

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

How Knowledge Interferes with Reasoning - Suppression Effects by Content and Context

Permalink

<https://escholarship.org/uc/item/18z7k2pr>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 21(0)

Authors

Neth, Hansjorg

Beller, Sieghard

Publication Date

1999

Peer reviewed

How Knowledge Interferes with Reasoning – Suppression Effects by Content and Context

Hansjörg Neth (neth@psychologie.uni-freiburg.de)
Sieghard Beller (beller@psychologie.uni-freiburg.de)
Department of Psychology; University of Freiburg
79085 Freiburg, Germany

Abstract

The suppression of logically valid inferences by the content or context of premises can be seen as an instance of knowledge having a detrimental influence on reasoning. Although Henle (1962) has claimed that invalid deductions are due to additional premises drawn from background knowledge, current research on content effects ignores the methodological implications of this claim. Elaborating on the suppression effect in conditional reasoning (Byrne, 1989), we present a knowledge-based approach that makes relevant features of background knowledge an integral part of the analysis. After identifying the sufficiency and necessity of conditions as the type of knowledge mediating the effect, we construct and validate task materials independently from any assessment of reasoning (Experiment 1). We then replicate and extend suppression effects in syllogism tasks (Experiment 2) and show that participants are able to couch their background knowledge in formally correct wordings (Experiment 3).

Suppose you were presented with the following premises:

- (1) If Ann is in New York, she visits the Guggenheim.
- (2) If it is still open, she visits the Guggenheim.
- (3) Ann is in New York.

Would you conclude that Ann visits the Guggenheim? If not, you fell prey to the *suppression effect* (Byrne, 1989) by not drawing the valid inference of Modus Ponens, warranted by premises (1) and (3). The demonstration that the valid inferences of Modus Ponens (MP) and Modus Tollens (MT) can be suppressed by an additional premise (2) was used against an earlier result that additional premises can also prevent the fallacious inferences ‘Affirmation of the Consequent’ (AC) and ‘Denial of the Antecedent’ (DA), thereby *facilitating* logical performance (Rumain, Connell, & Braine, 1983).

The ambiguous nature of this phenomenon provoked a vigorous debate about contemporary theories of human reasoning. Proponents of a mental models theory (Byrne, 1991; Johnson-Laird & Byrne, 1991) focus on semantic procedures of integrating formally identical premises and argue that contextual information facilitates the search for counterexamples to putative conclusions (Byrne, Espino, & Santamaria, 1998). Theorists who conceptualize the mind as equipped with tacit rules of inference claim that premise (2) leads participants to question the truth of premise (1), and to combine the two antecedent propositions in a conjunctive way

(Politzer & Braine, 1991). This conjunction yields “If Ann is in New York *and* it is still open, she visits the Guggenheim” and blocks the MP in question. As both theories refer to interpretive processes that are not within their scope, they suffer from the same shortcoming (Fillenbaum, 1993): neither can explain *why* premises of one and the same syntactic form are interpreted differently based on their content. To understand this process, it is necessary to embed the study of deductive reasoning in a knowledge-based approach (Chan & Chua, 1994; Beller, 1997; Beller & Spada, 1998).

A Knowledge-Based Approach

To introduce our point of view, let us analyze another version of the above example. Based on the premises

- (1) If Ann is in New York, she visits the Guggenheim.
- (4) Ann is *not* in New York.

many certainly would conclude that Ann does *not* visit the Guggenheim. But this conclusion seems logically unjustified, for concluding ‘not *q*’ from ‘if *p* then *q*’ and ‘not *p*’ means to commit the fallacy of DA. Yet, again an extra premise might be to blame for this apparent error in reasoning. For if someone *knows* that the Guggenheim is actually located in New York, he or she might introduce this spatial information as

- (5) If Ann visits the Guggenheim, she is in New York.

Based on this implicit premise it follows by MT that Ann does not visit the Guggenheim if she is not in New York. But even this explanation depends on background knowledge. For believing (5) means to be negligent of the fact that there is another Guggenheim museum in Bilbao, Spain.

The relation between knowledge and reasoning has puzzled philosophers ever since Plato. Experimental evidence indicating that additional information can suppress logical thinking seems to refute the naïve equation “the more knowledge the better the reasoning.” Henle (1962) sought to re-establish universal deductive competence by claiming that all inferences are perfectly valid but occasionally based on additional or distorted premises. Yet dissolving the issue of faulty reasoning by mere reference to interpretive processes introduces a second black box to account for the first. Even if

additional or misinterpreted premises may account for a variety of apparent errors, the claim that eventually all fallacies can be reduced to this mechanism could still be premature.

A crucial point in the debate on content effects is the dependence of logical validity on the premises *as understood by the reasoner*. But despite its appeal, the methodological consequences of this idea have widely been ignored. To account for content effects, at least four levels of analysis have to be distinguished: (a) the linguistic framing of the premises, (b) their formal properties as defined by the semantics of logic, (c) the inferential properties of background knowledge that is triggered by the premises' content, and finally, (d) the mental integration of all this in the reasoner's mind.

Within our knowledge-based framework we focus on the complementary levels (b) and (c): Separate formalizations of form and content yield two predictive models whose combination predicts the *facilitation* of formally correct inferences whenever the conclusions warranted by both models coincide. Likewise, *suppression* of formally correct inferences is expected to occur whenever both predictions diverge (as it was the case in both examples above). However, when form and content predict inconsistent conclusions, evaluation of logical validity does not depend on logic, but on the preferred point of view.

Methodologically, the relativity of conclusions to different predictive models suggests experimental research on content effects to proceed in a two step strategy:

1. When background knowledge is claimed to have effects on reasoning, it must not be consulted only *post hoc* to justify otherwise invalid answers. To avoid circular argumentation (Smedslund, 1970), relevant inferential features of content have to be specified in advance and tested independently from any judgment about the validity of inferences.

2. Once assumptions about the relevant type of background knowledge have been secured, its influence on reasoning can be addressed, e.g., by assessing the extent to which changes in content or context invite formally invalid inferences.

General Design of the Experiments

To apply this framework to account for effects of suppression and facilitation in conditional reasoning, we follow Thompson (1994) by taking knowledge about the *sufficiency* and *necessity* of conditions as key variables to capture the relevant features of content. Combining both dimensions results in four *patterns of dependence* between a condition p and a consequence q :

- S+N+: p is both sufficient and necessary for q .
- S+N-: p is sufficient but not necessary for q .
- S-N+: p is not sufficient but necessary for q .
- S-N-: p is neither sufficient nor necessary for q .

Sometimes it is only known that a condition p correlates with a consequence q but the precise nature of the relation is unknown. This can be represented as a fifth pattern (S?N?). Furthermore, several conditions may act additively or alternatively. Additive conditions can be understood as conjunctions of single factors (e.g. $p \wedge p_{add}$), and alternative conditions as disjunctions (e.g. $p \vee p_{alt}$). For instance, in the

Table 1: An example of a S-N+ scenario: According to general social knowledge p is necessary but insufficient for q .

primary condition p :	She has enough money.
consequent q :	She buys herself a dress.
alternative condition p_{alt} :	She has a credit card.
additional condition p_{add} :	The shops are open.

example at the very beginning, the two conditions ("Ann is in New York" and "It is still open") may be interpreted as conjunctive factors for her visit to the Guggenheim.

Having described how relevant background knowledge can be conceptualized, we were able to develop concrete task scenarios. The core of each scenario consists of a primary condition p and a consequence q , whose relation is one of the five patterns of dependence defined above. For each scenario we introduced an *alternative* condition p_{alt} and an *additional* condition p_{add} . As sufficiency and necessity are abstract concepts we did not expect effects due to highly domain-specific reasoning schemas (Cheng & Holyoak, 1985). But to allow for their possibility, each pattern of dependence was instantiated in three knowledge domains (*causal*, *social*, and *conceptual* relations). Combining both factors resulted in fifteen scenarios of which Table 1 presents an example.

In line with the methodological program outlined in the previous section, our first experiment assures that participants actually perceive our scenarios' presumed patterns of dependence as intended. In the following two experiments the influence of background knowledge on two different tasks is analyzed: Experiment 2 shows that suppression and facilitation in conditional syllogisms depends systematically on the premises' content and context. Experiment 3 demonstrates peoples' linguistic competence to frame their knowledge in formally correct wordings.

Experiment 1: Rating Patterns of Dependence

The goal of this experiment was to validate our materials and to test a prediction about the mechanism of context effects. If different contexts trigger different aspects of background knowledge, the introduction of an extra condition should systematically alter the perceived sufficiency and necessity of the primary condition p : Mentioning an alternative condition p_{alt} , should specifically reduce the perceived necessity \hat{n} of p , but not its perceived sufficiency \hat{s} . Likewise, we expect an additive condition p_{add} to reduce \hat{s} , but not \hat{n} .

Method

To assess perceived sufficiency \hat{s} and necessity \hat{n} we devised a *rating task*. Despite encouraging evidence by Thompson (1995), a pilot study showed that participants' conceptions of these notions varied considerably. The natural language usage of "sufficiency" in particular seemed susceptible to deviations from its logical semantics. In logic, the presence of a sufficient condition *necessitates* its consequence; in everyday-contexts, however, possession of a \$100 bill might be considered to be "sufficient" to buy a beer, yet does not necessitate its purchase. To avoid these linguistic issues, we

asked for expected contingencies: If p is sufficient for q , it is *not possible* that p is the case without q being the case as well; therefore, whenever one knows that p , one would expect q . Similarly, when p is necessary for q , p *must* be the case whenever q is the case. Thus, assessing contingencies is logically equivalent to asking about sufficiency and necessity, but can be done in a more manageable format.

Materials To create a situation similar to the syllogism tasks to be used in Experiment 2, we wanted participants to notice conditional relations without actually asserting them. Therefore, each scenario introduced a conditional as the statement of an alter ego: "Suppose someone stated that 'if p , then q .'" We then pointed out that the actual truth of the conditional was irrelevant to the task, which was to rate – based on background knowledge – the degree of confidence that q is the case, provided that p , and vice versa. Both ratings were performed on a scale from 0–100%. Context was introduced by adding a statement like "No information about p_{alt} is available." Thus, an extra condition was mentioned, but the rating task nonetheless addressed the relation between p and q .

Procedure Thirty volunteers were recruited on the campus of University of Freiburg and randomly assigned to one of three groups. Each participant rated all 15 scenarios in a randomized order. Whereas the first group received only scenarios without context (only p/q), in the other groups an extra condition was always present (either p_{alt} or p_{add} , counterbalanced across groups).

Results and Discussion

As the manipulation of specific knowledge domains yielded no significant results it will be excluded from this and all subsequent analyses.

Figure 1 shows the mean ratings on both dimensions in the absence of any contextual condition. Overall ratings of perceived sufficiency \hat{s} were significantly higher for S+ scenarios (with conditions that were meant to be sufficient) than for those of type S– (91.8% > 21.3%; Wilcoxon signed ranks: $z=6.7$, $n=60$, $p<.001$). Likewise, perceived necessity \hat{n} was larger for scenarios of type N+ than of type N– (93.3% > 27.0%; $z=6.6$, $n=60$, $p<.001$).

Similar analyses of perceived sufficiencies and necessities within each single pattern of dependence indicated a deviation from our predictions only in one case: A significantly higher rating of \hat{n} than of \hat{s} within scenarios of type S–N– suggested that they do not correspond as close to our intentions as the others (34.3% > 21.3%; $z=3.16$, $n=30$, $p=0.002$). But as both means still were within the negative range, this slight deviation is acceptable.

The overall effect of "context" can be seen in Table 2, aggregated over all five patterns of dependence. Whereas there was no difference in perceived sufficiency and necessity without extra conditions (53.2% \approx 58.3%; Wilcoxon: $z=.97$, $n=150$, $p=0.33$), the introduction of an alternative condition p_{alt} selectively reduced perceived necessity \hat{n} (58.3% > 46.5%; Mann-Whitney-U: $z=3.15$, $n=150$, $p=0.002$), but not perceived sufficiency \hat{s} (53.2% \approx 53.7%; U: $z=0.05$, $n=150$, $p=0.96$). Mentioning an additional condition p_{add} instead had just the opposite effect of selectively




Figure 1: The mean perceived sufficiency and necessity of p for the five patterns of dependence in absence of any contextual condition. The ellipses extend one standard deviation in either dimension. Each mean represents 60 ratings.

reducing \hat{s} (53.2% > 40.4%; U: $z=3.42$, $n=150$, $p=0.001$) but not \hat{n} (58.3% \approx 53.6%; U: $z=0.98$, $n=150$, $p=0.33$). As a whole, the pattern illustrated by Table 2 is a rating equivalent of Byrne's (1989) original suppression effect.

Both goals of Experiment 1 could be achieved: Participants' knowledge led them to perceive the scenarios as theoretically predicted, and the perceived patterns of dependence were selectively modulated by the introduction of extra conditions. This validates our materials and will be the basis to interpret the results of the following experiments.

Experiment 2: Drawing Conditional Inferences

Correct conditional reasoning, that is, drawing the valid inferences of MP and MT while abstaining from the invalid ones of AC and DA, is one of the hallmarks to assess human logical competence. In our second experiment, we combined the methodologies employed by Byrne (1989) and Thompson (1994, 1995) to replicate and generalize their effects and explain them within our knowledge-based framework. With regard to the two predictive models (based on form vs. con-

Table 2: The effects of context (p_{alt}/p_{add}) on the ratings of sufficiency and necessity aggregated over patterns of dependence. (Standard deviations in brackets.) Each mean represents 150 ratings on a 0–100 scale.

Rated dimension	Context		
	only p/q	+ p_{alt}	+ p_{add}
Sufficiency \hat{s}	53.2 (37.7)	53.7 (40.5)	40.4 (33.7)
Necessity \hat{n}	58.3 (37.1)	46.5 (33.1)	53.6 (37.4)




Figure 2: The percentages of inferences based on a single conditional ‘If p then q ’ without an extra conditional as its context. High white and low black bars signify high percentages of formally correct inferences. Every bar represents the mean percentage over 30 trials.

tent) we predicted an increase in formally justified inferences when the predictions of both models coincide, and a corresponding decrease in the case of diverging predictions. As we understand this effect to be mediated by background-knowledge, we expect that it can be experimentally evoked in the same fashion by variations of content and context.

Method

Materials The same fifteen scenarios validated in Experiment 1 were now used to develop conditional syllogisms. Each item consisted of a primary conditional ‘If p then q ’, a categorical premise (e.g. ‘ p ’ for MP) and a choice of three possible conclusions (e.g. ‘ q ’, ‘not q ’ and ‘nothing follows’). For each scenario all four conditional inferences were used (MP, MT, AC, and DA). In the context conditions, an extra conditional with an alternative or additive antecedent (‘If p_{alt}/p_{add} then q ’) was inserted after the primary one.

As the propositions p and q of the primary conditional still exemplified each of the five patterns of dependence, a note on possibly “odd” premises is in place here: According to textbook logic, relations between a necessary and insufficient condition p and a consequence q (S-N+) have to be formalized as ‘ $q \rightarrow p$ ’ (e.g. “If Jack goes skiing, there’s snow” – at least as long as alpine skiing is meant). In a natural language context, however, the inversion “If there’s snow, Jack goes skiing” does not seem peculiar at all. On the contrary: Whereas the first sentence only states a notorious fact of common knowledge, its inversion conveys new information about Jack.

Procedure 50 volunteers from the University of Freiburg participated in the experiment. Each completed two booklets. The first contained a random sequence of all 24 conditional syllogisms that can be constructed from the primary conditional of the six scenarios of two different patterns of dependence.

After finishing the first booklet, the volunteers received another 24 items made up by the six scenarios of two other patterns of dependence. Now each syllogism was presented in the context of an extra conditional, half of which had alternative, and half of which had additive antecedents. Thus, every participant took part in a “simple” and a “context” condition and responded to total of 12 scenarios from four different patterns of dependence with four different categorical premises. Position of patterns of dependence was counter-balanced across five groups, to which participants were randomly assigned.

Results and Discussion

The percentages of different conditional inferences with only a single conditional premise, ‘If p then q ’, are represented in Figure 2.

As scenarios of type S?N? yield no knowledge-based prediction we used the inferences drawn under this condition as a baseline to assess the influence of background knowledge by other patterns of dependence. The last bar chart in Figure 2 shows that this baseline consists of a standard result: MP is endorsed nearly universally and more often than MT, and both valid inferences are drawn more often than the two fallacies of AC and DA.

Whenever it follows from background knowledge that a condition p is sufficient for a consequence q to occur (S+), this content prediction coincides with the formal interpretation of the conditional ‘ $p \rightarrow q$ ’, so that an overall increase of MP and MT inferences should result. Vice versa, if the content model tags a condition p as being insufficient for a consequence q (S-), we expected a decrease. However, only the first prediction was supported by our data: Whereas the mean frequency of MP and MT inferences was significantly higher for S+ scenarios than for our reference S?N? ($91.7\% > 81.7\%$; $\chi^2=4.01$, $df=1$, $p < .05$), they were not significantly reduced by S- scenarios ($77.5\% \approx 81.7\%$; $\chi^2=.70$, $df=1$, n.s.).

Complementarily to the influence of perceived sufficiency on MP and MT, perceived necessity should influence the inferences of AC and DA. If, as in the sample scenario shown in Table 1, a condition p is known to be necessary for a consequence q (N+), the content-based prediction contradicts the formal properties of a conditional ‘ $p \rightarrow q$ ’, which defines p as being sufficient for q . As the content corresponds to a reversed formalization ‘ $q \rightarrow p$ ’, we expected an increase of inferences based on categorical premises ‘ q ’ and ‘not p ’, i.e., AC and DA with respect to ‘ $p \rightarrow q$ ’. Similarly, if condition p is known to be unnecessary (N-) a corresponding decrease should result. Again, only the first prediction was supported: The mean frequency of AC and DA inferences of both N+ scenarios was significantly higher than the corresponding frequency of scenario S?N? ($92.5\% > 40.0\%$; $\chi^2=68.91$, $df=1$, $p < .01$), but scenarios of type S- did not yield the opposite effect ($37.5\% \approx 40.0\%$; $\chi^2=.16$, $df=1$, n.s.).

These analyses show a twofold result: Whenever the content predicts a definite conclusion (as in the S+ and N+ cases) this conclusion is likely to be drawn, whether it coincides with the formal prediction (as in the case of MP and MT) or not (as in the case of AC and DA). However, a content

Table 3: The percentages of inferences (out of 150) in three different contexts, aggregated over patterns of dependence.

Context	MP	MT	AC	DA
only p/q	92.7	75.3	58.0	62.0
$+p_{alt}$	92.7	78.0	16.0	14.0
$+p_{add}$	64.0	46.7	50.0	47.3

model that predicts the indefinite answer “nothing follows” does not significantly affect reasoning.

Effects of *context*, i.e., the introduction of an extra conditional premise with an alternative or additive antecedent, are depicted in Table 3. Compared to the frequencies in the condition without an extra conditional (only p/q), the predicted effects of facilitation and suppression were found: The frequency of AC and DA inferences was significantly reduced by the introduction of an alternative condition p_{alt} ($15.0\% < 60.0\%$; $\chi^2=253.13$, $df=1$, $p < .01$), and also by an additive condition p_{add} ($48.7\% < 60.0\%$; $\chi^2=16.06$, $df=1$, $p < .01$), though not to the same extent. Vice versa, the mean frequency of MP and MT inferences was significantly reduced in the additive context $+p_{add}$ ($55.3\% < 84.0\%$; $\chi^2=183.43$, $df=1$, $p < .01$), but not in the alternative context $+p_{alt}$ ($85.3\% \approx 84.0\%$; $\chi^2=.40$, $df=1$, n.s.).

To sum up, the main predictions of this experiment were empirically supported: Through variations of content and context we were able to replicate and combine the results of Byrne (1989) and Thompson (1994, 1995), i.e., we found effects of suppression and facilitation in conditional syllogisms. By theoretically accounting for them within our general framework, we demonstrated that performance in conditional reasoning is modulated systematically by background knowledge.

Experiment 3: Selecting Appropriate Wordings

Psychological assessment of logical competence traditionally presupposes participants’ proper understanding of semi-formal premises, as used in verbal tasks of conditional reasoning. On the other hand, fallacious responses are frequently attributed to misinterpreted logical notions, as, for example, the confusion of conditional and biconditional relations (see Evans, Newstead, & Byrne, 1993, for an overview.) In a third experiment we shifted our focus from processes of language comprehension and reasoning to the selection of logically adequate formulations. If participants’ linguistic competence indeed includes a basic mastery of the conditional connective, we would expect them to select different wordings for different patterns of dependence. Moreover, their selections should be altered systematically by the introduction of extra conditions p_{alt} and p_{add} .

Method

Materials To limit the range of possible answers we preferred a *multiple choice* format to a free formulation task. For each scenario, we first presented two conditionals, which, as in Experiment 1, were introduced by an alter ego: “When asked for the proper relation between p and q , someone proposed ‘If p then q ’, and someone else ‘If q then p ’”. We em-

phasized that both, either one, or none of these propositions could express the actual relationship between p and q , and prompted participants to indicate which of these four possibilities was most appropriate according to their background knowledge. Due to the construction of the task, the four options corresponded to a biconditional ($p \leftrightarrow q$), conditional ($p \rightarrow q$ vs. $q \rightarrow p$), or non-conditional ($p \sim q$) interpretation. Manipulations of context were accomplished by mentioning an extra condition, as in “No information about p_{alt} is available.”

Procedure Identical to Experiment 1.

Results and Discussion

As indicated in Table 4, participants were very sensitive to variations in patterns of dependence. Overall, 68.2% of all responses matched the predicted category. This percentage was highest for context ‘only p/q ’ (79.2%), and lower for contexts ‘ $+p_{alt}$ ’ (64.0%) and ‘ $+p_{add}$ ’ (61.3%).

The correspondence between participants’ choices and the predictions is best illustrated by an example. For the sample scenario of type S–N+ presented in Table 1, the two proposed conditionals were “If she has enough money, she buys herself a dress” and “If she buys herself a dress, she has enough money.” Because having money (p) is a necessary but insufficient prerequisite for buying a dress (q), only the second sentence expresses the general relation between p and q appropriately ($q \rightarrow p$). 28 out of 30 participants had this intuition.

What should happen if extra conditions are mentioned? Stating that it is unknown whether or not “she has a credit card” provides participants with an *alternative* condition p_{alt} . The formally appropriate representation of this new situation ‘ $q \rightarrow p \vee p_{alt}$ ’ does not correspond to either conditional formulation. Therefore, if this new information is taken into account, the non-conditional category ‘ $p \sim q$ ’ ought to be

Table 4: The number of selected conditional relations by context and pattern of dependence. Each row represents a total of 30 choices. Predicted choices are **bold-faced**.

Context	Pattern of dependence	Selected category			
		$'p \rightarrow q'$	$'q \rightarrow p'$	$'p \leftrightarrow q'$	$'p \sim q'$
only p/q	S+N–	29**	0	1	0
	S–N+	0	28**	1	1
	S+N+	1	4	25**	0
	S–N– ^a	1	4	1	23**
	S?N?	5	7	5	13*
$+p_{alt}$	S+N–	28**	0	0	2
	S–N+	4	13*	2	11
	S+N+	15**	3	9	3
	S–N–	1	2	2	25**
	S?N?	3	3	7	17**
$+p_{add}$	S+N–	15**	2	1	12
	S–N+	1	27**	0	2
	S+N+	2	12	10	6
	S–N–	1	5	2	22**
	S?N?	2	8	1	19**

*: $P(X \geq 13, 30, .25) < .05$. **: $P(X \geq 14, 30, .25) < .01$.

^a: one missing value.

selected. As Table 4 ('+ p_{alt} ; S-N+') shows, 11 people followed this prediction. Vice versa, mere mentioning of the additional condition p_{add} "The shops are open" should not induce the same shift in responses, as the resulting formal representation ' $q \rightarrow p \wedge p_{add}$ ' still implies ' $q \rightarrow p$ ', and 27 participants indeed selected the latter category.

In general, participants selected formally appropriate wordings for the basic patterns of dependence, and tended to choose the predicted categories when extra conditions were introduced. Such specific shifts of preferred verbal expressions can be seen as another instance of the suppression effect. The fact that many participants selected the original option even when a change in context seemed to necessitate a shift suggests that some subjects chose to ignore the additional information.

Conclusions

The results of all three experiments back up our general approach to content effects that distinguishes between inferences based on the form of premises and inferences based on their semantic content. First, we confirmed that our experimental scenarios were perceived as theoretically intended. Second, both types of syllogism tasks (with and without an extra conditional premise) yielded the predicted effects of facilitation and suppression, which shows that reasoning performance depends systematically on the specific background knowledge triggered. Finally, a third experiment demonstrated that participants were able to select appropriate wordings for conditional relations although the syllogism tasks seemed to compromise their logical competence.

What do our results mean for the theoretical debate? Our analysis could be adapted to either a mental models framework or a mental rule theory. Presently, our knowledge representation is more similar to the mental models point of view, but a computer implementation uses formal rules to derive inferences from these representations (see Beller, 1997, for details). Thus, our account cannot resolve the debate between mental logic and mental models theories, but by addressing interpretive processes that are ubiquitous in human reasoning it fills a gap that has been neglected by both.

Suppression effects of content and context are multifaceted, knowledge-based phenomena. In order to explain their underlying mechanisms, syntactic and semantic accounts must be integrated rather than pitted against each other. A knowledge-based approach addresses questions of content without having to assert content-specific rules of inference or being committed exclusively to one theoretical framework. It adds predictive power to general theories of reasoning and helps to clarify the complex interplay of knowledge, reasoning, language, and logic.

Acknowledgements

We thank Hans Spada, Michael Scheuermann, and Josef Nerb for many helpful discussions on the topic of this paper. We are also grateful to Phil Johnson-Laird, Zachary Estes, Mary Newsome, and Yingrui Yang for their encouraging criticisms and suggestions.

References

- Beller, S. (1997). *Inhaltseffekte beim logischen Denken – Der Fall der Wason'schen Wahlaufgabe*. [Content effects in deductive reasoning: The case of Wason's selection task.]. Lengerich: Pabst Science Publishers.
- Beller, S., & Spada, H. (1998). Conditional reasoning with a point of view: the logic of perspective change. In M. Gernsbacher & S. Derry (Eds.), *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 138–143). Mahwah, NJ: Lawrence Erlbaum.
- Byrne, R. M. J. (1989). Suppressing valid inferences with conditionals. *Cognition*, 31, 61–83.
- Byrne, R. M. J. (1991). Can valid inferences be suppressed? *Cognition*, 39, 71–78.
- Byrne, R. M. J., Espino, O., & Santamaria, C. (1998). Context can suppress inferences. In A. C. Quelhas & F. Pereira (Eds.), *Cognition and context* (pp. 201–214). Lisboa: Instituto Superior de Psicologia Aplicada.
- Chan, D., & Chua, F. (1994). Suppression of valid inferences: Syntactic views, mental models, and relative salience. *Cognition*, 53, 217–238.
- Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, 17(4), 391–416.
- Evans, J. St. B. T., Newstead, S. E., & Byrne, R. M. J. (1993). *Human reasoning: The psychology of deduction*. Hove: Lawrence Erlbaum Associates.
- Fillenbaum, S. (1993). Deductive reasoning: What are taken to be the premises and how are they interpreted? *The Behavioral and Brain Sciences*, 16(2), 348–349.
- Henle, M. (1962). On the relation between logic and thinking. *Psychological Review*, 69(4), 366–378.
- Johnson-Laird, P. N., & Byrne, R. M. J. (1991). *Deduction*. Hove, England: Lawrence Erlbaum Associates.
- Politzer, G., & Braine, M. D. (1991). Responses to inconsistent premisses cannot count as suppression of valid inferences. *Cognition*, 38, 103–108.
- Rumain, B., Connell, J., & Braine, M. D. (1983). Conversational comprehension processes are responsible for reasoning fallacies in children as well as adults: *If* is not the biconditional. *Developmental Psychology*, 19(4), 471–481.
- Smedslund, J. (1970). On the circular relation between logic and understanding. *Scandinavian Journal of Psychology*, 11, 217–219.
- Thompson, V. A. (1994). Interpretational factors in conditional reasoning. *Memory and Cognition*, 22(6), 742–758.
- Thompson, V. A. (1995). Conditional reasoning: The necessary and sufficient conditions. *Canadian Journal of Experimental Psychology*, 49(1), 1–60.