prolonged when discourse context and pertinent knowledge rendered the target sentence false. Thus, individuals seem to routinely validate information they encounter in a discourse context.

Rationale of the Present Experiments

The present experiments concentrated on the assumptions of the epistemic view (a) that there are routine (non-strategic) validation processes and (b) that these processes are carried out fast and efficiently early in information processing. Unlike the notion of epistemic elaboration processes which seems to be commonly accepted, the idea of routine epistemic monitoring processes that require little cognitive effort is controversial. Experiment 1 scrutinized this idea by including the availability of validity-relevant background beliefs as a potential moderating variable in the experimental paradigm introduced by Gilbert et al. (1990). According to the view outlined here, epistemic monitoring depends crucially on the availability of accessible, certain, and relevant background beliefs. If knowledge with these characteristics is available, individuals are expected to use it for fast and efficient epistemic monitoring processes. As a consequence of these processes, individuals should be protected from representing false information as true even if working memory resources during comprehension are depleted by additional cognitive load.

Experiment 2 extended the perspective of Experiment 1 by investigating whether epistemic monitoring is dependent on specific processing goals of the individual. If validation of information is an optional and strategic component of the processing of incoming information, as most studies on resistance to persuasion and the dual-stage model seem to imply, it is likely that processing goals play a critical role in an individual's decision about whether to engage in effortful validation processes. According to the view that epistemic monitoring is a routine part of information processing, however, specific processing goals are of lesser importance.

Experiments 1 and 2 were designed to replicate and refine the results reported by Gilbert et al. (1990) with the primary goal to establish the availability of background beliefs as a potential moderator of these effects. For this reason, the procedures used in these experiments closely resembled those of the original experiment. In contrast, Experiments 3 and 4 went one step further by providing direct evidence for the routine and automatic character of epistemic validation processes. We used a Stroop-like paradigm in which participants performed speeded judgments of orthographical correctness on individual words. These words were part of simple assertions whose truth value varied. We expected that despite the fact that orthographical judgments are non-epistemic and non-semantic in nature, affirmative and negative responses in these judgments would be impaired by interfering validation processes that occurred automatically in the course of language processing (epistemic Stroop effect).

Experiment 1

Experiment 1 tested the assumption implied by the epistemic view that the availability of accessible, relevant, and certain background beliefs about an issue enables individuals to carry out validation processes fast and efficiently. For that purpose, Experiment 1 used realistic factual statements that were either strongly or weakly related to participant's knowledge.

The experiment was based on procedures introduced by Gilbert et al. (1990, Study 1). In an initial phase of the experiment, participants learned assertions together with their truth values. In some of the learning trials, participants were put under additional cognitive load. In a later phase of the experiment, they were asked to indicate in a verification task whether these assertions were true or false. The stimulus materials of the present experiment, however, differed in important ways from those used by Gilbert et al. (1990, Study 1). We used objectively true or false assertions related to which participants could be expected to have either strong or weak validity-related background beliefs. For assertions associated with weak background beliefs, we predicted the pattern of effects for true and false assertions that Gilbert et al. (1990, Study 1) found: True assertions should be verified accurately, and verification performance for true assertions should not be affected by additional cognitive load during comprehension. For false assertions, accuracy should be impaired by additional cognitive load. For assertions associated with strong background beliefs, in contrast, we expected that additional cognitive load would not impair the verification performance for true or false assertions. In these cases, we assumed that participants would be able to carry out fast and efficient validation processes that would not be affected by additional cognitive load.

Two additional features of Experiment 1 were meant to rule out alternative explanations for the hypothesized moderating role of background beliefs. First, the verification task included new assertions in addition to those that participants had seen in the learning phase. By comparing the results for new assertions and assertions presented in the learning phase, it was possible to delineate effects of validation processes in the learning phase more clearly and separate these effects from belief effects that took place in the verification phase. Second, participants had to provide their verification judgments within a specified time frame that varied in length. If the hypothesized moderator effect for background beliefs was indeed dependent on whether validation processes were carried out in the learning phase or not, it may be expected that the verification of assertions associated with strong background beliefs is affected less strongly by a shorter response time-frame than the verification of assertions associated with weak background beliefs. Alternatively, if the moderating role of background beliefs was dependent on the ability of participants to carry out resource-dependent validation processes during the verification task (a possibility compatible with Gilbert’s, 1991, dual-stage model), even the verification of assertions associated with strong background beliefs would be strongly affected by a shorter response time-frame.

Method

Participants. Thirty-six psychology undergraduates (29 women and 7 men) with an average age of 26.2 years (*SD* = 7.4) took part in Experiment 1.

Stimulus material. The experimental stimuli were assertions in the form of simple predications of the structure "[a] [*concept noun*] [is/has/causes/contains] [a] [*concept noun/adjective*]". The assertions contained only generic concepts and no proper names, and they were either true or false. A norming study was conducted to select assertions for which participants could be assumed either to have strong background beliefs (e.g., *Perfume contains scents*, *Soft soap is edible*) or to have no or only weak background beliefs (e.g., *Krypton is a noble gas*, *Tooth paste contains sulfur*). The participants of the norming study (12 psychology undergraduates not identical to the experimental samples) were presented with 288 assertions (144 true and 144 false). Participants read each assertion on a computer screen in a self-paced manner. On subsequent screens, they were asked to judge the truth vs. falsity of each assertion and to rate their certainty in this judgment on a 6-point-scale (ranging from 1=*very uncertain* to 6=*very certain*). We included these certainty ratings as a meta-judgmental indicator of belief strength (Gross et al., 1995). On the basis of these data, (a) 36 true assertions were chosen which were consistently judged as true (mean agreement: 100%) and were associated with a consistently high judgment certainty (*M* = 5.57; *SD* = 0.23; *true assertions/strong beliefs*), (b) 36 true assertions were chosen which were not consistently judged as either true or false (mean agreement: 50%) and were associated with an overall low judgment certainty (*M* = 3.75; *SD* = 3.12, *true assertions/weak beliefs*), (c) 36 false assertions were chosen which were consistently judged as false (mean agreement: 98%) and were associated with a consistently high judgment certainty (*M* = 5.45 ; *SD* = 0.26, *false assertions/strong beliefs*), and (d) 36 false assertions were chosen which were not consistently judged as either true or false (mean agreement: 50%) and were associated with an overall low judgment certainty (*M* = 3.92; *SD* = 3.00, *false assertions/weak beliefs*).

Procedure. The experiment consisted of a learning phase and a test phase. During the *learning phase*, participants were presented with 96 assertions and information about whether the assertion was true or false. Before the actual learning phase started, there was a practice phase involving four assertions. Participants were wearing headphones throughout the experiment. They were instructed that their task was to keep in mind the assertions shown on the screen and the validity of these assertions as they would be tested for both kinds of information in a later phase of the experiment. Participants were also told that when they hear a tone over their headphones, they should press the space bar on the computer keyboard as quickly as possible. Presentation of the statements followed a procedure similar to the one used by Gilbert et al. (1990). Each assertion was presented for a duration that corresponded to the maximum reading time determined in a pilot study (*M* = 3388 ms; *Min* = 1390 ms for *Epidemics are dangerous* [in German: *Epidemien sind gefährlich*]; *Max* = 5892 ms for *A fast food restaurant offers delicacies* [*Ein Imbiss hat Delikatessen*]). The assertions were presented in black letters (height 2 cm, font type arial) in a white 20 X 4 cm square placed in the middle of the screen against a teal background. The viewing distance was approximately 60 cm. After a blank screen, which appeared for two seconds, the German word for *true* appeared for three seconds when the preceding assertion was true and the German word for *false* appeared when the preceding assertion was false. On half of the trials, a 1000-hz tone was sounded during the presentation of the validity information until the participant pressed the space bar (interference trials). During the *test phase*, participants saw 96 yes/no-questions corresponding to the assertions which they encountered during the learning phase plus 48 yes/no-questions corresponding to assertions not presented in the learning phase (e.g., *Does perfume contain scents?*). Following a fixation cross presented for 500 ms, the questions appeared in the middle of the screen either for 2423 ms (long time-frame), 1721 ms (medium time-frame), or 1375 ms (short time-frame). These values correspond to quartile boundaries of the reading times distribution as determined in a pilot study. They were followed by a blank screen for 2000 ms. Participants were instructed to respond to the questions as quickly as possible by pressing one of two response keys for *yes* (marked green) or *no* (marked red). They were told that each question would be presented for a limited time, and that they should try to respond within this time frame. Only accurate responses that were provided within the response-time frame were counted as correct responses. Before the actual test trials, there were three practice trials to make participants familiar with the length of the different response time-frames.

Design. The design was a 2(*validity*: true vs. false) X 2(*background beliefs*: strong vs. weak) X 3(*learning condition*: no interference vs. interference vs. no learning) X 3(*response time-frame*: long vs. medium vs. short) within-subjects design. The assignment of assertions to the levels of the factors learning condition and response time-frame was counterbalanced across participants by nine different lists based on Latin squares. In both the learning phase and the test phase, the assertions were presented in random order.

Results and Discussion

The reported *F*-values are based on by-subjects analyses and a multivariate approach to repeated-measures ANOVA which is robust against violations of the sphericity assumption. For all significance tests, the α-level was set at .05. We report partial η^2 as a measure of effect size (Cohen, 1988). The most complex effect tested was a four-way interaction of all independent variables. For the test of this effect, the design of Experiment 1 yielded a power (1-β) of .79 under the assumptions of a medium effect size (f = .25 according to Cohen, 1988) and medium correlations between the levels of the independent variables ($\rho = .5$) in the population (power calculations were done with the software GPower, Erdfelder, Faul, & Buchner, 1996).

The two views on the nature of validation processes contrasted in Experiment 1 implied different patterns of interaction effects. According to the common view that validation is confined to slow and effortful processes, additional cognitive load in the learning phase should impair the proportion of correct verification responses for false assertions but not for true assertions, regardless of whether the assertions were associated with weak or strong background beliefs. According to the epistemic view proposed here, in contrast, this pattern should occur only in assertions associated with weak background beliefs but not in assertions associated with strong background beliefs. In a repeated measurements ANOVA, the two-way interaction of validity and learning condition turned out to be significant, $F(2,34)=5.8$, $p<.01$, $\eta^2=.25$. Similar to Gilbert et al. (1990, Study 1), the additional cognitive load in the interference trials lowered the accuracy of responses in the verification task for false assertions (no interference: $M=.69$, $SE_M=.03$; interference: $M=.62$, $SE_M=.03$) but not for true assertions (no interference: $M=.75$, $SE_M=.02$; interference: $M=.76$, $SE_M=.02$). However, as predicted by the epistemic view of efficient validation processes, this interaction was qualified by a three-way interaction with the availability of background beliefs, $F(2,34)=3.3$, $p<.05$, $\eta^2=.16$. Four planned single degree of freedom a priori contrasts (Abelson & Prentice, 1997) were conducted to test whether the impact of the cognitive load manipulation on the verification of true and false assertions associated with strong and weak background beliefs did indeed vary with the availability of background knowledge. For false assertions associated with weak background beliefs, learning with additional cognitive load ($M=.45$, $SE_M=.03$) caused a significantly lower accuracy of verification responses than learning without additional cognitive load ($M=.55$, $SE_M=.04$), $F(1,35)=11.2$, $p < .01$, $\eta^2=.24$ (Figure 1a). However, as predicted by the epistemic view, for false assertions associated with strong background beliefs, the verification performance did not differ significantly between associations learned with additional cognitive load (M=.79, SE_M=.03) and those learned without cognitive load (M=.83, SE_M=.02), $F(1,35)=3.7$, $p > .05$ (Figure 1b). Similarly, for true assertions associated with weak background beliefs, the comparison between learning with additional cognitive load and learning without additional cognitive load was not significant, $F(1,35)=0.1$, $p > .05$. Finally, for true assertions associated with strong background beliefs, the comparison between learning with additional cognitive load and learning without additional cognitive load was not significant either, $F(1,35)=0.1$, $p > .05$.

Additionally, we set up two single degree of freedom contrasts that tested our focal prediction in terms of the difference between the interference effects for true vs. false assertions (Judd & McClelland, 1989; see Table A1 in the Appendix). The first contrast referred to learned assertions associated with weak background knowledge. It incorporated the null hypothesis that the effects of interference would be equal for true and false items. For the learned assertions associated with weak background knowledge, both the common view of slow and effortful validation processes and the epistemic view predict that interference by the secondary task affects validation of false items more than true items, implying that the null hypothesis should be rejected. In line with this prediction, the contrast turned out to be significant, $F(1,35)=7.8$, $p < .01$, $\eta^2=.18$. The second contrast incorporated the analogous null hypothesis for learned assertions associated with strong background knowledge (Table A2 in the Appendix). For these assertions, the common view of slow and effortful validation processes again predicts greater effects of interference for false compared to true items, implying rejection of the null hypothesis, whereas the epistemic view predicts no such difference, implying that the null hypothesis can be maintained. In line with the epistemic view, the contrast did not reach significance, $F(1,35)=1.7$, $p = .20$.

Another way to phrase the predictions of the epistemic view is to state that verification performance for the learned items should be good in all conditions except for the condition in which false assertions associated with weak background beliefs were learned with additional cognitive load. Only in this specific condition, validation during initial comprehension should fail, with the result of impaired verification performance in the test phase. Accordingly, we also set up a single degree of freedom contrast comparing this condition (coded with -1) against the mean of all other seven conditions with learned items (coded with -1/7) as an alternative way to test the epistemic view. This contrast was significantly different from zero, $F(1,35)=156.9$, $p < .001$, $\eta^2=.82$.

Thus, cognitive load during learning impaired verification performance only in false assertions associated with weak background beliefs but not in false assertions associated with strong background beliefs. These results suggest that even when individuals are placed under a high cognitive load during comprehension, they are not necessarily forced to accept everything they comprehend as being true. In contrast to the assumption of effortful validation processes, the availability of strong background beliefs seems to enable people to reject invalid information in a fast and efficient manner.

Two additional interaction effects of background beliefs and response-time frame back up the interpretation that participants indeed carried out efficient validation processes during the learning phase (as opposed to resource-dependent validation processes in the verification phase). First, availability of background beliefs interacted with response-time frame, $F(2,34)=3.9$, $p<.05$, $\eta^2=.19$. The proportion of correct responses in the verification task declined less strongly from the long to the short time frame in assertions associated with strong background beliefs (long: $M = .90$, $SE_M = .02$, medium: $M = .83$, $SE_M = .02$, short: $M = .69$, $SE_M = .03$) compared to those associated with weak background beliefs (long: $M = .59$, $SE_M = .02$, medium: $M = .48$, $SE_M = .03$, short: M = .31, SE_M = .03). This bolstering effect of background beliefs was indeed due to the fact that participants were able to validate assertions in the learning phase, as indicated by a three-way interaction effect of availability of background beliefs, response-time frame and learning condition, $F(4,32)=6.7$, $p<.05$, $\eta^2=.46$: Only if participants verified assertions that were associated with strong background beliefs *and* presented in the learning phase, the decline from the long to the short response-time frame was relatively moderate (simple main effects of response-time frame: $\eta^2=.28$ and $\eta^2=.41$ for assertions learned with/without additional cognitive load). In all other conditions, there was a steep decline from the long to the short response-time frame (simple main effects from $\eta^2=.60$ to $\eta^2=.64$). This pattern of effects suggests that, as expected, the effects of background beliefs were located already in the learning phase, i.e. during initial encounter with the assertions. In addition to these theoretically relevant ordinal interactions, there were main effects of all four independent variables that primarily underscore the validity of the experimental manipulations.[1]

Experiment 2

Experiment 2 was conducted to consolidate and extend the conclusions drawn from the previous experiment. The primary focus was on the question whether validation processes are strategic, i.e. depend on specific processing goals. According to the epistemic view advocated here, epistemic monitoring processes are routinely involved in the construction of a referential representation of the communicated information. As a consequence, individuals are expected to engage in these processes non-strategically, i.e. regardless of which specific processing goal they follow. Alternatively, validation processes could be voluntary and strategic, i.e. depend on particular processing goals. The latter view is akin to the assumption of the dual-stage model that the effortful "unbelieving" of initially accepted information in the validation phase is optional (Gilbert, 1991). We tested these contrary predictions in Experiment 2 with a design similar to the one of the previous experiment by setting two different processing goals for the learning phase, one that did and one that did not require the validation of information: Participants either studied the assertions with the goal to perform well in a later verification task (validation goal) or with the goal to perform well in a later recognition task (memorization goal).

Method

Participants. Thirty-eight psychology undergraduates (24 women and 14 men) took part in Experiment 2. Their average age was 28.1 years (*SD* = 6.5).

Stimulus material. In Experiment 2, the stimulus assertions were drawn from the same pool of pretested assertions that were used in Experiment 1. However, only assertions for which participants could be assumed to have strong background beliefs were used in Experiment 2. In particular, (a) 72 true assertions were used which were consistently judged as true (mean agreement: 95%) and which were associated with a consistently high judgment certainty (*M* = 5.45; *SD* = 0.35; *true assertions*), and (b) 72 false assertions were chosen which were consistently judged as false (mean agreement: 93%) and which were associated with a consistently high judgment certainty (*M* = 5.21; *SD* = 0.62, *false assertions*).

Procedure. The procedure of Experiment 2 was similar to Experiment 1, including a learning phase, in which 96 of the 144 assertions were presented and a secondary task (reacting to a 1000-hz tone by a key press) was given in half of the trials, and a test phase with 96 assertions already presented in the learning phase and 48 new assertions. However, there were three major differences to Experiment 1. The first difference was that no information on the validity of assertions was provided in the learning phase. In each learning trial, one assertion was presented for a duration that corresponded to the maximum reading time determined in a pilot study. Then a blank screen appeared for two seconds and the next trial started. The second difference was that one out of two instructions was given at the beginning of the learning phase. Participants were instructed either that their task was to keep in mind the validity of the presented assertions (validation goal) or that their task was to keep in mind the presented assertions (memorization goal). In both cases, they were also told that they would be tested for the respective information in a later phase of the experiment. The third difference was that for all participants, the test phase consisted of a verification task and a recognition task. Each of the two tasks was based on 72 assertions (48 presented in the learning phase and 24 not presented in the learning phase). The procedure of the verification task was the same as in Experiment 1, with yes/no-questions presented for either a long, a medium, or a short time-frame (2423 ms vs. 1721 ms vs. 1375 ms) during which participants were requested to provide their response. The procedure of the recognition task was identical to that of the verification task, except for the fact that participants had to indicate whether an assertion had been presented in the learning phase or not.

Design. The design was a 2(*processing goal*: validation vs. memorization) X 2(*validity*: true vs. false) X 3(*learning condition*: no interference vs. interference vs. no learning) X 3(*response time-frame*: long vs. medium vs. short) design, with processing goals varied between subjects and the other three independent variables varied within subjects. The assignment of assertions to the levels of the factors learning condition, response time-frame and to the verification or recognition task as well as the order of the verification and the recognition tasks were counterbalanced across participants by 18 different lists based on Latin squares. In both the learning phase and the test phase, the assertions were presented in a random order.

Results and Discussion

For the test of the hypothesized interaction of the between-subjects variable processing goal and the within-subjects variables learning condition and validity, the design of Experiment 2 yielded a power of .58 for the verification task under the assumptions of medium effect sizes ($f = .25$) and medium correlations between the levels of the independent variables ($\rho = .5$) in the population. For the most complex possible effect involving only within-subjects variables, the power was .79 under the same assumptions. We report results for the verification task only because only these data are relevant to our hypotheses.

According to the view that validation processes are voluntary and strategic, the extent to which participants validate information should critically depend on whether they follow a validation or a memorization goal. In particular, the rejection of false information in the interference trials should be improved when a validation goal is followed. In contrast, according to the epistemic view, participants should routinely validate true as well as false information regardless of their processing goal, given that validity-relevant beliefs are available. In line with the latter view, independent of participant's processing goal, the verification of assertions associated with strong background beliefs was not impaired by the additional cognitive load in the interference trials. Assertions that participants had learned without interference (M=.82, SE_M=.02) were verified as accurately as those they had learned with interference (M=.82, SE_M=.02), but both types of assertions were verified more accurately than those that had not been presented in the learning phase (M=.73, SE_M=.02), with $F(2,35)=9.6$, $p<.001$, $\eta^2=.35$ for the global main effect of learning condition. As expected, the planned contrast comparing learning with interference and learning without interference was not significant, $F(1,36)=0.0$, but the contrast comparing assertions from these two conditions to assertions that had not been presented in the learning phase was significant, $F(1,36)=19.5$, $p<.001$, $\eta^2=.35$. Apparently, participants were able to validate true as well as false assertions quite reliably even when they were put under additional cognitive load.

The discussed pattern of effects occurred independently of the processing goal that participants were given in the learning phase, and the validation goal did not enhance verification performance in general, with $F(1,36)=0.6$ for the main effect, or the rejection of false information in particular, with $F(1,36)=0.6$ for the interaction of processing goal and validity (Figure 2). Thus, given that the power for the interaction effects with the between-subjects-factor was acceptable, it seems that there was little support for the assumption that participants following the goal to validate information should be less susceptible to false information than participants following the goal to memorize this information. Nevertheless, we further scrutinized the competing views of routine vs. strategic validation processes by testing two separate one degree of freedom contrasts for the validation goal and the memorization goal conditions (see Table A3 and Table A4 in the Appendix). The first contrast incorporated the null hypothesis that regardless of processing goal, any effects of interference during learning on verification performance would be equal for true and false items. This is exactly what the epistemic view predicts, implying that the contrast should not be significant. In line with this expectation, the contrast was not significant, $F(1,36)=0.2$, $p=.69$. The second contrast represented the null hypothesis that any moderator effect of validity on the difference between the verification of true vs. false items would not depend on processing goal. Again, the epistemic view predicts that the null hypothesis holds whereas the view of voluntary and strategic validation processes implies that the null hypothesis must be rejected. Consistent with the epistemic view, the second contrast failed to reach significance as well, $F(1,36)=1.0$, $p=.33$. Thus, tests of contrasts representing the predictions of the two competing views corroborate the predictions of the epistemic view. The results suggest that individuals routinely engaged in validation processes during the learning phase as no specific processing goal seems to be required to initiate these processes.

One other aspect of the results should be noted. There was an interaction of learning condition and validity, $F(2,35)=4.8$, $p<.05$, $\eta^2=.22$. This interaction was due to the fact that the differences in the proportions of correct responses to true and false assertions were small for assertions presented in the non-interference trials (true assertions: $M=.85$, $SE_M=.02$; false assertions: M=.78, SE_M=.02) as well as the interference trials (true assertions: M=.85, SE_M=.02; false assertions: M=.79, SE_M=.03) but much larger for assertions that had not been presented in the learning phase (true assertions: M=.82, SE_M=.02; false assertions: M=.65, SE_M=.03). While the results for the learned items are fully consistent with the idea of fast and efficient validation processes, the difference between true and false items that had not been presented in the learning phase seems to be puzzling at first sight. Does the advantage for non-learned true over non-learned false items reflect an affirmation bias that contradicts the assumptions of the epistemic view? There are good reasons speaking against this interpretation. First of all, the two-phase structure of Experiments 1 and 2 (which is adopted from Gilbert et al., 1990, Study 1) allows dissociating comprehension and verification effects only for items that are presented both in the learning phase and in the test phase. In contrast to the learned items, any effects found for the non-learned items can be due to differences in the ease of comprehension, verification, or both. For instance, it is well conceivable that the true assertions used in the experiments were easier to comprehend because they referred to situations that were more familiar to participants. The situation described in the true assertion *Perfume contains scents*, e.g., might be more familiar to participants than the situation described in the false assertion *Soft soap is edible*. In line with this speculation, the mean reading times measured in the pilot study for true assertions were shorter (M = 1846 ms, SE_M = 58) compared to false assertions (M = 2051 ms, SE_M = 58; $t(142) = 2.5$, $p < .05$, $d = 0.42$). Thus, the reading time data suggest that the differences in the verification performance for true vs. false items that were not presented in the learning phase actually reflect content-dependent differences in the ease of comprehending these assertions. Importantly, however, these differences do not affect the interpretation of the results for assertions learned with and without interference which are central for our hypotheses.[2]

Despite its clear and consistent results, Experiment 2 leaves some doubts concerning the routine character of validation processes, as implied by the concept of epistemic monitoring. Strictly speaking, the absence of any effects of the processing goal manipulation on verification performance does not preclude the possibility that participants in the recognition condition still applied some kind of validation strategies. In addition, it might be argued that some aspects of the paradigm used in Experiments 1 and 2, such as the two-phase procedure with separate learning and test phases render the experiments suboptimal for the study of fast and efficient processes. Thus, the results of these experiments establish the availability of background beliefs as a moderator of the affirmation bias found by Gilbert et al. (1990) but they do not provide direct evidence for the hypothesized routine and efficient character of epistemic validation processes. For this reason, we conducted two further experiments based on a different paradigm that is more informative in this respect.

Experiment 3

One way to investigate the routine character of epistemic monitoring processes more directly is to establish that these processes interfere with the execution of a completely unrelated task. A classical and powerful example for a paradigm that produces such interference effects is the Stroop task (Stroop, 1935), which has been adapted for the study of automatic processes in a large variety of fields (McLeod, 1991). Experiment 3 used a new variant of the Stroop task to provide a strict and direct test of the automaticity of validation processes. Participants performed orthographical judgments on individual words that appeared one by one on a computer screen. Sequences of words formed simple assertions that were either true or false. The critical trials were those in which the target word was the last word of an assertion associated with strong background beliefs. If individuals routinely perform epistemic monitoring processes, as suggested by the epistemic view, they should automatically develop a negative response tendency when a false assertion is encountered. The negative response tendency, in turn, should interfere with affirmative responses in the orthographical judgment task. Thus, we predicted an increase of response latencies and error rates when the orthographical judgment task required an affirmative response (i.e., the last word was spelled correctly) but the assertion was false. Given that orthographical judgments are unrelated to the validity of the presented information, the predicted effect may be regarded as a Stroop-like interference effect (epistemic Stroop effect).

In designing Experiment 3, we took care to minimize the possibility that the strength with which the words in the assertions are associated in the mental lexicon could account for the predicted effects. Due to the sequential character of the presentation, words in the subject position of the experimental assertions possibly prime words in the object position, and the magnitude of these priming expects is likely to depend on the strength of the association between the two kinds of words. To control for such associations, we made sure that the semantic associations between the words in the subject and the object positions of the assertions did not differ between valid and invalid experimental assertions. As a measure of the strength of semantic association, we computed the cosine of the words in the subject and object positions of each assertion in Latent Semantic Analysis (LSA) space (Landauer & Dumais, 1997). LSA uses co-occurrences of words in large corpora of written language to derive a high-dimensional semantic space (usually around 300 dimensions) via a principal components analysis of the co-occurrence matrix. In the semantic space derived by this procedure, words (and more complex linguistic expressions) can be represented as vectors. The cosine of the vectors of either two linguistic expressions may be interpreted as a measure of their association in semantic space. For the present purpose, cosines computed in LSA space seem to be well suited because they represent a relatively pure measure of semantic relatedness in terms of associative relationships. Among other things, it has been shown that the cosines computed in LSA space contribute to predicting the magnitude of semantic and associative priming effects (M.N. Jones, Kintsch, & Mewhort, 2006).

Method

Participants. Thirty-two psychology undergraduates (25 women and 7 men) took part in Experiment 3. Their average age was 24.4 years (*SD* = 8.5). All participants were native speakers of German.

Stimulus material. Stimuli were assertions taken from the pool of pretested assertions that were also used in the previous experiments. As experimental items, we used 64 three-word assertions for which participants could be assumed to have strong background beliefs. Thirty-two of these assertions had consistently been judged as being true (mean agreement: 100%) with a high judgment certainty (*M* = 5.70; *SD* = 0.15, *true assertions*), and the other 32 assertions had consistently been judged as being false (mean agreement: 98%) with a high judgment certainty (*M* = 5.65; *SD* = 0.20, *false assertions*). For all pairs of words on the subject and object position of the experimental assertions, we computed the cosine of the in LSA space to check whether true and false assertions were equivalent in terms of the strength of semantic associations between their component words. Two different LSA spaces were used to cross-validate the results. First, we used an LSA space (300 dimensions) based on a corpus of German informational texts (Lenhard, Baier, Hoffman, & Schneider, 2006, see http://www.summa.psychologie.uni-wuerzburg.de/). Second, we translated the experimental assertions into English and replicated the analyses with the semantic space based on the standard LSA corpus of English texts (General Reading up to 1st year in college, 300 dimensions, see http://lsa.colorado.edu/). According to both types of LSA cosines, there were no systematic differences between true assertions (German LSA space: $M = .11$, $SE_M = .02$; English LSA space: $M = .17$, $SE_M = .03$) and false assertions (German

LSA space: $M = .12$, $SE_M = .03$; English LSA space: $M = .22$, $SE_M = .03$) in terms of the strength of semantic associations (for both comparisons: $t(62) < 1.4$, $p > .17$).

For each of the experimental items, an orthographically incorrect version was constructed by exchanging, inserting, or removing one letter in the third word of the assertion. In addition to the experimental items, 80 assertions were used as filler items. Among the filler items, there were 16 true and 16 false assertions associated with strong background beliefs. Half of these filler items contained an orthographical error in the first or the second word. Moreover, the filler items included 24 true and 24 false assertions associated with weak background beliefs. Of these 48 items, 24 were spelled correctly. Eight items contained an orthographical error in the first word, 8 items contained an error in the second word and 16 items contained an error in the third word.

Procedure. Participants performed judgments of orthographical correctness for the 64 experimental and 80 filler items that were presented word-by-word on the computer screen. They wore headsets throughout the experiment. The words were presented at a fixed rate of 750 ms per word with a blank screen of 1000 ms after each trial of three words. At one target word in each item, the rapid serial visual presentation (RSVP) halted and participants were prompted by a 1000-Hz tone to judge whether the word was spelled correctly or incorrectly. The tone was initiated 400 ms after the onset of the presentation of the target word. Participants provided their judgments by pressing one of two response keys. Responses and response latencies were recorded. Immediately after the response, the next word of the assertion or the blank screen, which was shown between trials, appeared. In experimental trials, participants were always prompted to respond to the third word. Half of the experimental trials presented orthographically correct items, requiring a yes-response, and half presented incorrect items, requiring a no-response. In filler items, participants were prompted to respond to the first word (24 trials), the second word (24 trials), or the third word (48 trials). The filler trials with the first or the second word as target words were included to prevent participants from focusing their attention exclusively on the third word in each trial. Again, half of the filler items were orthographically correct, requiring a yes-response, and half were orthographically incorrect, requiring a no-response.

Design. The design was a 2(*validity*: true vs. false) X 2(*orthographical correctness*: correct vs. incorrect) within-subjects design. Dependent variables were the response latencies and the accuracy of the orthographical correctness judgments. Assignments of true and false experimental assertions to the levels of the factor orthographical correctness were counterbalanced across participants by four item lists. Experimental and filler items were presented in random order.

Results and Discussion

Given the design and the sample size of Experiment 3, the power for detecting the focal interaction of validity and orthographical correctness was .97 under the assumptions of a medium effect size ($f = .25$) and medium correlations between the levels of the independent variables ($\rho = .5$) in the population.

Response latencies. Response latencies deviating more than two standard deviations from the mean of the experimental condition (less than 2.9% of all latencies) were treated as outliers and removed from the data set (Ratcliff, 1993). According to the epistemic Stroop effect predicted by the epistemic view, yes-responses in the orthographic judgment task should be slower in false compared to true assertions. In line with this prediction, a repeated measurements ANOVA of the latencies of correct responses revealed a large interaction of validity and orthographical correctness, $F(1,31) = 13.9$, $p < .001$, $\eta^2 = .31$. Planned contrasts showed that yes-responses in the orthographical judgment task were considerably slower for invalid assertions ($M = 892.9$, $SE_M = 53.4$) compared to valid assertions ($M = 743.9$, $SE_M = 39.2$), $F(1,31) = 44.6$, $p <$

.001, $\eta^2 = .59$. No-responses were also slower for invalid assertions ($M = 728.8$, $SE_M = 42.0$) compared to valid assertions ($M = 682.4$, $SE_M = 33.1$), but this difference was much smaller than the difference for yes-responses, $F(1,31) = 6.4$, $p < .05$, $\eta^2 = .17$. Thus, the latency data exhibited a strong epistemic Stroop effect.

In addition to the predicted interaction, there was also a main effect of the validity of assertions, $F(1,31) = 41.9$, $p < .001$, $\eta^2 = .58$, reflecting the fact that responses to invalid assertions were generally slower. In addition, there was a main effect of orthographical correctness, $F(1,31) = 39.6$, $p < .001$, $\eta^2 = .56$. No-responses, i.e. judgments that a word was spelled incorrectly, were generally faster than yes-responses, i.e. judgments that a word was spelled correctly.

Error rates. If there is indeed an automatic tendency to reject invalid information, this tendency might not only yield slower yes-responses in the orthographic judgment task but also an increase in erroneous no-responses. We tested this assumption in a repeated measurements ANOVA on the proportion of inaccurate responses. Considering that the overall proportion of inaccurate responses was near 0 (error rate: $M = .05$, $SD = .03$), we re-ran all analyses with arcsine-transformed proportions to double-check for artifactual results due to restricted variances. In line with our predictions, we found an interaction of validity and orthographical correctness, $F(1,31) = 4.6$, $p < .05$, $\eta^2 = .13$ (arcsine-transformed proportions: $F(1,31) = 3.4$, $p = .08$, $\eta^2 = .10$). Planned contrasts revealed that valid assertions were judged less accurately when the critical word was spelled incorrectly and required a no-response ($M = .08$, $SE_M = .02$) compared to when the last word was spelled correctly and required a yes-response ($M = .01$, $SE_M = .00$), $F(1,31) = 18.9$, $p < .001$, $\eta^2 = .38$ (arcsine-transformed proportions: $F(1,31) = 23.8$, $p < .001$, $\eta^2 = .44$). In contrast, for invalid assertions, there was no significant difference between correctly spelled words (M = .04, SE_M = .01) and the incorrectly spelled words (M = .06, SE_M = .01), $F(1,31)$ = 3.3, p = .08 (arcsine-transformed proportions: $F(1,31)$ = 3.0, p = .10). In sum, the accuracy data exhibited an interference effect for valid assertions that may be interpreted as an epistemic Stroop effect but not the expected epistemic Stroop effect for invalid assertions. In addition to the hypothesis-relevant effects, there was a main effect for orthographical correctness, $F(1,31)$ = 20.5, $p < .001$, $\eta^2 = .40$ (arcsine-transformed proportions: $F(1,31)$ = 21.0, $p < .001$, $\eta^2 = .40$), indicating that orthographically correct words were judged more accurately than orthographically incorrect words.

Effects of semantic associations. In order to explore whether the strength of the semantic association between the words in the experimental assertions affects the hypothesized effects, we reran the analyses reported in the previous paragraphs with the cosines computed in the German LSA space as an additional predictor. Given that this measure is a continuous variable varying between experimental trials, it is not possible to incorporate it in a traditional repeated-measurements ANOVA. Rather, a multilevel approach with experimental items as the units of observation on level 1 nested under participants as the units of observation on level 2 is required (Nezlek, 2001; Raudenbush & Bryk, 2002; Richter, 2006). For the response latencies as dependent variable, we estimated a hierarchical linear model with random intercept (random variation between participants) and fixed effects of validity (contrast-coded: 1 = true, -1 = false), orthographical correctness (contrast-coded: 1 = correct, -1 = incorrect), semantic association (z-standardized), and the two- and three-way interactions of these variables as predictors on level 1 (experimental items). The model did not contain any predictors on level 2 (participants). Parameters were estimated with the Restricted Maximum Likelihood/Generalized Least Squares (RML/GLS) algorithm implemented in HLM 6 (Raudenbush, Bryk, Cheong, & Congdon, 2004). For the accuracy data as the dependent variable, an identically structured model was used in combination with a logit link function that allows incorporating binary dependent variables into the Generalized Hierarchical Linear Models framework (e.g., Quené & van den Bergh, in press; Raudenbush & Bryk, 2002, Ch. 10). For this model, parameters were estimated with the Penalized Quasi-Likelihood (PQL) algorithm implemented in HLM 6 (Raudenbush et al., 2004).

If the alternative explanation in terms of strength of the semantic associations holds, there should be an interaction of semantic relatedness and orthographical correctness that parallels the hypothesized interaction of validity and orthographical correctness. Contrary to this prediction, the interaction of semantic relatedness and orthographical correctness failed to reach significance for the response latencies ($\gamma = -2.76$, $SE = 6.70$, $t(1892) = -0.4$, $p = .68$) and for the accuracy data as well ($\gamma = -0.14$, $SE = 0.13$, $t(2040) = -1.1$, $p = .26$). However, the hypothesized interaction of validity and orthographical correctness remained intact in the hierarchical linear model that included semantic relatedness both for response latencies ($\gamma = -24.52$, $SE = 6.79$, $t(1892) = -3.6$, $p < .001$) and for the accuracy data ($\gamma = 0.44$, $SE = 0.14$, $t(2040) = 3.3$, $p < .01$). Thus, it seems unlikely that the epistemic Stroop effects found in the ANOVAs for response latencies and the error rates are actually effects based on semantic relatedness rather than validity.

Taken together, the epistemic Stroop effect that we obtained for the latency data suggests that participants validated the implicitly presented assertions even though neither comprehension nor validation of these assertions was required by the task. These findings are consistent with the assumption of routine epistemic monitoring processes. In particular, the results for the response latency data of Experiment 3 demonstrated that comprehenders routinely detect and reject false information. However, it remains an open question whether they also routinely accept information as well that has passed the epistemic gatekeeper. The error rates but not the latency data suggest such a conclusion. With the following experiment, we sought to answer this question.

Experiment 4

In Experiment 3, response latencies increased when participants had to provide an affirmative response to words that were part of false assertions. However, the complementary interference effect for negative responses to words that were part of true assertions was found only for the error rates. The presence of this effect for the error rates as well as for the latency data would provide strong support for the assumption that comprehenders not only routinely reject false information that has been detected as being false but that they also tend to accept information that has been found true as a result of an epistemic monitoring process. One possible explanation for the failure to find such an effect for the response latencies in Experiment 3 relates to the fact that the presentation rates for the stimulus words were relatively long. In typical experiments that use the RSVP technique for detecting automatic comprehension processes, the presentation times are usually shorter (e.g., 300 ms, Potter, 1984). Thus, it is possible that the long presentation times gave rise to strategic and controlled processes that might have undermined the epistemic Stroop effect in the case of valid assertions. For this reason, Experiment 4 used shorter presentation rates with the aim at to replicate and extend the results of the previous experiment. One group of participants received the stimulus words at a rate of 300 ms per word (comparable to typical RSVP experiments on automatic comprehension processes, Greene, McKoon, & Ratcliff, 1992) and a second group received the words at a rate of 600 ms. The 600 ms presentation rate, which is still in the range for which automatic processes may be expected, was included because more time might be needed to comprehend the German stimulus words which are longer than the average English words used in typical RSVP experiments. In addition to increasing the presentation rate, we increased the proportion of filler items associated with weak background beliefs as a further measure to prevent participants from responding strategically.

Method

Participants. Forty-two psychology undergraduates took part in Experiment 4. Five of them were excluded from the analyses because they were not native speakers of German. The average age of the remaining 37 participants (30 women and 7 men) was 24.1 years (SD = 5.0).

Stimulus materials and procedure. The stimulus materials and procedure of Experiment 4 were identical to that of Experiment 3 except for three changes. First, participants received the stimulus words in the rapid serial visual presentation either at a rate of 600 ms per word or 300 ms per word. Second, a visual rather than an auditory response prompt appeared on the screen when participants had to provide a response. Third, the proportion of filler trials was increased. There were 48 experimental trials (24 valid and 24 invalid assertions) and 104 filler trials (52 valid and 52 invalid assertions). Half of the experimental assertions had consistently been judged as being true (mean agreement: 100%) in the pilot study with a high judgment certainty (M = 5.73; SD = 0.18, *true assertions*), and the other half had consistently been judged as being false (mean agreement: 98%) with a high judgment certainty (M = 5.78; SD = 0.10, *false assertions*). Half of the experimental trials presented orthographically correct items, requiring a yes-response, and half presented orthographically incorrect items, requiring a no-response. According to the LSA cosines between the content words, there were no systematic differences between true assertions (German space: M = .11, SE_M = .03; English space: M = .24, SE_M = .03) and false assertions (German space: M = .10, SE_M = .03; English space: M = .17, SE_M = .04) in terms of the strength of semantic associations (for both comparisons: $t(46) < 1.5$, $p > .15$).

In filler items, participants were prompted to respond to the first word (28 trials), the second word (28 trials), or the third word (48 trials). Half of the filler items were orthographically correct, requiring a yes-response, and half were orthographically incorrect, requiring a no-response. Thirty-two of the filler items were associated with strong background knowledge, and 72 were associated with weak background knowledge.

Design. The design was a 2(presentation rate: 600 ms vs. 300 ms) X 2(*validity*: true vs. false) X 2(*orthographical correctness*: correct vs. incorrect) design with presentation rate varied between subjects and the other two independent variables varied within subjects. Dependent variables were the response latencies and the accuracy of the orthographical correctness judgments. Participants were randomly assigned to one of the two presentation rates. Assignments of true and false experimental assertions to the levels of the factor orthographical correctness were counterbalanced across participants by four item lists. Experimental and filler items were presented in random order.

Results and Discussion

Given the design and the sample size of Experiment 4, the power for detecting the focal interaction of validity and orthographical correctness was .99 under the assumptions of a medium effect size (f = .25) and medium correlations between the levels of the independent variables (ρ = .5) in the population.

Response latencies. Response latencies deviating more than two standard deviations from the mean of the experimental condition (5.2 % of all latencies) were treated as outliers and removed from the data set. An ANOVA of the response latencies revealed a strong cross-over interaction effect consistent with the predictions of the epistemic view, $F(1,35) = 21.8$, $p < .001$, $\eta^2 = .38$. Planned contrasts revealed that in invalid assertions, yes-responses were slower (M = 912, $SE_M = 54$) compared to no-responses (M = 790, $SE_M = 43$), $F(1,35) = 10.0$, $p < .01$, $\eta^2 = .22$. However, in valid assertions, no-responses were slower (M = 758, $SE_M = 37$) compared to yes-responses (M = 697, $SE_M = 38$), $F(1,35) = 5.2$, $p < .05$, $\eta^2 = .13$ (Figure 4). Thus, the epistemic Stroop effect found in Experiment 3 for invalid assertions was replicated. Moreover, Experiment

4 also revealed a complementary epistemic Stroop effect for valid assertions. In addition to the theoretically relevant interaction of validity and orthographical correctness, there was main effect for validity, $F(1,35) = 38.5$, $p < .001$, $\eta^2 = .52$. Overall, responses to invalid assertions were slower (M = 851 ms, $SE_M = 45$) that responses to valid assertions (M = 728 ms, $SE_M = 35$). There was also an interaction of validity and presentation rate, $F(1,35) = 4.2$, $p < .05$, $\eta^2 = .11$. This interaction was due to the fact that the latency difference between valid and invalid assertions was greater for the shorter presentation rate (165 ms) compared to the longer presentation rate (83 ms). Importantly, none of the additional effects affected the interpretation of the epistemic Stroop effect that is central for the epistemic view. It is likely that the double epistemic Stroop effect for invalid as well as valid assertions (as opposed to only the invalid assertions) is due to the shorter presentation rates and the higher proportion of filler items in Experiment 4. The fact that the epistemic Stroop effects are not moderated by presentation rate suggests that both presentation rates used in Experiment 4 (600 ms and 300 ms) were brief enough to prevent participants from responding strategically. This seems plausible given that the German words used in Experiment 4 were considerably longer than the English words used in typical RSVP experiments.

Error rates. An ANOVA performed on the proportion of incorrect responses revealed a strong interaction of validity and orthographical correctness that mirrored the one that we found for the response latencies, $F(1,35) = 12.3$, $p < .001$, $\eta^2 = .26$ (arcsine-transformed proportions: $F(1,35) = 12.5$, $p < .001$, $\eta^2 = .26$). In line with the predictions of the epistemic view, responses to words in valid assertions were more accurate when the target word was spelled correctly (error rates: $M = .03$, $SE_M = 0.01$) and required a yes-response compared to incorrectly spelled target words ($M = .08$, $SE_M = 0.01$), $F(1,35) = 6.3$, $p < .05$, $\eta^2 = .15$ (arcsine-transformed proportions: $F(1,35) = 6.5$, $p < .05$, $\eta^2 = .16$). In invalid assertions, however, responses were more accurate when the word was spelled incorrectly (M = .04, SE_M = 0.01) and required a no-response, compared to correctly spelled target words (M = .07, SE_M = 0.02), $F(1,35) = 4.4$, $p < .05$, $\eta^2 = .11$ (arcsine-transformed proportions: $F(1,35) = 3.6$, $p = .07$, $\eta^2 = .10$). Thus, similar to the response latencies, the error rates revealed epistemic Stroop effects both for valid and invalid assertions.

Effects of semantic associations. As in the previous Experiment, we conducted additional analyses based on multilevel models (experimental items nested within participants) with the cosines derived from the German LSA space as predictor on the item level (level 1) to rule out an alternative explanation of the detected effects in terms of strength of semantic associations. The model specifications and estimation procedures were identical to those used for the data from Experiment 3, except for the fact that presentation rate was included as level 2 (participant level) predictor and the cross-level interactions of this predictor with all predictors on level 1 were included in the models as well. Contrary to what the alternative explanation predicts, the interaction of semantic relatedness and orthographical correctness failed to reach significance for the response latencies ($\gamma = -0.14$, $SE = 7.62$, $t(1570) = -0.02$, $p = .99$) as well as for the accuracy data ($\gamma = 0.004$, $SE = 0.07$, $t(1768) = 0.06$, $p = .95$). In contrast, the interaction of validity and orthographical correctness predicted by the epistemic view remained intact for the response latencies ($\gamma = -28.01$, $SE = 10.47$, $t(1570) = -2.7$, $p < .01$) and the error rates ($\gamma = 0.21$, $SE = 0.07$, $t(1768) = 3.0$, $p < .01$). These results imply that the epistemic Stroop effects that we found in Experiment 4 cannot be attributed to an interaction of semantic relatedness and orthographical correctness.

General Discussion

Four experiments were conducted to test the assumptions that individuals use relevant background knowledge to validate communicated information in a fast and efficient manner, and that they do so routinely, i.e. without following specific processing goals. In Experiments 1 and

2, participants studied simple assertions that were marked as either true or false, and verified these assertions in a later phase. The results of these experiments support the idea of fast and efficient validation processes. In Experiment 1, participants were able to verify assertions when they had strong background beliefs that they could use for a validation of assertions in the learning phase. Both true and false assertions of this kind were verified reliably, even if participants had been put under additional cognitive load in the learning phase. The verification of false assertions that were associated with only weak background beliefs, in contrast, was impaired by additional cognitive load in the learning phase. Apparently, the availability of background beliefs enabled participants to validate information fast and efficiently in the course of comprehending the information. Only for assertions associated with weak background beliefs, verification of false but not of true assertions was impaired, similar to the affirmation bias that Gilbert et al. (1990) has found for assertions about fantasy facts. This finding is consistent with an interpretation in terms of the schema-plus-tag model of negation (Clark & Chase, 1972).

Experiment 2 basically replicated the findings of Experiment 1 for assertions associated with strong background beliefs and shed light on the non-strategic character of validation processes. There was no indication that epistemic validation processes are strategic in the sense that a specific processing goal (validation) is required to trigger these processes. Quite to the contrary, these processes seem to be effective even when comprehenders follow a processing goal (memorization) that does not require them. This finding supports the theoretical proposition that validation through epistemic monitoring processes is a routine companion of the comprehension processes. In the same vein, Experiment 3 established an epistemic Stroop effect that may be regarded as strong evidence for the proposed routine character of validation processes. The automatic tendency to reject false information slowed down affirmative responses in the completely unrelated, non-epistemic task to provide orthographical judgments. Experiment

4 replicated this finding and revealed a complementary epistemic Stroop effect for true assertions, implying that there is also an automatic tendency to accept information that has successfully passed the epistemic monitoring process.

General Implications for Social Cognition

The most important conclusion that may be drawn from the present experiments is the existence of efficient cognitive mechanisms that guard beliefs against belief-inconsistent and false information. When individuals have knowledge available that is accessible, integrated, held with a high subjective certainty, and relevant, they are by no means forced to accept every piece of information communicated to them as being true. Quite to the contrary, they are able to identify and reject knowledge-incongruent information through efficient epistemic monitoring processes. We conceptualize epistemic monitoring processes as general cognitive mechanisms that might play a role in a range of specific phenomena. First of all, the concept coheres well with recent suggestions that there is a fast and efficient but nonetheless cognitive (i.e., belief-based) route to resistance against persuasive communication (Wegener et al., 2004). The theoretical view developed in this article provides an account of the cognitive mechanisms that might underlie this route to resistance, and specifies some of the conditions that might govern these mechanisms. Second, the notion of epistemic monitoring is closely linked to a number of other approaches that assign validation a role in early phases of information processing. In this regard, the experiments by Schul et al. (2004) on encoding under trust and distrust and the work by Dodd and Bradshaw (1980) on suspicion effects on memory for presupposed facts are particularly interesting because they suggests that despite being non-strategic in nature, epistemic monitoring can be modulated by epistemic mindsets. Further research could pursue this line of thinking by including other epistemic mindsets or epistemic motivations such as curiosity or need for cognitive closure (for the concept of epistemic motivations, see Kruglanski & Semin, 2007). A special case of an epistemic mindset that prevents individuals from epistemic monitoring is the state of transportation that is characteristic for readers of fictional narratives (e.g., Appel & Richter, 2007; Green & Brock, 2002; Prentice et al., 1997). Readers who are mentally transported into the fictitious world of the narrative are subject to a temporary suspension of disbelief that makes them susceptible for implicit persuasion.

Implications for the Dual-Stage Model

What are the implications of the present findings for Gilbert's (1991) dual-stage model of comprehension and validation? Our experiments certainly do not falsify this model. Rather, they contribute to a clarification of its scope by highlighting the availability of strong background knowledge as an important moderator of routine, fast, and efficient validation processes. Apparently, the assumption of the dual-stage model that validation – if there is any – is based on resource-dependent validation processes alone applies when comprehenders have little or no strong background knowledge to bear on the validation of a message. This can happen, for example, when lay people process medical information (Skurnik et al., 2005) or episodic information about events they have not witnessed themselves (Gilbert et al., 1993). In these and similar cases, information that is explicitly communicated as being false (e.g., *"Aspirin destroys tooth enamel" is false*) cannot be validated in an efficient manner, the construction of an alternative state of affairs is impossible, and comprehenders have to spend considerable cognitive effort to construct a tagged propositional representation (Clark & Chase, 1972; Mayo, Schul, & Burnstein, 2004). As a consequence and as predicted by the dual-stage model, false information is often represented as being true, particularly when cognitive resources are depleted.

Apart from implying that validation usually relies on resource-dependent, strategic processes, the dual-stage model portrays comprehension and validation as strictly sequential stages of processing. This claim was not in the focus of the present experiments, and our data do not allow any conclusions concerning its validity. Theoretically, however, the epistemic view advocated here implies the contrary assumption that comprehension and validation dynamically interact in the construction of a referential representation. As soon as a new entity, property, or relation is integrated into an existing situation model, its consistency with the current state of the situation model and other available knowledge is evaluated. In research on language processing, the dynamics of this process have been termed situation model updating (Albrecht & O'Brien, 1993; O'Brien & Albrecht, 1992). The theoretical framework by Wyer and Radvansky (1999) is the first attempt to apply these ideas to the processing of social information. One promising avenue of future research would be to combine this framework with the notions developed here and examine the interplay of comprehension and validation in social information processing more directly.

The Adaptive Value of Routine Epistemic Validation Processes

What are the benefits of epistemic monitoring? Knowledge is fallible as a matter of principle and deception, non-intentional misinformation, and other instances of inaccurate information are ubiquitous phenomena. For these reasons, epistemic monitoring serves a vital function in everyday social life. In the face of deception and non-intentional forms of misinformation, epistemic monitoring is clearly necessary to construct accurate referential representations of state of affairs in the world. Accurate referential representations, in turn, are a prerequisite for successful individual actions and social interactions. Against this background, one may speculate that the human ability to validate information fast and efficiently has developed in the course of evolution along with other linguistic and communicative skills (Pinker, 1994). From a slightly different angle, fast and efficient epistemic monitoring processes may also be regarded as an aspect of a more general human capability to detect cheating on social contracts. In their research program on the evolutionary foundations of social interaction, Cosmides and Tooby (1989, 1992) have repeatedly demonstrated that people are able to detect violations of conditionals of the form *If P then Q* very accurately if these rules are framed as instances of generic social rules such as *If you pay the cost then you take the benefit*. This finding is remarkable because ordinary people usually have great difficulties in reasoning with the same conditionals when these are framed differently or presented as problems of formal logic. Cosmides and Tooby explain this discrepancy with the adaptive value of detecting acts of cheating that in the long run, might undermine the foundations of social life. Communicating false and inaccurate information also represents a violation of a generic social rule, the conversational maxim of quality (Grice, 1975). This is most obvious in deception and lying, which are universally considered to be immoral ways of behavior. From this perspective, the adaptive value of fast and efficient validation processes is twofold: They assist individuals in the construction of accurate referential representations that can be used for situated action and they contribute to the stability of social groups by detecting cheaters who violate the conversational maxim of quality.

How Epistemic Validation Assists Comprehension

How does epistemic validation assist the construction of a referential representation of the communicated information? At any rate, the detection of an inconsistency between incoming information and prior knowledge through epistemic monitoring enables individuals to reject knowledge-inconsistent information, thereby preventing it from becoming part of the situation model. Apart from this general characteristic, three major cases may be distinguished, depending on the nature of the detected inconsistency and the knowledge involved in resolving the inconsistency.

The first case pertains to situations when the available knowledge allows inferring a specific alternative to the rejected information. In this case, the specific alternative becomes part of the current situation model, and the asymmetry in the effort needed to process true and false statements is small. For example, most people are able to reject the statement *Fire trucks are green* fast and efficiently. At the same time, their world knowledge allows them to generate a referential representation of the alternative fact that fire trucks are red. Most likely, a referential representation of a red fire truck is constructed as soon as the first concept mentioned in the sentence is processed. This way of conceptualizing the processing of false or belief-inconsistent information is related to the fusion model of negation encoding (Mayo et al, 2004). According to the fusion model, comprehenders use their knowledge to integrate the negated proposition and the negation marker into one meaningful unit (e.g., the statement *Jim is not guilty* is encoded as a representation of the state of affairs that Jim is innocent).

In the second case, an inconsistency may be detected but neither the discourse context nor the available knowledge allows generating a specific alternative to the rejected information. For example, most people are able to determine that the statement *Apple trees bloom in December* is probably false but not many would be able to construct a precise situation model of an alternative state of affairs. In that case, comprehenders may construct a fuzzy referential representation, e.g., a situation model encoding the fact that apple trees bloom sometime in the period from spring to early summer, or a family of hypothetical models that encode all possible alternatives. Both possibilities increase the computational complexity and, as a consequence, the cognitive effort that is needed to comprehend the false information (Johnson-Laird, 1983). To give an example, experiments by Hovland and Weiss (1953) on conceptual learning from positive and negative instances demonstrated that learning from negative instances is more difficult than learning from positive instances, even if the amount of information transmitted by the two sets of instances is held constant. Thus, the second case entails an asymmetry in the processing of true (or belief-consistent) and false (or belief-inconsistent) information. However, this asymmetry is not due to a fundamental inability to reject false information.

In the third case, an inconsistency is detected but the available information does not permit the construction of a situation model, forcing comprehenders to retreat to a propositional representation in which the rejected information is marked as false (similar to the schema-plus-tag model of negation, Clark & Chase, 1972). Experiments by Hasson, Simmons, and Todorov (2005) illustrate this case and its differences to the case that entails the construction of a referential representation. In these experiments, participants learned information about a fictitious person, and the information was marked as being true or false. In one experiment, participants were able to reject information learned as false in a verification task if its purported falsehood allowed them to infer a specific alternative (e.g., *This person is a liberal – false*) even when they had to perform a secondary task during comprehension. Only if such an inference was not possible (e.g., *This person walks barefoot to work – false*), the additional cognitive load induced by the secondary task caused an affirmation bias, i.e. a higher proportion of errors in the verification task for assertions learned as false compared to those learned as true. Thus, when knowledge is available that allows inferring a specific alternative, individuals are able to reject false information reliably even if they are put under cognitive load. One possible explanation for this finding is that knowledge permits the construction of a referential representation of the true state of affairs that can later be used for verification judgments (e.g., Kaup & Zwaan, 2003). However, when neither available knowledge nor the pragmatic context allows inferring a specific alternative state of affairs, it is not possible to construct a referential model. In that case, epistemic monitoring processes lead to a tagged propositional representation that makes it more difficult to reject false information in later verification tasks (see Mayo et al., 2004, for analogous findings for negated statements). Notably, this case also applies to the stimuli used in the experiments in support of the dual-stage model (e.g., Gilbert et al., 1990, 1993).

Characteristics of Knowledge as Moderators of Epistemic Validation

It is important to note that despite the effortless and routine character of epistemic monitoring, inconsistencies of incoming information and prior knowledge are not always detected. Given that epistemic monitoring rests on memory-based processes, inconsistencies with knowledge that is available in principle but that is less salient may go unnoticed, with the consequence that belief-inconsistent information and even information that an individual knows to be false can find its way into the situation model. Thus, the *accessibility* of knowledge may be regarded as a crucial moderator of its use in epistemic monitoring, similar to the moderating role of attitude accessibility for the attitude-behavior relationship (Fazio, 1995). A related property is the *degree of integration* in larger knowledge structures (Naumann & Richter, 2000). A high degree of integration should enhance the likelihood that a particular piece of knowledge is passively cued by incoming information and automatically activated by memory-based processes. In addition, beliefs that are strongly integrated in a network of other supportive beliefs provide a broader basis for detecting and rejecting belief-inconsistent information (e.g.,Petty, Haugtvedt, & Smith, 1995). As a third aspect that has also been described in the literature on attitude strength (e.g., Gross, Holtz, & Miller, 1995), *belief certainty* may be assumed to moderate the relationship of knowledge and epistemic monitoring. Cognitive conflicts of knowledge with incoming information can arise only if this knowledge is held with a minimum of subjective certainty. Finally, the available knowledge must bear some *relevance* to the validity of incoming information, that is, it must have implications for the truth value of the incoming information. Singer's research program on causal inferences has established that comprehenders activate relevant information to validate implicit premises that bridge cause and effect (Singer et al., 1992; Singer, 1993). However, due to the memory-based character of epistemic monitoring processes, there are limits to the scope of these processes. Apart from explicitly communicated information, only inferences that are regularly drawn during comprehension and integrated into the situation model can be subject to epistemic monitoring.

Negative Effects of Routine Epistemic Validation Processes

Despite its essential role for achieving accurate referential representations, epistemic monitoring processes can also lead to the acceptance of false information. Biasing effects of epistemic monitoring are documented by various approaches in social and cognitive psychology that have examined effects of comprehension on validation processes. In studies based on the debriefing paradigm (Ross, Lepper, & Hubbard, 1975) and the related misinformation paradigm (Johnson & Seifert, 1994), information given in an earlier part of the experiment continued to influence participants' judgments and beliefs after it had been explicitly discredited or contradicted by information provided at a later point. In particular, the influence of false information persevered if participants embedded this information in a causal chain of events (Johnson & Seifert, 1994) or inferred causal explanations for the information which was discredited later (Anderson, Lepper, & Ross, 1980). The generation of knowledge-based explanations and inferences may even alter the subjective likelihood of events that are explicitly marked as hypothetical (Ross, Lepper, Strack, & Steinmetz, 1977). These findings go well with a situation model account of the relationship of comprehension and validation: When comprehenders are able to construct a referential representation that is consistent with their prior knowledge about an issue, the likelihood of accepting this information as true will be increased (Schroeder et al., in press). In this case, the referential representation itself serves as the knowledge base for epistemic monitoring processes. It is important to note that this explanation of affirmation bias differs from the mechanism of initial acceptance that is put forward by Gilbert's (1991) dual-stage model. Research by Fiedler and colleagues on biases in social judgments clearly elucidates the differences between both accounts (Fiedler et al., 1996; Fiedler, Walther, Armbruster, Fay & Naumann, 1996). In their experiments, participants' ratings of a target person's behavior were biased towards the contents of questions or assertions about this behavior presented earlier, even if participants had explicitly rejected this information. However, these effects occurred only when the propositional information contained in the questions matched participants' knowledge structures, and they occurred after participants had ample opportunity to think about the validity of this information. Thus, these effects are best described as constructive, knowledge-based biases that result from an interplay of comprehension and validation.

References

Abelson, R.P., & Prentice, D.A. (1997). Contrast tests of interaction hypotheses. *Psychological Methods, 2,* 315-328.

Albrecht, J. E. & O'Brien, E. J. (1993). Updating a mental model: Maintaining both local and global coherence. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 19, 1061-1070.

Appel, M., & Richter, T. (2007). Persuasive effects of fictional narratives increase over time. *Media Psychology, 10*, 113-134.

Appel, M., & Richter, T. (2008). The role of transportation and need for affect in persuasion through fictional narratives: A mediated moderation model. *Manuscript submitted for publication*.

Anderson, C.A., Lepper, M.R., & Ross, L.R. (1980). Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology, 39*, 1037-1049.

Bless, H., Bohner, G., Schwarz, N., & Strack, F. (1990). Mood and persuasion: A cognitive response analysis. *Personality and Social Psychology Bulletin, 16*, 331-345.

Carpenter, P.A., & Just, M.A. (1975). Sentence processing: A psycholinguistic processing model of verification. *Psychological Review, 82*, 45-73.

Chen, S., & Chaiken, S. (1999). The heuristic-systematic model in its broader context. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 73-96). New York: Guilford.

Chen, Y.W., & Blanchard-Fields, F. (2000). Unwanted thought: Age differences in the correction of social judgments. *Psychology and Aging, 15*, 475-482.

Clark, H.H., & Chase, W.G. (1972). On the process of comparing sentences against pictures. *Cognitive Psychology, 3*, 472-517.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.

Cosmides, L., & Tooby, J. (1992). Cognitive adaptions for social exchange. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind* (pp. 163-228). New York: Oxford University Press.

Davidson, D. (2001). *Inquiries into truth and meaning* (2nd ed.). Oxford: Oxford University Press.

Descartes, R. (1644/2006). *Selections from the principles of philosophy* (translated by J. Veitch). Retrieved March 7, 2006, from http://www.gutenberg.org/etext/4391

Dodd, D.H., & Bradshaw, J.M. (1980). Leading questions and memory: Pragmatic constraints. *Journal of Verbal Learning & Verbal Behavior, 19*, 695-704.

Dummett, M. (1981). *Frege: Philosophy of language* (2nd ed.). Cambridge, MA: Harvard University Press.

Erdfelder, E., Faul, F., & Buchner, A. (1996). GPOWER: A general power analysis program. *Behavior Research Methods, Instruments, & Computers, 28*, 1-11.

Fazio, R.H. (1995). Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R.E. Petty & J.R. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 247-282). Mawah, NJ: Erlbaum.

Fiedler, K., Armbruster, T., Nickel, S., Walther, E., & Asbeck, J. (1996). Constructive biases in social judgment: Experiments on the self-verification of question contents. *Journal of Personality and Social Psychology, 71*, 861-873.

Fiedler, K., Walther, E., Armbruster, T., Fay, D., & Naumann, U. (1996). Do you really know what you have seen? Intrusion errors and presupposition effects on constructive memory.

Journal of Experimental Social Psychology, 32, 484-511.

Gilbert, D.T. (1991). How mental systems believe. *American Psychologist, 46*, 107-119.

Gilbert, D.T. (1997). Ordinary personology. In D.T. Gilbert, S.T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., pp. 89-150). New York: McGraw Hill.

Gilbert, D.T., & Malone, P.S. (1995). The correspondence bias: The what, when, how and why of unwarranted dispositional inference. *Psychological Bulletin, 117*, 21-38.

Gilbert, D.T., Krull, D.S., & Malone, P.S. (1990). Unbelieving the unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology, 59*, 601-613.

Gilbert, D.T., Tafarodi, R.W., & Malone, P.S. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology, 65*, 221-233.

Green, M.C., & Brock, T.C. (2002). In the mind's eye: Transportation-imagery model of narrative persuasion. In M.C. Green, J.J. Strange, & T.C. Brock (Eds.), *Narrative impact: Social and cognitive foundations* (pp. 315-342). Mahwah, NJ: Erlbaum.

Greene, S.B., McKoon, G. & Ratcliff, R. (1992). Pronoun resolution and discourse models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 266-283.

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

Gross, S.R., Holtz, R., & Miller, N. (1995). Attitude certainty. In R.E. Petty & J.A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 215-246). Mahwah: Erlbaum.

Gueraud, S., & O'Brien, E.J.(2005). Components of comprehension: A convergence between memory-based processes and explanation-based processes. *Discourse Processes, 39*, 123-124.

Halldorson, M., & Singer, M. (2002). Inference processes: Integrating relevant knowledge and text information. *Discourse Processes, 33*, 145-162.

Hasson, U., Simmons, J.P., & Todorov, A. (2005). Believe it or not: On the possibility of suspending belief. *Psychological Science, 16*, 566-571.

Hovland, C.I., & Weiss, W. (1953). Transmission of information concerning concepts through positive and negative instances. *Journal of Experimental Psychology, 45*, 175-182.

Jacoby, L.L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513-541.

Johnson, H.M., & Seifert, C.M. (1994). Sources of the continued influence effect: When misinformation in memory affects later inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 1420-1436.

Johnson-Laird, P.N. (1983). *Mental models: Towards a cognitive science of language, inferences, and consciousness*. Cambridge: Cambridge University Press.

Jones, E.E. (1979). The rocky road from acts to dispositions. *American Psychologist, 34*, 107-117.

Jones, M.N., Kintsch, W., & Mewhort, D.J.K. (2006). High-dimensional semantic space accounts of priming. *Journal of Memory and Language, 55*, 534-552.

Judd, C.M., & McClelland, G.H. (1989). *Data analysis: A model comparison approach.* New York: Harcourt Brace Jovanovich.

Kaup, B., & Zwaan, R.A. (2003). Effects of negation and situational presence on the accessibility of text information. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 439-446.

Kaup, B., Zwaan, R.A., & Lüdtke, J. (in press). The experiential view of language comprehension: How is negated text information represented? In F. Schmalhofer & C.A. Perfetti (Eds.), *Higher level language processes in the brain: Inference and comprehension processes*.

Mahwah, NJ: Erlbaum.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review, 95*, 163-182.

Knowles, E.S., & Condon, C.A. (1999). Why people say ‘yes’: A dual-process theory of acquiescence. *Journal of Personality and Social Psychology, 77*, 379-386.

Kruglanski, A.W., & Semin, G. (2007). The epistemic bases of interpersonal communication. In M. Hewstone, Schut, H.A.W., de Wit, J.B.F., van den Bos, K., & Stroebe, M. (Eds.), *The scope of social psychology: Theory and applications* (pp. 107-120). New York: Psychology Press.

Landauer, T.K., & Dumais, S.T. (1997). A solution to Plato’s problem: The Latent Semantic Analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 104*, 211-240.

Lea, R.B. (1995). On-line evidence for elaborative logical inferences in text. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 18-26.

Lea, R.B., Mulligan, E.J., & Walton, J.L. (2005). Accessing distant premise information: How memory feeds reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 387-395.

MacDonald, M.C., & Just, M.A. (1989). Changes in activation levels with negation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 633-642.

Mayo, R., Schul, Y., & Burnstein, E. (2004). "I am not guilty" vs "I am innocent": Successful negation may depend on the schema used for its encoding. *Journal of Experimental Social Psychology, 40*, 433-449.

McGuire, W.J. (1964). Inducing resistance to persuasion: Some contemporary approaches. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 1, pp. 191-

229). New York: Academic Press.

McGuire, W.J., & Papageorgis, D. (1961). The relative efficacy of various types of prior belief-defense in producing immunity against persuasion. *Journal of Abnormal & Social Psychology, 62*, 327-337.

McKoon, G., & Ratcliff, R. (1995). The minimalist hypothesis: Directions for research. In C.A. Weaver, S. Mannes, & C.R. Fletcher (Eds.), *Discourse comprehension: Essays in honor of Walter Kintsch* (pp. 97-116). Hillsdale, NJ. Erlbaum.

McLeod, C.M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin, 109*, 163-203.

Montague, R. (1973). The proper treatment of quantification in ordinary English. In K. Hintikka, J. Moravcsik, & P. Suppes (Eds.), *Approaches to natural language* (pp. 221-242). Dordrecht, The Netherlands: Riedel.

Myers, J.L., & O'Brien, E.J. (1998). Accessing the discourse representation during reading. *Discourse Processes, 26*, 131-157.

Naumann, J. & Richter, T. (2000). Response times in attitudinal items as indicators of the continuous accessibility of knowledge related to attitudes. *Psychological Reports, 87*, 355-366.

Nezlek, J.B. (2001). Multilevel random coefficient analyses of event and interval contingent data in social and personality psychology research. *Personality and Social Psychology Bulletin, 27*, 771-785.

O’Brien, E.J., & Albrecht, J.E. (1992). Comprehension strategies in the development of a mental model. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 777-784.

Partee, B. (1975). Montague grammar and transformational grammar. *Linguistic Inquiry, 6*, 203-300.

Petty, R.E., & Wegener, D.T. (1999). The elaboration likelihood model: Current status and controversies. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 41-72). New York: Guilford.

Petty, R.E., Haugtvedt, C.P., & Smith, S.M. (1995). Elaboration as a determinant of attitude strength. In R.E. Petty & J.A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 93-130). Mawah, NJ: Erlbaum.

Pinker, S. (1994). *The language instinct*. London: Penguin.

Potter, M.C. (1984). Rapid serial visual presentation (RSVP): A method for studying language processing. In D.E. Kieras & M.A. Just (Eds.), *New methods in reading comprehension research* (pp. 91-118) . Hillsdale, NJ: Erlbaum.

Prentice, D.A., Gerrig, R.J., & Bailis, D.S. (1997). What readers bring to the processing of fictional texts. *Psychonomic Bulletin & Review, 4*, 416-420.

Quené, H. & van den Bergh, H. (in press). Examples of mixed-effects modeling with crossed random effects and with binomial data. *Journal of Memory and Language*.

Raudenbush, S.W., & Bryk, A.S. (2002). *Hierarchical linear models: Applications and data analysis methods* (2nd ed.). Newbury Park, CA: Sage.

Raudenbush, S.W., Bryk, A.S., Cheong, Y.F., & Congdon, R.T. (2004). *HLM 6: Hierarchical linear and nonlinear modeling*. Chicago, IL: Scientific Software International.

Richter, T. (2006). What is wrong with ANOVA and multiple regression? Analyzing sentence reading times with hierarchical linear models. *Discourse Processes, 41*, 221-226.

Ross, L., Lepper, M.R., Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. *Journal of Personality and Social Psychology, 32*, 880-892.

Ross, L., Lepper, M.R., Strack, F., & Steinmetz, J. (1977). Social explanation and social expectation: Effects of real and hypothetical explanations on subjective likelihood. *Journal of Personality and Social Psychology, 32*, 880-829.

Schroeder, S., Richter, T., & Hoever, I. (in press). Getting a picture that is both accurate and stable: Situation models and epistemic validation. *Journal of Memory and Language*.

Schul, Y., Mayo, R., & Burnstein, E. (2004). Encoding under trust and distrust: The spontaneous activation of incongruent cognitions. *Journal of Personality and Social Psychology, 86*, 668-679.

Singer, M. (1993). Causal bridging inferences: Validating consistent and inconsistent sequences. *Canadian Journal of Experimental Psychology, 47*, 340-359.

Singer, M. (2006). Verification of text ideas during reading. *Journal of Memory and Language, 54*, 574-591.

Singer, M., & Halldorson, M. (1996). Constructing and validating motive bridging inferences. *Cognitive Psychology, 30*, 1-38.

Singer, M., Halldorson, M., Lear, J. C., & Andrusiak, P. (1992). Validation of causal bridging inferences. *Journal of Memory and Language, 31*, 507-524.

Skurnik, I., Yoon, C., Park, D.C., & Schwarz, N. (2005). How warnings about false claims become recommendations. *Journal of Consumer Research, 31*, 713-724.

Spinoza, B. (1667/1997). *Ethics demonstrated in geometric order* (translated by R.H.M. Elwes). Retrieved March 7, 2006, from http://www.mtsu.edu/~rbombard/RB/Spinoza/ethica-front.html

Stroop, J.R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18*, 643-662.

Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review, 93*, 239-257.

Trope, Y., & Gaunt, R. (2000). Processing alternative explanations of behavior: Correction or integration? *Journal of Personality and Social Psychology, 79*, 344-354.

Trope, Y., Cohen, O., & Maoz, Y. (1988). The perceptual and inferential effects of situational inducements on dispositional attributions. *Journal of Personality and Social Psychology, 55*, 165-177.

van Dijk, T.A., & Kintsch, W. (1983). *Strategies of discourse comprehension*. New York: Academic Press.

Wegener, D.T., Petty, R.E., Smoak, N.D., & Fabrigar, L.R. (2004). Multiple routes to resisting attitude change. In E.S. Knowles & J.A. Linn (Eds.), *Resistance and persuasion* (pp. 13-38). Mawah, NJ: Erlbaum.

Wyer, R.S., & Radvansky, G.A. (1999). The comprehension and validation of social information. *Psychological Review, 106*, 89-118.

Zuwerink, J., & Devine, P.G. (1996). Attitude importance and resistance to persuasion: It's not just the thought that counts. *Journal of Personality and Social Psychology, 52*, 677-688.

Zwaan, R.A., & Radvansky, G.A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin, 123*, 162-185.

Appendix: Contrasts Tested in Experiments 1 and 2

Table A1:

Contrast 1 (Experiment 1) Representing the Null Hypothesis that Interference Affects Validation of True and False Assertions Associated with Weak Background Beliefs to the Same Degree.

	Validity					
	True assertions			False assertions		
	Learning condition			Learning condition		
Background Beliefs	No interference	Interference	No learning	No interference	Interference	No learning
Strong	0	0	0	0	0	0
Weak	-1	1	0	1	-1	0

Note. Contrast weights did not differ between the levels of the independent variable Response time window. For ease of presentation, this independent variable is omitted from the table.

Table A2:

Contrast 2 (Experiment 1) Representing the Null Hypothesis that Interference Affects Validation of True and False Assertions Associated with Strong Background Beliefs to the Same Degree.

	Validity					
	True assertions			False assertions		
	Learning condition			Learning condition		
Background Beliefs	No interference	Interference	No learning	No interference	Interference	No learning
Strong	-1	1	0	1	-1	0
Weak	0	0	0	0	0	0

Note. Contrast weights did not differ between the levels of the independent variable Response time window. For ease of presentation, this independent variable is omitted from the table.

Table A3:

Contrast 1 (Experiment 2) Representing the Null Hypothesis that the Effect of Interference during Learning Does not Differ Between True and False Assertions (Regardless of Processing Goal)

	Validity					
	True assertions			False assertions		
	Learning condition			Learning condition		
Processing goal	No interference	Interference	No learning	No interference	Interference	No learning
Recognition	-1	1	0	1	-1	0
Validation	-1	1	0	1	-1	0

Note. Contrast weights did not differ between the levels of the independent variable Response time window. For ease of presentation, this independent variable is omitted from the table.

Table A4:

Contrast 2 (Experiment 2) Representing the Null Hypothesis that There is No Three-Way Interaction of Learning with vs. without Interference, Validity and Processing Goal

	Validity					
	True assertions			False assertions		
	Learning condition			Learning condition		
Processing goal	No interference	Interference	No learning	No interference	Interference	No learning
Recognition	-1	1	0	1	-1	0
Validation	1	-1	0	-1	1	0

Note. Contrast weights did not differ between the levels of the independent variable Response time window. For ease of presentation, this independent variable is omitted from the table.

Footnotes

[1] Overall, true assertions were verified more accurately (M=.67, SE_M=.02) than false assertions (M=.59, SE_M=.02), $F(1,35)=23.9$, $p<.001$, $\eta^2=.41$. Assertions associated with strong background beliefs (M=.80, SE_M=.02) were verified more accurately than those associated with weak background beliefs (M=.46, SE_M=.02), $F(1,35)=875.4$, $p<.001$, $\eta^2=.96$. Assertions that participants had learned without interference (M=.72, SE_M=.02) were verified more accurately than those they had learned with interference (M=.69, SE_M=.02), while assertions learned with interference in turn were verified more accurately than those that had not been presented in the learning phase (M=.48, SE_M=.02), $F(2,34)=94.9$, $p<.001$, $\eta^2=.85$. Finally, assertions were verified more accurately within the long response time-frame (M=.74, SE_M=.01) than in the medium time-frame (M=.65, SE_M=.02) and the short time-frame (M=.50, SE_M=.03), $F(2,34)=66.1$, $p<.001$, $\eta^2=.80$.

[2] In addition to the reported effects, the ANOVA on the verification data of Experiment 2 revealed a number of other significant effects. Crucially, none of these additional results affects the interpretation of the theoretically relevant effects. First, as in Experiment 1, the main effects of validity and response-time frame turned out to be significant. True assertions were verified more accurately (M=.84, SE_M=.02) than false assertions (M=.74, SE_M=.02), $F(1,36)=43.7$, $p<.001$, $\eta^2=.55$. An overall higher proportion of correct responses was given within the long response time-frame (M=.86, SE_M=.02) compared to the medium time-frame (M=.83, SE_M=.02) and the short time-frame (M=.68, SE_M=.03), $F(2,35)=30.6$, $p<.001$, $\eta^2=.64$. Second, there was a significant interaction effect involving processing goal, learning condition, and response time-frame, $F(4,33) = 2.8$, $p < .05$, $\eta^2=.25$. Regardless of their processing goal, participants were able to verify assertions that they had seen in the learning phase reliably within the long response

time-frame (memorization goal: M=.87, SE_M=.03; validation goal: M=.89, SE_M=.03) and within the medium time-frame (memorization goal: M=.87, SE_M=.03; validation goal: M=.84, SE_M=.03). When they had to respond within the short response time-frame, the group given the memorization goal had a slightly lower proportion of correct responses (M=.70, SE_M=.04) than the group given the validation instruction (M=.75, SE_M=.04). For assertions that had not been presented in the learning phase, in contrast, the proportion of correct responses declined steadily in both processing goal conditions from the long time-frame (memorization goal: M=.83, SE_M=.04; validation goal: M=.84, SE_M=.03) over the medium time-frame (memorization goal: M=.79, SE_M=.04; validation goal: M=.78, SE_M=.04) to the short time-frame (memorization goal: M=.63, SE_M=.05; validation goal: M=.53, SE_M=.05).

Figure Captions

Figure 1. Effects of validity and learning condition in assertions associated with strong background beliefs (a) and weak background beliefs (b) on the proportion of correct verification responses (Experiment 1).

Figure 2. Effects of validity and learning condition for assertions learned with a verification goal (a) and a recognition goal (b) on the proportion of correct verification responses (Experiment 2).

Figure 3. Effects of validity and required response in the orthographical judgments task in Experiment 3 (presentation rate 1 word/800 ms) on response latencies (a) and error rates (b).

Figure 4. Effects of validity and required response in the orthographical judgments task in Experiment 4 (presentation rates 1 word/600 ms and 1 word/300 ms) on response latencies (a) and error rates (b).

a)

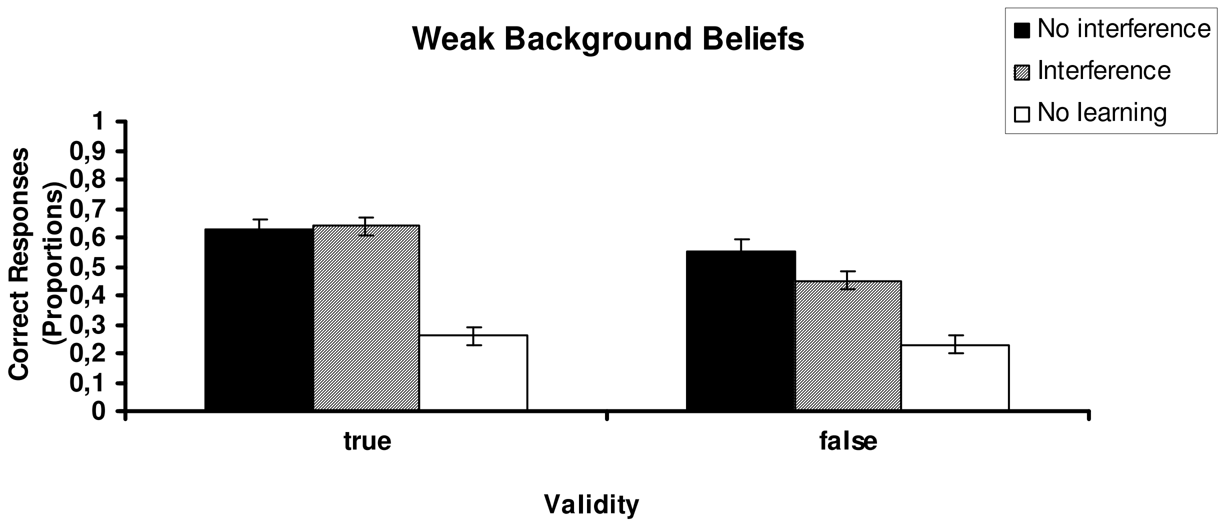
Weak Background Beliefs
No interference
Interference
No learning
Correct Responses (Proportions)
1
0,9
0,8
0,7
0,6
0,5
0,4
0,3
0,2
0,1
0
true
false
Validity

b)

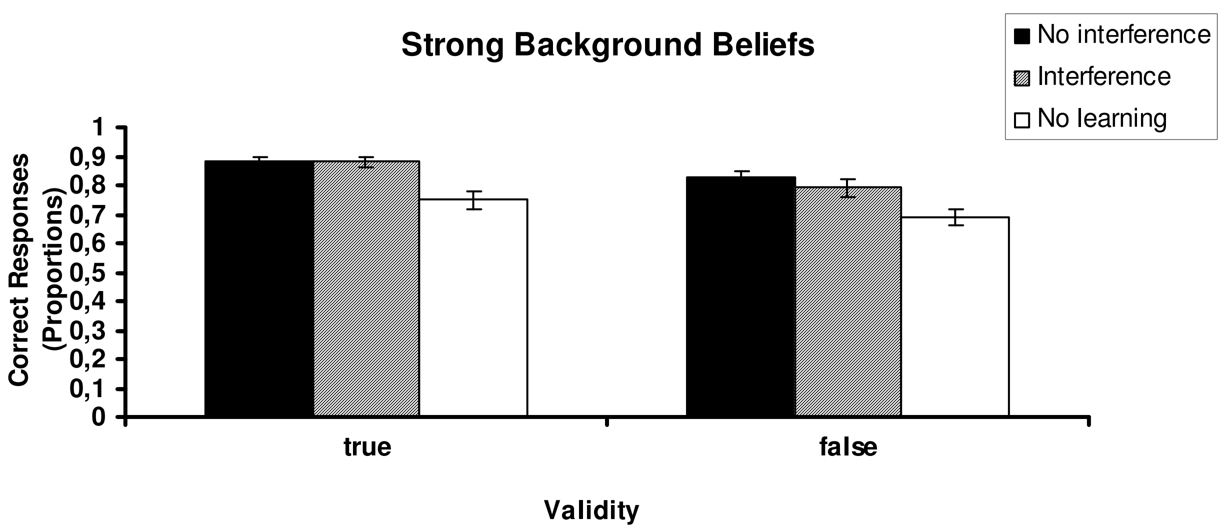
Strong Background Beliefs
No interference
Interference
No learning
Correct Responses (Proportions)
1
0,9
0,8
0,7
0,6
0,5
0,4
0,3
0,2
0,1
0
true
false
Validity

a)

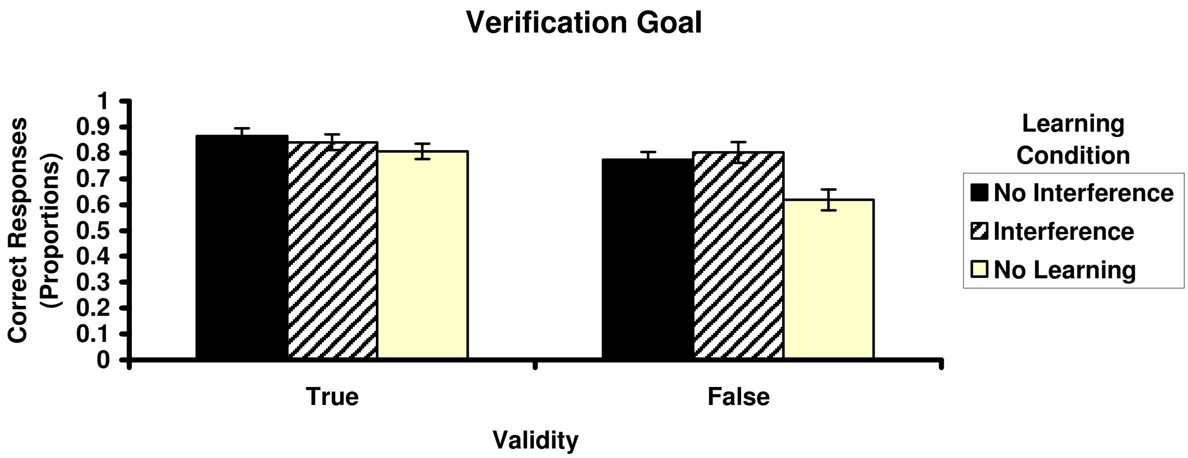
Verification Goal
Correct Responses (Proportions)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
True
False
Validity
Learning Condition
No Interference
Interference
No Learning

b)

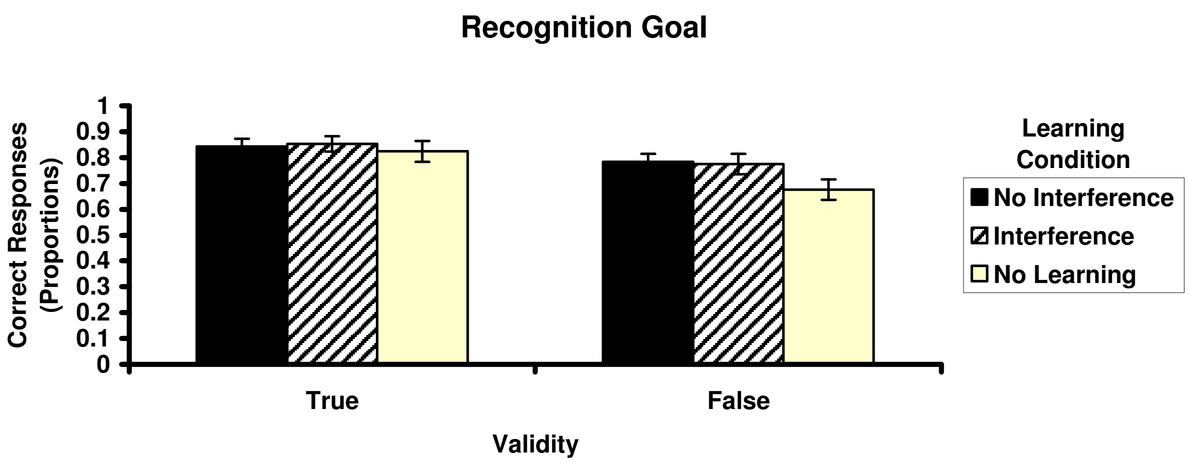
Recognition Goal
Correct Responses (Proportions)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
True
False
Validity
Learning Condition
No Interference
Interference
No Learning

a)

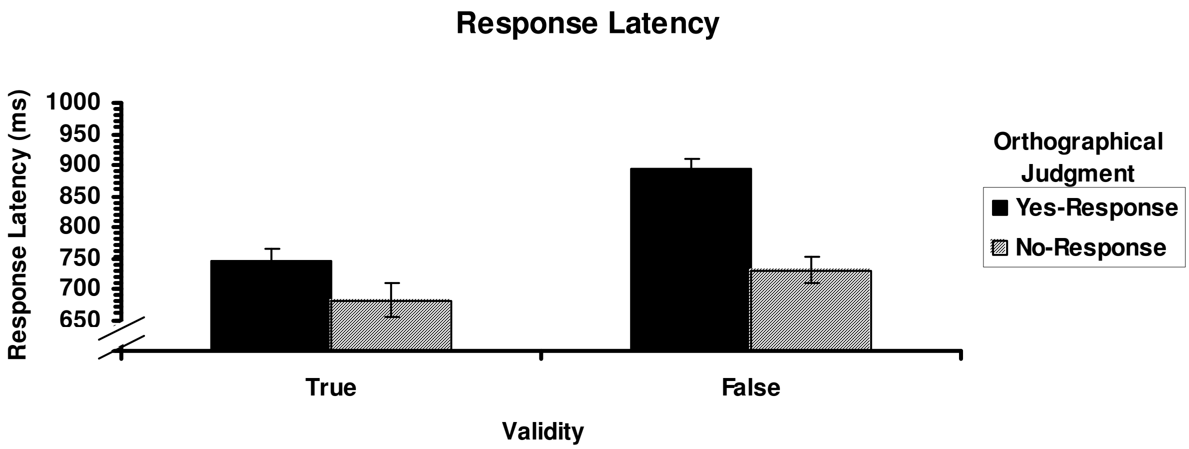
Response Latency
Response Latency (ms)
1000
950
900
850
800
750
700
650
True
False
Validity
Orthographical Judgment
Yes-Response
No-Response

b)

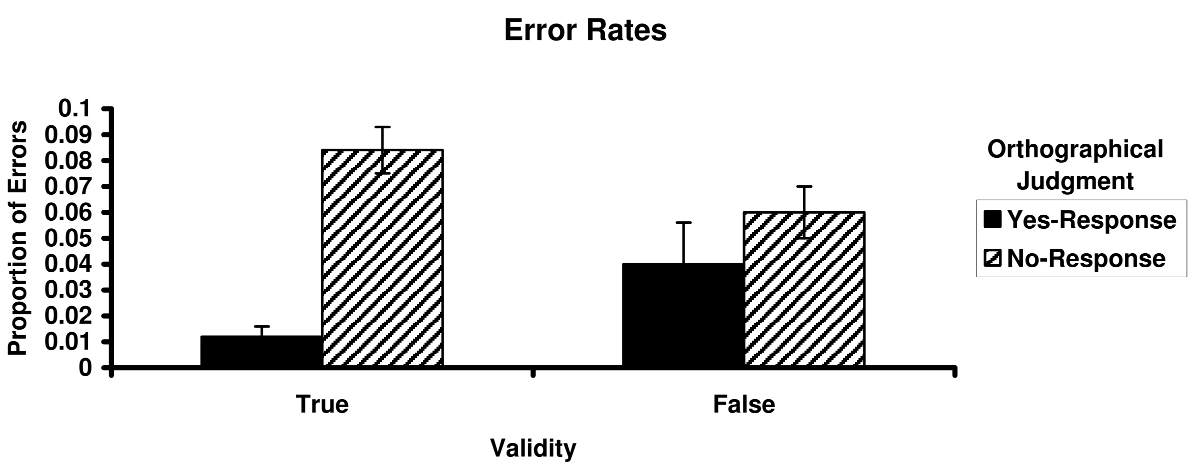
Error Rates
Proportion of Errors
0.1
0.09
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0.01
0
True
False
Validity
Orthographical Judgment
Yes-Response
No-Response

a)

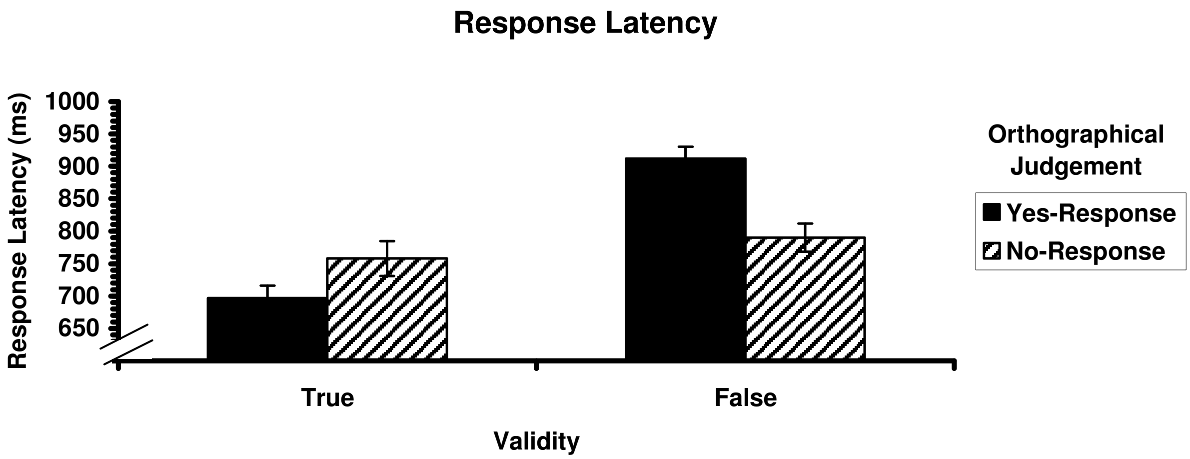
Response Latency
Response Latency (ms)
1000
950
900
850
800
750
700
650
True
False
Validity
Orthographical Judgement
Yes-Response
No-Response

b)


Error Rates
Proportion of Errors
0.1
0.09
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0.01
0
True
False
Validity
Orthographical Judgment
Yes-Response
No-Response