EUROPEAN JOURNAL OF COGNITIVE PSYCHOLOGY, 2003, *15* (2), 161–176

Causal conditional reasoning and strength of association: The disabling condition case

Wim De Neys, Walter Schaeken, and Géry d'Ydewalle

K.U. Leuven, Department of Psychology, Belgium

Cummins (1995) has shown that reasoning with conditionals involving causal content is affected by the relative number of available alternative and disabling conditions. More recent evidence (Quinn & Markovits, 1998) indicates that, beside the number of stored conditions, the relative strength of association of the alternative conditions with the consequent term is another important factor that affects causal conditional reasoning. In this study we examined the effect of the strength of association for the disabling conditions. We identified causal conditionals for which there exists only one highly associated disabler. With these conditionals we constructed conditional inference problems in which the minor premise was expanded with the negation of a strongly or weakly associated disabler. Results of two experiments indicate that strength of association of stored disabling conditions is affecting reasoning performance: Acceptance of Modus Ponens and Modus Tollens increased when there was no strongly associated disabler available.

Cognitive psychologists studying human reasoning have devoted a great deal of research to conditional reasoning. This kind of reasoning consists in making inferences on the basis of ''if p then q'' sentences. In a standard conditional inference task people are asked to assess arguments of the following four kinds:

Modus Ponens (MP)	If p then q, p therefore q
Modus Tollens (MT)	If p then q, not q therefore not p
Denial of the antecedent (DA)	If p then q, not p therefore not q
Affirmation of the consequent (AC)	If p then q, q therefore p

Requests for reprints should be addressed to W. De Neys, K.U. Leuven, Dept. of Psychology, Tiensestraat 102, B-3000 Leuven, Belgium. Email: Wim.Deneys@psy.kuleuven.ac.be

Preparation of the manuscript was supported by grants from the Fund for Scientific Research-Flanders (FWO) and the Belgian Program on Interuniversity Poles of Attraction, Convention number P4/19. We thank Nick Chater, Mike Oaksford, and Jonathan Evans for their thoughtful comments on drafts of this paper. Parts of this study were presented at the 23rd annual conference of the Cognitive Science Society, Edinburgh, UK, 2001.

http://www.tandf.co.uk/journals/pp/09541446.html DOI:10.1080/09541440244000058

Under the material implication interpretation of propositional logic, MP and MT are considered valid inferences, whereas DA and AC are regarded as fallacies. Much of the work on conditional reasoning has tried to identify the factors that influence performance on these four problems (for a review, see Evans, Newstead, & Byrne, 1993).

A growing body of evidence is showing that people's knowledge about the relation between the p (antecedent) and q (consequent) part of the conditional has a considerable effect on the underlying reasoning process (e.g., see Byrne, 1989; Byrne, Espino, & Santamaria, 1999; Markovits, 1984; Newstead, Ellis, Evans, & Dennis, 1997; Rumain, Connell, & Braine, 1983; Thompson, 1994).

In the case of reasoning with conditionals involving causal content (e.g., ''If cause p, than effect q'') seminal work has been done by Cummins and her colleagues (1995; Cummins, Lubart, Alksnis, & Rist, 1991). Following Byrne (1989), Cummins examined the effect of the alternative and disabling conditions of a causal conditional. An alternative condition is a possible cause that can produce the effect mentioned in the conditional, whereas a disabling condition prevents the effect from occurring despite the presence of the cause. Consider the following conditional:

If the brake is depressed, then the car slows down.

Possible alternative conditions for this conditional are:

running out of gas, having a flat tyre, shifting the gear down . . .

The occurrence of these conditions will result in the car slowing down. The alternatives make it clear that it is not necessary to depress the brake in order to slow the car down. Other causes are also possible. Possible disabling conditions are:

a broken brake, accelerating at the same time, skid due to road conditions . . .

If such disablers are present, depressing the brake will not result in the slowing down of the car. The disablers make it clear that it is not sufficient to depress the brake in order to slow down the car. There are additional conditions that have to be fulfilled.

When people (fallaciously) accept DA and AC inferences, they fail to see that there are other causes that may lead to the occurrence of the effect beside the original stated one. Cummins (1995) and Cummins et al. (1991) found that people's acceptance of DA and AC inferences decreased for conditionals with a high number of possible alternative conditions. This showed that a crucial factor in making the fallacious inferences is the number of alternative causes people can think of. In addition, she found that the number of disabling conditions affected the acceptance of the valid MP and MT inferences: If there were many conditions that could disable the relation between antecedent and consequent, people also tended to reject the valid inferences.

Quinn and Markovits (1998) have identified another factor that may influence reasoning with causal conditionals. They showed that not only the number of alternative conditions is important, but also what they call the ‘‘strength of association’’ of the alternative conditions. Quinn and Markovits developed a framework (see also Markovits, Fleury, Quinn, & Venet, 1998) where reasoning performance is being linked to the structure of semantic memory. In this framework it is assumed that, when confronted with a causal ‘‘if p then q’’ conditional, reasoners will access a causal structure in semantic memory that corresponds to ‘‘ways of making q happen’’ (i.e., alternative conditions). Within the structure, there will be causes that will be more strongly associated with q than others. The more strongly associated a specific cause is, the higher the probability that it will be retrieved by the semantic search process.

Quinn and Markovits (1998) measured strength of association by frequency of generation: In a pretest, participants were asked to write down as many potential causes for a certain causal consequent (effect, e.g., ‘‘a dog scratches constantly’’). This allowed the construction of conditionals with a strongly (e.g., ‘‘If a dog has fleas, then it will scratch constantly’’) and weakly (e.g., ‘‘If a dog has skin disease, then it will scratch constantly’’) associated cause. Because the consequent is the same in both conditionals, the number of possible alternative conditions is kept constant.

Quinn and Markovits (1998) reasoned that, with both the ‘‘strong’’ and ‘‘weak’’ type of conditional, people would try to activate and retrieve ‘‘alternative ways of making q happen’’. However, reasoners given the weak conditional will be able to activate the strongly associated cause, whereas for the strong conditional they will have to activate some other, less closely associated term. Thus, it will be more difficult to retrieve an alternative condition in case of the strong conditional, which would lead to a greater acceptance of DA and AC inferences. The results of the study were consistent with the predicted response pattern.

The identification of the strength of association effect raises the question whether this effect is also present for the disabling conditions. Indeed, although knowledge of disabling conditions is also stored in semantic memory, Quinn and Markovits (1998) restricted their case to an analysis of the alternative conditions. Cummins (1995) already showed that both the number of alternatives and disablers is affecting reasoning performance. In addition, Elio (1998) has shown that the process of disabler retrieval is not only important in conditional reasoning but also in the field of belief revision and non-monotonic reasoning: Belief in a conditional after contradiction was lower when people could find many disablers. Thus, both for reasoning psychologists and the psychological and AI community studying belief revision, examining the effect of associative strength of disablers can identify a new factor affecting the crucial disabler retrieval. The present study focused on this topic.

The framework developed by Quinn and Markovits (1998) was adopted and extended to the disabling conditions. It was assumed that when presented a causal conditional, people will not only access a causal structure with alternative conditions but also one that corresponds to ''ways that prevent q to occur'' (see Markovits, 2000; Vadeboncoeur & Markovits, 1999). When such disabling conditions are retrieved, p will no longer be perceived as a sufficient condition for q what renders the MP and MT conclusions uncertain.

In a generation task we identified strongly and weakly associated disablers for a number of conditionals. We constructed experimental items by expanding the original antecedents of the conditionals with the negation of the strongly or weakly associated disabler. Suppose that for a certain conditional we find that S is a strongly associated disabler, whereas W is a weak one. This allows the construction of the expanded conditionals: ''If P and not S, then Q'' (strongly expanded conditional) and ''If P and not W, then Q'' (weakly expanded conditional). These expanded conditionals have an equal number of possible disablers (i.e., the original number minus one). However, reasoners presented ''If P and not W, then Q'' will still be able to activate the strongly associated disabler S, whereas with ''If P and not S, then Q'' they will have to activate a less closely associated one. Thus, it will be harder to access and retrieve disablers for the strong conditionals. This access-to-disablers manipulation rests solely on the strength of association of the disablers and not on the number of accessible disablers.

Retrieving disablers from semantic memory will decrease the acceptance of MP and MT inferences. Therefore, we predict that acceptance ratings for MP and MT inferences will be higher for the strongly expanded conditionals than for the weakly ones.

In the present experiment we did not manipulate the access to alternative conditions. The classical findings of Cummins (1995) indicate that retrieving disablers has no effect on DA and AC. Given these findings, one might expect that the access-to-disablers manipulation will have no effect on DA and AC acceptance.

EXPERIMENT 1

Pretest

The material for the present experiment was selected from previous pilot work (see De Neys, Schaeken, & d'Ydewalle, 2000), where 20 participants wrote down as many disabling conditions as possible for a set of 20 causal conditionals (with 1.5 min generation time for each conditional). The set included the 16 conditionals from Cummins (1995, Exp. 1) and four additional ones. The generation protocols were scored by two independent raters in order to identify unrealistic items and items that were simple variations of a single idea (e.g., for the previous example ''skid due to water on road'', ''skid due to mud'', ''skid due to ice on road''). Fewer than 5% of the generated disablers were disallowed by the raters (interrater reliability was .83).

For every conditional we established the relative frequency of appearance of the disablers that participants wrote down. We needed conditionals with a set of disablers in which there was one specific disabler that was very frequently generated. The expanded conditionals manipulation also forced us to take an additional criterion into account. We could not allow disablers that express a quantification of the original antecedent (e.g., ''brake not depressed hard enough''). Expanding the original with this kind of disablers would result in inconsistencies for some problems (e.g., DA, ''The brake was not depressed and the brake was depressed hard enough''). We selected three conditionals that met these criteria. From each set of disablers one infrequently generated disabler was selected. This weakly associated disabler had to meet the non-quantification criterion. Furthermore, if the strongly expanded conditional contained an explicit negation (e.g., ''If the apples are ripe and they are not picked''), we opted to express the selected weakly associated disabler in an explicit negated way too. The negation criterion should guarantee that the strongly and weakly expanded conditionals have comparable lexical complexity. Finally, the selected disablers had to sound as natural (according to our intuitions) as possible (e.g., ''not too little wind'' was not accepted). Table 1 presents the material that was selected for the experiment.

TABLE 1
Relative frequency of generation of the most frequently mentioned disablers for the three selected conditionals

If the apples are ripe, then they fall from the tree
Picked (65%)
Too little wind (25%)
Not enough weight (20%)
Not ripe enough (20%)
Apples caught in branches (10%)
If John grasps the glass with his bare hands, then his fingerprints are on it
Hands not greasy (50%)
Grasped glass with palms only (35%)
Prints wiped off (30%)
Glass was wet (25%)
If water is heated to 100°C, then it boils
Not pure water (75%)
Not normal pressure (30%)
Bad temperature measure (30%)

The disablers are given in order of frequency (%). Selected strongly and weakly associated disablers are italicised.

One could remark that our pretest allowed 90 s generation time (as in Cummins, 1995), whereas Quinn and Markovits (1998) allowed only 30 s. It could be argued that this might confound the strength of association classification. However, De Neys et al. (2000) also used a 30 s generation task for some of the conditionals. When generation frequency for the disablers in the 1.5 min and 30 s generation task was compared, results indicated that although the shorter generation time decreased the absolute frequency of generation, the crucial relative ranking of the disablers was not affected.[1]

We also note that the pilot study supplemented Quinn and Markovits (1998) by addressing some possible reservations about the use of generation frequency as strength of association measure. We looked at the relation between generation frequency and other possible associative strength measures such as plausibility and generation order. After the generation task we asked participants to judge the plausibility of the generated disablers. We also calculated the probability that a certain disabler was generated as the first one for a specific conditional. It was shown that the frequency of generation measure correlated with the other strength of association measures: More frequently generated disablers were judged more plausible and tended to be generated first.[2]

Method

Participants. Eighty-nine first-year university students participated in the experiment.

Material. Participants received a four-page booklet. Page 1 included the instructions for the task. On the top of each of the next three pages appeared the selected conditionals. Each conditional was embedded in the four inference types (MP, DA, MT, AC). So, each of the three pages included one conditional with four inference problems. For each conditional there was a specific presentation order of the four inferences (AC, MT, DA, MP or MP, MT, DA, AC or MP, DA, MT, AC). The three pages were bound into booklets in randomised order. Below each inference problem appeared a 7-point rating scale. This resulted in the item format shown in Figure 1. This is an example of the MP problem. On the same page participants would also find the MT, DA, and AC

[1] Generation frequency of the disablers generated for eight different conditionals could be compared. The Spearman rank order correlation reached .84, $t(82) = 13.71$, $p < .001$.

[2] For every disabler generated for a specific conditional we calculated the mean plausibility rating and the overall (# generated first/20) and relative (# generated first/# generated) probability that this disabler was the first generated one. Frequency of generation was associated with plausibility, Spearman rank order test, $R_s = .36$, $t(162) = 4.83$, $p < .00001$; $R_s = .50$, $t(35) = 3.41$, $p < .005$ when only disablers generated by at least 50% of participants were considered, and overall, $R_s = .63$, $t(162) = 10.22$, $p < .001$, and relative, $R_s = .40$, $t(162) = 5.52$, $p < .001$, probability that the disabler is generated first.

Rule: If water is heated to 100°C, then it boils

Fact: The water is heated to 100°C and the water is pure
Conclusion: The water boils

1	2	3	4	5	6	7
Very sure	Sure	Somewhat sure		Somewhat sure	Sure	Very sure

That I CANNOT draw this conclusion — That I CAN draw this conclusion

Figure 1.

problem. The access to disablers manipulation consisted in the presentation of two different minor premises (the information under the heading ''Fact''); the example in Figure 1 would belong to the strongly associated group where the original information was expanded with the negation of the strongest associated disabler. Similarly, in the weakly associated group, the negation of the selected weakly associated disabler was added to the ''Fact'' information. In both expanded groups appeared the original conditionals on top of the item pages. Thus, participants were not presented explicit expanded conditionals but rather conditional inference problems with expanded minor premises. All the items in a single booklet belonged to the same group. Table 2 gives an overview of the different material in the two groups (for an MP problem).

Procedure. The booklets were randomly given out to students who agreed to participate in the experiment. No time limits were imposed. The instructions presented on the first page of the booklet were read aloud by the experimenter. The instructions explained the specific item format of the task. Participants were told that the task was to decide whether or not they could accept the different

TABLE 2
Different content in the groups of Experiment 1

Expanded strongly associated:
(a) Water is heated to 100°C and the water is pure
(b) The apples are ripe and they are not picked
(c) John grasps the glass with his bare hands and his hands are greasy

Expanded weakly associated:
(a) Water is heated to 100°C and the pressure is normal
(b) The apples are ripe and they are not caught in the branches
(c) John grasps the glass with his bare hands and the glass is dry

Both groups only differ by the kind of information that is presented in the minor premise. This is an example of the minor premises for the MP problem.

conclusions. The instruction page showed an example problem (always standard MP) together with a copy of the rating scale. Care was taken to make sure that participants understood the precise nature of the rating scale. As in Cummins (1995), participants were NOT specifically instructed to accept the premises as always true and to endorse only conclusions that follow necessarily. With Cummins we assume that this encourages people to reason as they would in everyday circumstances. However, one should note that strictly speaking the task is therefore not a deductive inference task (see Evans, 2000). Thus, accepting the MP/MT and AC/DA inferences should not be considered correct or incorrect reasoning. When we refer to the standard nomenclature, a nominalist stance is adopted towards the use of the terms ''valid inferences'' and ''fallacies''.

Results and discussion

Participants rated each of the four inference types three times. For every inference type the mean of these three ratings was calculated. This resulted in a 4 (inference type, within-subjects) $\times$ 2 (group, between-subjects) design. All hypotheses were tested with planned comparison tests and rejection probability of .05.

Table 3 shows the overall mean acceptance ratings for the four inference types in the expanded weakly and strongly associated group. Acceptance ratings in both groups differed significantly, $F(1, 87) = 4.55$, $MSe = 3.85$, $p < .04$. As expected, we obtained higher MT ratings in the strongly associated group, $F(1, 87) = 4.99$, $MSe = 2.67$, $p < .03$, where the strongly associated disabler was eliminated. A similar tendency was observed for the MP inference, but the effect did not reach significance.

TABLE 3
Mean acceptance rating for the four inference types in the strongly associated and weakly associated groups

	Group	
Inference type	*Expanded weakly associated (n = 45)*	*Expanded strongly associated (n = 44)*
MP	5.70	5.92
DA	4.78	5.11
MT	4.37*	5.14*
AC	4.98	5.44

* Planned contrast $p < .05$.

The rating scale ranged from 1 (*very sure cannot draw this conclusion*) to 7 (*very sure can draw this conclusion*) with 4 representing *cannot tell*.

In line with Cummins' findings both expanded groups did not significantly differ in terms of AC and DA acceptance. However, we neither observed a significant interaction between the different inferences types. All inferences tended to be accepted more when no strongly associated disabler was available. Thus, a possible impact of the strength of association of stored disablers on the fallacies can not be discarded.

The MT findings do provide some preliminary support for the hypothesis that retrieving disablers from semantic memory is affected by their strength of association. As predicted, people's acceptance of MT inferences increased when there was no strongly associated disabler available.

We suspect that the non-significance of the expected effect on MP may be due to some specific task characteristics of the experiment. In the pretest, relatively few disablers were generated (less than the overall mean) for the three conditionals that were adopted for the experiment. Cummins (1995) already obtained high MP acceptance ratings for these conditionals. The ''expansion'' manipulation in the present experiment then further decreased the available number of disablers. This may have resulted in a ceiling effect. It could be the case that MP acceptance was already at the top in the weakly associated group. Mean acceptance for MP in the weakly associated group (mean = 5.7, see Table 3) indeed tended to the ''Sure that I can draw this conclusion'' rating, located at the upper end of the scale. As in Cummins (1995), acceptance ratings on the (more difficult) MT inference were lower, what allowed the associative strength effect to show up.

Since there is a strong tendency to accept MP in both expanded groups, it is also possible that the 7-point rating scale was not sensitive enough to detect an effect. In addition, the strength of association factor was manipulated between groups of subjects which may have further hampered the detection of the MP effect.

In order to avoid these possible problems we decided to look at the effect of the strength of association factor in a second experiment.

EXPERIMENT 2

Experiment 2 further examined the effect of the strength of association of the disabling conditions. We tried to avoid the highlighted problems in Experiment 1 by making a number of changes to the materials and design used there. First, in the pretest, we looked for conditionals with a large number of disablers (higher than the overall mean). Since there are more disablers available, acceptance for these conditionals will be lower, which should reduce an eventual ceiling effect. Second, by using an 11-point rating scale instead of the 7-point scale in Experiment 1 we tried to offer participants the possibility for a more detailed discrimination at the extreme levels of the scale. Finally, all participants received both a strongly and weakly expanded conditional (with different content) in order to allow a within-subjects test of the strength of association effect.

Participants were presented MP, MT, and AC problems with expanded minor premises based on the selected conditionals. Since the problems possibly resulting in the non-significance of the MP effect in Experiment 1 were avoided, we expected to find an effect of the associative strength manipulation on both MP and MT.

The AC problem was included in order to further examine the effect of associative strength of disablers on the fallacies. In Experiment 1 we found no significant effects on the AC and DA inferences. However, we neither observed a significant interaction between the different inference types. Especially for AC, acceptance ratings tended to increase when no strongly associated disabler was available. By presenting participants AC inferences in addition to the MP and MT inferences we could examine the generality of this trend.

Method

Participants. Thirty-seven first-year psychology students volunteered to participate in the experiment.

Material. For the strength of association manipulation it was crucial to find conditionals with a set of disablers in which there is only one that is very frequently generated. Although most conditionals with many disablers in the pretest had typically also a larger number of very frequently generated (>50% of participants) disablers, two conditionals (i.e., ''If Jenny turns on the air conditioner, then she feels cool'' and ''If the brake is depressed, then the car slows down'') could still be selected that were appropriate for the present purpose. The most frequently generated disabler was mentioned by 95% and 85% of the participants, whereas the second most frequently generated disabler was mentioned by only 50% and 45% of the participants, respectively. Both conditionals have also a comparable number of possible alternative conditions (see Cummins, 1995; De Neys et al., 2000). Strongly and weakly associated problem versions were constructed by expanding the minor premise with the negation of the strongest and second strongest associated disabler, respectively. The different versions are presented in Table 4 (for an MP problem).

As in Experiment 1, the original conditional was presented on top of an item page. The expanded MP, MT, and AC inferences were then presented on the same page. Below each inference problem appeared the 11-point rating scale shown in Figure 2. Participants received a three-page booklet with instructions on the first page. On the second and third page appeared the two selected conditionals, one with strongly expanded inferences problems and the other with weakly expanded ones. In half of the booklets the strongly expanded inference

TABLE 4
Material for Experiment 2

Expanded strongly associated:	
	(a) Jenny turns on the air conditioner and the air conditioner is not broken (95%)
	(b) The brake is depressed and the brake is not broken (85%)
Expanded weakly associated:	
	(a) Jenny turns on the air conditioner and she has no fever (50%)
	(b) The brake is depressed and the road is not slippery (45%)

The generation frequency (%) of the corresponding disablers in the pretest is presented between brackets.

problems were presented first, and the order was reversed for the other half. Each conditional was used for the weakly expanded inference problems in approximately half of the booklets and for the strongly expanded problems in the other half.

Procedure. The procedure was similar to that used in Experiment 1.

Results and discussion

An ANOVA was performed on the acceptance ratings with inference type (MP, MT, AC) and strength of association (weakly or strongly expanded) as within-subject factors. Due to the within-subjects manipulation of the strength of association factor, each participant rated both weakly and strongly expanded inferences. Although the three strongly and weakly expanded inferences were based on different conditionals, it is still possible that the order in which the strength of association manipulation was presented in the booklets (weakly expanded inference problems presented before or after the strongly associated ones), biased the ratings. Therefore, presentation order was entered as a between-subjects factor in the ANOVA.

The acceptance ratings corresponding to the numbers 5 to 1 on the left hand of the 11-point rating scale were recoded and assigned the values −5 to −1 such that increasing numbers corresponded to increased acceptance.

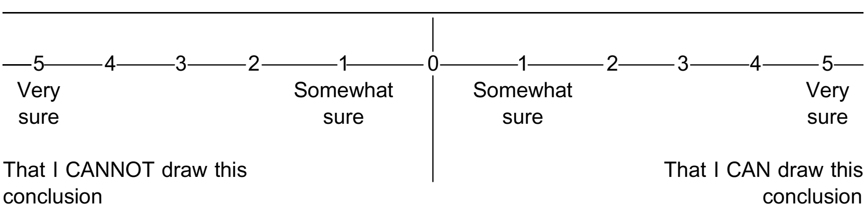

Figure 2.

Results showed that the inference type and strength of association factor interacted significantly, $F(2, 70) = 3.66$, $MSe = 2.22$, $p < .05$. Planned contrast tests indicated that the acceptance ratings for MT, $F(1, 35) = 4.54$, $MSe = 3.76$, $p < .05$, were higher in the strongly associated group, whereas the strength of association factor had no effect on AC acceptance ratings (see Table 5). We also found a significant effect of the strength of association manipulation on the MP acceptance ratings: As MT, MP was more accepted when there was no strongly associated disabler available, $F(1, 35) = 12.41$, $MSe = .53$, $p < .005$.

The presentation order factor did not affect the strength of association effect for any of the three inference types, what confirms that the results were not biased by the order in which participants rated the expanded conditional inferences.

One might note that, although not significant, AC acceptance in Experiment 2 tended to be somewhat lower in the expanded strongly associated group. In Experiment 1 however, there was a trend in the opposite direction with higher AC and DA acceptance when there was no strongly associated disabler available. Detailed examination of these (reversed) AC and DA trends indicated they could be attributed to a complication due to the specific nature of the presented thematic material.

The complication lies in the fact that some disablers can also qualify as possible alternatives (e.g., see Manktelow & Fairley, 2000). In Experiment 1, this was the case for one of the strongly associated disablers we eliminated (i.e., ''If the apples are ripe and they are not picked, then they fall from the tree''). In an alternatives generation task (see De Neys et al., 2000), 55% of the participants generated ''apples dropped by picker'' as a possible alternative. Now, if the apples are not picked, they can not be dropped by the picker. Thus, the manipulation also eliminated a possible alternative. Therefore, the number of available alternatives was smaller in the expanded strongly associated group,

TABLE 5
Mean acceptance ratings for the strongly and weakly expanded inference problems

Inference type	*Expansion type*	
	Expanded weakly associated	*Expanded strongly associated*
MP	3.78**	4.38**
MT	2.03*	2.97*
AC	2.76	2.43

*Planned contrast $p < .05$, **$p < .01$.

The rating scale ranged from –5 (*very sure cannot draw this conclusion*) to 5 (*very sure can draw this conclusion*) with 0 representing *cannot tell*.

what should result in slightly increased DA and AC acceptance. For the other presented disablers the complication was not present. In fact, when the conditional ''If the apples are ripe, then they fall from the tree'' was removed from the analysis in Experiment 1 the trends on AC and DA dissolved, whereas the trends on MP and MT were unaffected.[3] In Experiment 2, a similar complication was present but now in the expanded weakly associated group. One of the presented disablers (i.e., ''If Jenny turns on the air-conditioning and she has no fever, then she feels cool'') indeed qualified as possible alternative. In an alternative generation task 35% of participants generated ''having cold fever'' as a possible alternative. Consequently, in Experiment 2, the number of available alternatives was somewhat smaller in the expanded weakly associated group. Thus, the DA and AC trends in Experiment 1 and 2 can be explained by an impact on the availability of alternatives. In line with what Cummins (1995) found for the number of disablers, this indicates that associative strength of a stored disabler has no effect on the fallacies *per se*.

GENERAL DISCUSSION

The experiments in this study allow us to conclude that MP and MT acceptance increases when there is no strongly associated disabler available. The results from the second experiment established that, as expected, the manipulated availability of disabling conditions affected both MP and MT acceptance. These results support the hypothesis that in addition to the number of disabling conditions (Cummins, 1995), retrieving disablers from semantic memory is affected by their strength of association.

Taken together, the outcome of the two experiments suggests there is no impact of the associative strength of stored disablers on the DA and AC inferences. The observed trends could be attributed to an impact of the disabler manipulation on the availability of the alternatives. This is consistent with Cummins' (1995) findings, where the number of stored disablers had no effect on AC or DA.

Nevertheless, one should note that deviations of Cummins' findings have been reported. Some studies (e.g., Liu, Lo, & Wu, 1996; Markovits & Potvin, 2001, Exp. 3) did report an impact of disabler retrieval on AC and DA acceptance. This would not be surprising when the outcome of the disablers search process could affect the efficiency of the alternatives search process. In that case one could hypothesise that successful disabler retrieval will be resource demanding what might burden the search for alternatives. However, at present, such a possible interplay between the search processes has not yet been examined. Together with the ''disablers qualifying as alternative'' complication,

[3] Mean acceptance ratings in strongly vs weakly expanded group: MP = 6.08 vs MP = 5.77, DA = 5.60 vs DA = 5.60, MT = 5.38 vs MT = 4.66, AC = 5.96 vs AC = 5.80.

this makes it clear that more research is needed to establish the precise relation between the disablers and alternatives retrieval processes. In the absence of such more specific studies, the present DA and AC conclusions should be interpreted with some caution.

In this study we adopted Quinn and Markovits' (1998) notion of a semantic search process and extended it to the disabling conditions. We should note that Quinn and Markovits (see also Markovits et al., 1998) incorporated the postulated semantic search process in the mental models theory (Johnson-Laird & Byrne, 1991). They argued that successful retrieval of an alternative would lead to the construction of an additional [not-P Q] model. This model would represent the possibility that the effect specified in the conditional occurs without the occurrence of the specified cause. Therefore, AC and DA would no longer be supported.

This account can easily be extended to incorporate the present findings. When the semantic search process retrieves a stored disabler an additional [P not-Q] model would be constructed. This model will represent that it is possible that occurrence of the antecedent will not result in occurrence of the consequent. With such a model MP and MT will no longer be supported. Since the probability of successful retrieval is lower when there is no strongly associated disabler available, people will be less likely to construct the additional [P not-Q] model. This would account for the higher MP and MT acceptance. In mental models theory, the AC and DA inferences are only affected by models where the consequent does occur [Q] or the antecedent does not occur [not-P]. Thus, whether or not the [P not-Q] model is constructed is not important for AC and DA. This would explain why the associative strength of disablers has no effect here.

However, it is important to notice that the present study only focused on the semantic search process during conditional reasoning. As Quinn and Markovits (1998), we examined a factor that affects successful memory retrieval. This research does not address how the retrieved disabler is actually incorporated in the reasoning process. We therefore refrained from making specific claims about the nature of the basic inferential principles (i.e., mental models or mental inference rules). The general semantic search process can be incorporated in other reasoning theories like mental logic (Braine & O'Brien, 1998; Rips, 1994) or the probabilistic approach (Oaksford & Chater, 1998; Oaksford, Chater, & Larkin, 2000).

In the probabilistic approach for example, acceptance of MP and MT depends upon the value of an ''exceptions parameter'' (i.e., the probability of not-q given p; see also Stevenson & Over, 1995). This parameter represents the probability that exceptions (disablers) will occur. The higher the exceptions value the less MP and MT will be accepted. Now, it is reasonable to assume that a reasoner determines this probability by searching his/her memory for known exceptions. Thus, one can also predict that successful disabler retrieval, by increasing the exceptions parameter, will decrease MP and MT acceptance. Comparing these different implementations is not within the scope of the present experiment or the Quinn and Markovits study.

We mentioned the relevance of the present study for the work of Elio (1997, 1998) and other researchers in the domain of belief revision and non-monotonic reasoning. Elio established that the number of stored disabling conditions affected people's belief revisions and stated that conditional reasoning and belief revision are guided by the same memory search process. Our results show that successful retrieval is not only affected by the number of stored disabling conditions but also by their strength of association.

Finally, the present study can be related to the work of Chan and Chua (1994). They examined the effect of ''relative salience'' of disabling conditions. This factor can be interpreted as strength of association. Chan and Chua presented inference problems with two conditionals (e.g., ''If p then q, If r then q, p, thus q?''). The second conditional mentioned a possible disabling condition, although the categorical premise was not expanded (see Byrne, 1989). Acceptance of MP and MT decreased with the strength of association of the mentioned disabler. However, a crucial difference with our study is that the present manipulation specifically affected the retrieval of disablers from semantic memory. In Chan and Chua's experiment, reasoning was affected by the strength of association of the mentioned disabler *per se*. The expansion of the categorical premise in the present experiment eliminated a strongly or weakly associated disabler and thereby affected the strength of association in the residual disabler set.

In sum, our study indicated that the conditional inferences people make are influenced by the strength of association of the disabling conditions. This complements Quinn and Markovits' (1998) contention that the strength of association of elements in semantic memory is an important factor in predicting conditional reasoning performance.

Manuscript received May 2001
Revised manuscript received August 2001

REFERENCES

Braine, M.D.S., & O'Brien, D.P. (Eds.). (1998). *Mental logic*. Mahwah, NJ: Lawrence Erlbaum Associates Inc.

Byrne, R.M.J. (1989). Suppressing valid inferences with conditionals. *Cognition*, *31*, 61–83.

Byrne, R.M.J., Espino, O., & Santamaria, C. (1999). Counterexamples and the suppression of inferences. *Journal of Memory and Language*, *40*, 347–373.

Chan, D., & Chua, F. (1994). Suppression of valid inferences: Syntactic views, mental models, and relative salience. *Cognition*, *53*, 217–238.

Cummins, D.D. (1995). Naive theories and causal deduction. *Memory and Cognition*, *23*, 646–658.

Cummins, D.D., Lubart, T., Alksnis, O., & Rist, R. (1991). Conditional reasoning and causation. *Memory and Cognition*, *19*, 274–282.

De Neys, W., Schaeken, W., & d'Ydewalle, G. (2000). *Causal conditional reasoning, semantic memory retrieval, and mental models: A test of the ''semantic memory framework''* (Psychological Rep. No. 270). Leuven, Belgium: University of Leuven, Laboratory of Experimental Psychology.

Elio, R. (1997). What to believe when inferences are contradicted: The impact of knowledge type and inference rule. In *Proceedings of the nineteenth annual conference of the Cognitive Science Society* (pp. 211–216). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Elio, R. (1998). How to disbelieve p → q: Resolving contradictions. *Proceedings of the twentieth meeting of the Cognitive Science Society* (pp. 315–320). Mahwah, NJ: Lawrence Erlbaum Associates Inc.

Evans, J. St. B.T. (2000). What could and could not be a strategy in reasoning. In W. Schaeken, G. De Vooght, A. Vandierendonck, & G. d'Ydewalle (Eds.), *Deductive reasoning and strategies*. Mahwah, NJ: Lawrence Erlbaum Associates Inc.

Evans, J. St. B.T., Newstead, S.E., & Byrne, R.M.J. (1993). *Human reasoning: The psychology of deduction*. Hove, UK: Lawrence Erlbaum Associates Ltd.

Johnson-Laird, P.N., & Byrne, R.M.J. (1991). *Deduction*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

Liu, I., Lo, K., & Wu, J. (1996). A probabilistic interpretation of ''if–then''. *Quarterly Journal of Experimental Psychology*, *49A*, 828–844.

Manktelow, K.I., & Fairley, N. (2000). Superordinate principles in reasoning with causal and deontic conditionals. *Thinking and Reasoning*, *6*, 41–65.

Markovits, H. (1984). Awareness of the ''possible'' as a mediator of formal thinking in conditional reasoning problems. *British Journal of Psychology*, *75*, 367–376.

Markovits, H. (2000). A mental model analysis of young children's conditional reasoning with meaningful premises. *Thinking and Reasoning*, *6*, 335–347.

Markovits, H., Fleury, M., Quinn, S., & Venet, M. (1998). The development of conditional reasoning and the structure of semantic memory. *Child Development*, *69*, 742–755.

Markovits, H., & Potvin, F. (2001). Suppression of valid inferences and knowledge structures: The curious effect of producing alternative antecedents on reasoning with causal conditionals. *Memory and Cognition*, *29*, 736–744.

Newstead, S.E., Ellis, C.E., Evans, J. St. B.T., & Dennis, I. (1997). Conditional reasoning with realistic material. *Thinking and Reasoning*, *3*, 49–96.

Oaksford, M., & Chater, N. (1998). *Rationality in an uncertain world: Essays on the cognitive science of human reasoning*. Hove, UK: Psychology Press.

Oaksford, M., Chater, N., & Larkin, J. (2000). Probabilities and polarity biases in conditional inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 883–899.

Quinn, S., & Markovits, H. (1998). Conditional reasoning, causality, and the structure of semantic memory: Strength of association as a predictive factor for content effects. *Cognition*, *68*, B93–B101.

Rips, L.J. (1994). *The psychology of proof*. Cambridge, MA: MIT Press.

Rumain, B., Connell, J., & Braine, M.D.S. (1983). Conversational comprehension processes are responsible for reasoning fallacies in children as well as adults. *Developmental Psychology*, *19*, 471–481.

Stevenson, R.J., & Over, D.E. (1995). Deduction from uncertain premises. *Quarterly Journal of Experimental Psychology*, *48A*, 613–643.

Thompson, V.A. (1994). Interpretational factors in conditional reasoning. *Memory and Cognition*, *22*, 742–758.

Vadeboncoeur, I., & Markovits, H. (1999). The effect of instructions and information retrieval on accepting the premises in a conditional reasoning task. *Thinking and Reasoning*, *5*, 97–113.