What We Can Learn from Causal Conditional Reasoning about the Naïve Understanding of Causality

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

Abstract

Causal conditional reasoning means drawing inferences from a conditional statement that refers to causal content. It is argued that data on causal conditional reasoning not only tell us something about how people draw deductive inferences from conditionals, but also provide us with information about how they understand causal relations. In particular, three principles emerge from existing data: the modal principle, the exhaustive principle, and the equivalence principle. An experiment sheds new light on how people interpret and use conditionals in causal contexts, and reveals evidence for the proposed representational principles.

Keywords: Causality, conditionals, deduction, content effects, dual source approach.

Introduction

Causal conditional reasoning means drawing inferences from a conditional premise that refers to causal content. By evaluating these inferences against the yardstick of propositional logic, the hope is that causal conditional reasoning tasks will provide a deeper insight into people's understanding of conditionals (e.g., Johnson-Laird & Byrne, 2002). Consider the following example:

(1) If a car is involved in a serious accident,
then the airbag inflates.

According to the conditional relation "*If* P, *then* Q", this statement first of all expresses that a condition "P" is sufficient for a consequent "Q". This justifies two inferences, *Modus Ponens* and *Modus Tollens*:

MP: "*If* P, *then* Q" and "P" – infer "Q"
MT: "*If* P, *then* Q" and "*not* Q" – infer "*not* P"

In the other two cases, *Denial of the Antecedent* and *Affirmation of the Consequent*,

DA: "*If* P, *then* Q" and "*not* P"
AC: "*If* P, *then* Q" and "Q"

it depends on the necessity of "P" whether or not definite conclusions are possible: If "P" is not necessary for "Q" (conditional interpretation), then definite conclusions are not justified. Only if "P" is regarded as necessary for "Q" (biconditional interpretation) can we definitively infer "*not*-Q" from "*not*-P" and "P" from "Q".

The content of the Ps and Qs is not important to obtain these interpretations; the only important aspect is the relation between the antecedent "P" and the consequent "Q", which is established by the terms "if-then".

In causal conditional reasoning tasks, however, people follow this content-independent interpretation very rarely. Instead, they interpret the conditional premise content-specifically as a statement referring to a cause and its effect. In addition, they even consider what they know about alternative causes for this effect, or about inhibitory factors that need to be absent for the effect to occur (e.g., Cummins, 1995; Cummins, Lubart, Alksnis & Rist, 1991; de Neys, Schaeken & D'Ydewalle, 2002, 2003; Quinn & Markovits, 2002; Thompson, 1994).

The data show the following general pattern: If people can think neither of alternative causes nor of inhibitory factors, then the *biconditional* interpretation is typically strengthened, and definite inferences are drawn in all four cases (MP, MT, DA, and AC). If people can easily think of alternative causes, then the *conditional* interpretation is strengthened, that is, MP and MT increase, while the definite DA and AC inferences decrease. If, instead, people can easily think of inhibitory factors, then the *reversed* interpretation is strengthened, that is, MP and MT decrease, while the definite DA and AC inferences increase. Sometimes, the retrieval of an alternative cause also activates knowledge about inhibitory factors (Markovits & Potvin, 2001). Analogous effects of alternative causes and inhibitory factors were also found in causal selection tasks that require people to evaluate whether or not a causal conditional is violated (e.g., Beller & Spada, 2003; Fairley, Manktelow & Over, 1999).

Typically, these findings are looked at from the perspective of conditional reasoning. Thereby, it is asked: What do we learn from causal conditional reasoning about people's understanding of and deductions from *conditional statements*? From this perspective, the experiments show that the core meaning of the conditional is modulated by semantic or pragmatic information retrieved from long-term memory (Johnson-Laird & Byrne, 2002) and that people flexibly include such information in the reasoning process (de Neys et al., 2002, 2003; Markovits & Quinn, 2002; Quinn & Markovits, 2002), depending on working memory resources (de Neys, et al., 2005). The phenomenon of causal conditional reasoning, however, is like a coin with two sides: the form of the premise (a conditional) and its content (a causal relation). These should, at least initially, be examined separately in order to understand how they mutually influence one another (Beller & Spada, 2003). This perspective triggers a complementary question: What do we learn from causal conditional reasoning about people's understanding of and deductions from *causal relations*? As detailed in the next section, at least three general principles can be identified.

Some Principles of Causal Understanding

Three basic principles of causal understanding emerge from the causal conditional reasoning data, which may be called the modal principle, the exhaustive principle, and the equivalence principle (Beller & Kuhnmünch, 2006).

The modal principle was formulated by Goldvarg and Johnson-Laird (2001). It states that causal knowledge represents possible and impossible states of the world making such knowledge particularly useful for prediction and intervention. The use of conditional statements "*If* P *then* Q" is compatible with the modal principle as conditionals also define which states of affairs are possible and which are not. According to the conditional interpretation, for example, only the combination of "P" and "*not*-Q" is impossible.

According to **the exhaustive principle** people try to consider all causally relevant factors *exhaustively* when they reason about an effect. Causal factors may contribute to the effect in different ways, necessitating different ways of integration: Alternative causes require a disjunction of factors, while a compound cause requires a conjunction of factors. The causal conditional reasoning data also support this second principle. As explained above, people tend to go beyond the causal factors mentioned in the conditional and integrate additional knowledge about alternative causes or inhibitory factors. In other words: They consider what they regard as causally relevant for the task at hand. The exhaustive principle does not imply that people in fact have complete causal knowledge about the causes for an effect; it only means that they assume – until there is evidence to the contrary – that they have considered all relevant factors, as otherwise causal inferences would not be possible at all.

Finally, **the equivalence principle** qualifies the relation between an effect and the whole set of causes for this effect as an *equivalence* (Beller & Spada, 2003): Every effect has a cause (a position going back to the Scottish philosopher David Hume, 1711-1776; cf. Wilson, 1937), and without a cause there will be no effect. This corresponds with the assumption that the world is causally determined and reacts in a causally reliable manner. While the first two principles are so basic that everyone swiftly agrees, the equivalence principle is more critical because it implies that people presume a strong conception of causation. Goldvarg and Johnson-Laird's (2001) mental model theory, for example, also allows for weak causation (i.e., sufficient but not necessary causes). Furthermore, the causal conditional reasoning data seem to support the equivalence principle only in parts: As stated above, people in fact apply a *biconditional* interpretation to the conditional (in correspondence with a strong causal relation) but only if they are sure that there are no other causal factors. If alternative causes or inhibitory factors come into play, the inference patterns change suggesting either a conditional interpretation if alternative causes are thought of, or the reversed interpretation if inhibitory factors are considered. It is argued in this paper that even in these latter cases people presume a strong causal relation between an effect an its causes, but that we need other than conditional reasoning tasks to detect it.

To illustrate this point, let us consider three examples that all refer to the causal connection between a car accident and the reaction of an airbag. Provided that a person considers *only* the accident causally relevant, according to the modal and equivalence principle this can be represented as follows[1]:

(2) $\texttt{accident}_{\texttt{Car}} \leftrightarrow \texttt{airbag\_inflation}_{\texttt{Car}}$

An accident is necessary and sufficient for the inflation of the airbag justifying definite inferences in all four conditional inference tasks (MP, MT, DA, and AC) with reference to conditional (1): "If a car is involved in a serious accident, then the airbag inflates."

If further causal factors come into play, these need to be considered according to the exhaustive principle. In the airbag scenario, for example, the sensor that triggers the airbag reaction may come to mind. It may react in an oversensitive way or not sensitively enough. An oversensitive sensor opens up the possibility that the airbag is inflated in the absence of an accident by an alternative cause. This situation includes a disjunction of causes:

(3) $(\texttt{accident}_{\texttt{Car}} \vee \texttt{oversensitive\_sensor}_{\texttt{Car}}) \leftrightarrow \texttt{airbag\_inflation}_{\texttt{Car}}$

An accident is still sufficient, but is no longer necessary. Therefore, definite inferences are possible only in the case of MP and MT (conditional pattern).

The situation is different if, instead, a person assumes that the sensor is not sensitive enough. This defect corresponds to an inhibitory factor that prevents the airbag reaction even though a car is involved in an accident. The absence of the defect enters into the causal relation conjunctively:

(4) $(\texttt{accident}_{\texttt{Car}} \wedge \neg \texttt{insensitive\_sensor}_{\texttt{Car}}) \leftrightarrow \texttt{airbag\_inflation}_{\texttt{Car}}$

Because an accident is necessary for the airbag reaction, but is no longer sufficient, definite inferences are possible only in the case of DA and AC (reversed conditional pattern).

Note that the equivalence principle can be applied to all three cases equally. It characterizes the relation between the effect and the *whole set of causes* as a strong causal relation. However, if more than one causal factor is involved as in (3) and (4), the strong causal relation cannot be detected with conditional reasoning tasks, as these do not require people to draw inferences about the whole set of causes but only about one factor ("P") for the effect ("Q"). It is the aim of the following experiment to shed light on people's interpretational strategies in causal conditional reasoning and to find evidence for the proposed causal principles.

Experiment

The experiment uses the airbag scenario described above. The basic causal connection is introduced by means of conditional (1). Then, the salience of an alternative or an inhibitory factor is manipulated. Finally it is assessed which conditional inferences people draw and how the complete

[1] Logical operators used and their approximate linguistic form:

$P \leftrightarrow Q$	equivalence	"If and only if P then Q"
$P \vee Q$	disjunction	"P or Q or both"
$P \wedge Q$	conjunction	"P and Q"
$\neg P$	negation	"not P"

causal situation is reformulated – as a conditional or a biconditional statement. It is expected that the inference tasks will replicate the typical effects of causal conditional reasoning (partial support for the equivalence principle), while the reformulation task will fully support this principle.

Method

Materials. Three airbag scenarios were constructed. The first scenario mentioned only *one cause*, the second additionally mentioned a further *alternative cause*, while the third scenario additionally mentioned an *inhibitory factor.*

The relevant causal information was presented in an introductory section, which described the reason for inventing the airbag and introduced the causal conditional [additional parts for the *alternative cause* and *inhibitory factor* version are printed in square brackets]:

> *The airbag was invented to increase the security of car passengers. It protects persons in the front seats from frontal impacts in the case of an accident. [The airbag inflates triggered by a sensor. Normally,] The airbag functions according to the following simple causal rule:*
>
> *If a car is involved in an accident,*
> *then the airbag inflates.*

The three scenarios proceeded as follows. In the "one cause" version, no additional causal factor was mentioned:

> *Imagine that the airbag of the following cars A to D is functioning correctly. What can be inferred in this context for the cars A, B, C, and D?*

In the two other versions, a second factor was introduced:

> *Imagine that the sensor that triggers the inflation of the airbag of the following cars A to D*
> *might react in an oversensitive way* {"alternative cause"}/
> *might not be sensitive enough* {"inhibitory factor"}.
> *However, you do not know for sure whether this defect is actually present. What can be inferred in this context for the cars A, B, C, and D?*

Then, four inference tasks were presented, corresponding to MP, DA, AC, and MT with reference to the causal conditional. In each case, three answer options were presented from which participants were required to choose one:

MP: *Car A is involved in an accident.*
What follows from this? (a) The airbag inflates. (b) The airbag does not inflate. (c) It cannot be decided whether or not the airbag inflates.

DA: *Car B is <u>not</u> involved in an accident.*
What follows from this? (a) The airbag inflates ...

AC: *The airbag of car C inflates.*
What follows from this? (a) The car is involved in an accident. (b) The car is not involved in an accident. (c) It cannot be decided whether or not the car is involved in an accident.

MT: *The airbag of car D does <u>not</u> inflate.*
What follows from this? (a) The car is involved ...

The inference tasks were followed by an *evaluation task*, which aimed at examining the perceived sufficiency and necessity of the antecedent clause of the conditional. The task repeated the information of the introductory section, including the conditional rule. Then, two questions were posed (each to be answered with *Yes* or *No*):

> *Is an accident sufficient for the inflation of the airbag?*
>
> *Is an accident necessary for the inflation of the airbag?*

Finally, a *reformulation task* required participants to choose the best reformulation of the causal situation. Its aim was to assess the necessity and sufficiency status of the causal factor(s) as well as its (their) relation to the effect, that is, whether it is interpreted as a conditional or as a biconditional. The task repeated the information of the introductory section, including the conditional rule. Following this, several reformulations were presented, together with the question:

> *Which of the following formulations do you consider to be most appropriate for explaining to another person the exact causal relation between an accident [, a sensor that reacts in an oversensitive way / a sensor that is not sensitive enough] and the airbag reaction?*

In the "one cause" scenario, participants had to choose between a conditional and a biconditional reformulation:

> *If a car is involved in an accident, then the airbag inflates; otherwise you do not know whether or not the airbag inflates.* {conditional}
>
> *If a car is involved in an accident, then the airbag inflates; otherwise the airbag does not inflate.* {biconditional}

In the scenarios with two causal factors, four rules were given from which to choose. They were constructed by integrating the additional factor into the rules presented above either conjunctively or disjunctively, resulting in two conditional and two biconditional rules. In the "alternative cause" version, the antecedent clauses read as follows:

> *If a car is involved in an accident <u>and</u> the sensor reacts oversensitively, then ...*
>
> *If a car is involved in an accident <u>or</u> the sensor reacts oversensitively, then ...*

In the "inhibitory factor" version, these clauses read:

> *If a car is involved in an accident <u>and</u> the sensor is sensitive enough, then ...*
>
> *If a car is involved in an accident <u>or</u> the sensor is sensitive enough, then ...*

Participants. 177 students from introductory courses on cognitive psychology at the University of Freiburg participated in the experiment. 65 students were male and 111 female; the mean age was $M = 23.8$ years ($SD = 5.5$; *range*: 19-51; one person did not indicate his or her sex and age).

Design and procedure. Participants were randomly assigned to one of three experimental conditions ($n = 59$) corresponding to three scenarios (*one cause, alternative cause*, and *inhibitory factor*). Each participant received a booklet containing general instructions and six tasks. The four inference tasks were presented in one block to ensure that participants made their inferences on the basis of their spontaneous interpretation of the scenario. This block was followed by the evaluation task and the reformulation task.

Results

In order to determine whether the experimental manipulation of the necessity and sufficiency of the primary cause had the intended effect, the evaluation task is analyzed first.

Evaluation task: In this task, participants were required to directly evaluate the necessity and sufficiency of the conditional's antecedent clause ("a car is involved in an accident"). According to the typical effects of causal conditional reasoning, this evaluation should not depend on the conditional *form*, which does not vary across the three scenarios, but on the causal *content* that systematically changes the necessity and sufficiency status of the accident as represented in the expressions (2), (3), and (4).

Participants' evaluations are shown in Table 1. The perceived sufficiency of an accident varied as expected across the three conditions ($\chi^2 = 79.2$; $df = 2$; $N = 177$; $p < .001$) and also the perceived necessity ($\chi^2 = 79.1$; $df = 2$; $N = 177$; $p < .001$). In the "one cause" condition, there was no difference between sufficiency and necessity ($z = -.35$; $p = .724$; Wilcoxon test). According to the equivalence principle, the majority of participants evaluated an accident as sufficient ($\chi^2 = 6.1$; $df = 1$; $n = 59$; $p = .013$) and as necessary for the airbag reaction ($\chi^2 = 3.8$; $df = 1$; $n = 59$; $p = .051$). In the "alternative cause" condition, an accident was mostly regarded as sufficient but not necessary ($z = -7.1$; $p < .001$; Wilcoxon test), whereas in the "inhibitory" factor condition, it was regarded as not sufficient but necessary ($z = -5.9$; $p < .001$; Wilcoxon test). Altogether, the experimental manipulation had the intended effects.

Inference tasks: These tasks assess the four inferences MP, DA, AC, and MT from causal conditional (1). If people apply the expected causal interpretations (2), (3), and (4), the inferences should correspond with the evaluation results. The influence of the scenarios was first checked for each task separately by means of a χ^2 test comparing the frequencies of the two critical answers that are predicted by the conditional or the causal interpretation (the third answer option was omitted in these tests, since it was chosen only in a few cases). In all tasks, significant scenario effects were found, indicating a strong influence of the causal context (Table 2).

Next, two aggregated measures were computed for each participant: the mean percentage of *conditional* inferences and the mean percentage of *causal* inferences. Inspecting the number of conditional inferences, we find the typical effects of causal conditional reasoning: Compared to the "one cause" scenario (61.4%), the number of conditional inferences increased to 91.5% in the "alternative cause" scenario, and decreased to 10.2% in the "inhibitory factor" scenario; $F_{(2,174)} = 195.6$; $p < .001$. Altogether, 54.4% of all inferences follow the conditional interpretation. Assuming a causal interpretation fits the data much better: 84.6% of all inferences are consistent with this interpretation, though with a significant difference between the "one cause" scenario (74.6%) and the other two scenarios (91.5% and 88.1%); $F_{(2,174)} = 9.5$; $p < .001$. An examination of the par-

Table 1: Evaluation of the conditional's antecedent clause (%; $n = 59$ in each group; causally expected answers **bold**-faced).

An accident is ...	*One cause*	*Alternative Cause*	*Inhibitory factor*
sufficient	**66.1**	**94.9**	15.3
necessary	**62.7**	6.8	**86.4**

Table 2: Causal conditional inferences (%; $n = 59$ in each group; causally expected answers **bold**-faced).

	One cause	*Alternative Cause*	*Inhibitory factor*
MP: The car is involved in an accident. The airbag ...			
... inflates.	**89.8**^{C}	**98.3**^{C}	8.5^C
... does not inflate.	0.0	0.0	3.4
undecidable	10.2	1.7	**88.1**
DA: The car is not involved in an accident. The airbag ...			
... inflates.	1.7	1.7	1.7
... does not inflate.	**64.4**	11.9	**89.8**
undecidable	33.9^C	**86.4**^{C}	8.5^C
AC: The airbag inflates. The car ... in an accident.			
... is involved	**61.0**	6.8	**88.1**
... is not involved	0.0	0.0	0.0
undecidable	39.0^C	**93.2**^{C}	11.9^C
MT: The airbag does not inflate. The car ... in an accident.			
... is involved	0.0	3.4	1.7
... is not involved	**83.1**^{C}	**88.1**^{C}	11.9^C
undecidable	16.9	8.5	**86.4**
Inferences aggregated across all four tasks			
conditional	61.4^C	91.5^C	10.2^C
causal	**74.6**	**91.5**	**88.1**
Consistent inference patterns across all four tasks			
biconditional	**49.2**	1.7	3.4
conditional	27.1^C	**74.6**^{C}	1.7^C
reversed cond.	6.8	0.0	**71.2**
indeterminate	1.7	0.0	5.1
ambiguous	15.3	23.7	18.6

^{C} Logically correct according to the material conditional.
MP: $\chi^2 = 126.1$; $df = 2$; $n = 175$; $p < .001$
DA: $\chi^2 = 77.1$; $df = 2$; $n = 174$; $p < .001$
AC: $\chi^2 = 81.1$; $df = 2$; $n = 177$; $p < .001$
MT: $\chi^2 = 93.2$; $df = 2$; $n = 174$; $p < .001$

ticipants' *patterns of inferences* across all tasks shows the reason why.

For this final analysis, four types of interpretations were distinguished: biconditional, conditional, reversed conditional (i.e. complementary to the conditional), and indeterminate (i.e. giving the answer "undecidable" in all tasks). The inference pattern of a participant was assigned to a particular category if all four inferences matched; otherwise, the pattern was classified as ambiguous. The frequencies of the various patterns are shown at the bottom of Table 2.

The three conditions did not differ with regard to the total number of consistent patterns; $\chi^2 = 1.4$; $df = 2$; $N = 177$; $p = .501$. Consistent patterns clearly predominated (80.8% vs. 19.2% ambiguous on average; $\chi^2 = 67.1$; $df = 1$; $N = 177$; $p < .001$). This suggests that the majority of participants applied one particular interpretation and solved the four inference tasks accordingly. Altogether, 115 of the 143 consistent patterns (80.4%) fell into the causally predicted categories. An accident was quite uniformly regarded as sufficient, but not necessary in the "alternative cause" condition (conditional pattern), and as necessary, but not sufficient in the "inhibitory factor" condition (reversed conditional pattern). In the "one cause" condition, however, only about half of the participants followed the causally predicted biconditional pattern, while a substantial proportion followed the conditional interpretation suggested by the form of the premise. The reason might be that, in the "one cause" condition without any further causal context, the task is open to both interpretations, while taking the alternative cause or the inhibitory factor into account *necessitates* the causal interpretation. Altogether, this analysis reveals that people sometimes follow different, but consistent interpretations, which are hidden if we analyse each task separately.

Reformulation task: In the final task, participants had to choose a statement that describes the complete causal situation. To do this, they needed not only to decide how to integrate the alternative cause and inhibitory factor respectively (disjunctively vs. conjunctively), but also to determine the relation between the causal factor(s) and the effect as either conditional or biconditional. Especially this latter aspect should give information about the proposed equivalence principle. The results are shown in Table 3.

Table 3: Reformulation of the causal situation (%; $n = 59$; causally expected answers **bold**-faced).

		One cause	*Alternative Cause*	*Inhibitory factor*[1]
(A)	*Integration of causes*			
	disjunction	–	**96.6**	3.4
	conjunction	–	3.4	**96.6**
(B)	*Type of conditional*			
	conditional	44.1	15.3	32.8
	biconditional	**55.9**	**84.7**	**67.2**

[1] One missing answer; $n = 58$.

First, the additional causal factor was considered in accordance with the evaluation and inference results (Table 3A): The alternative cause was integrated by using a disjunction and the inhibitory factor was integrated by using a conjunction; $\chi^2 = 101.5$; $df = 1$; $n = 117$; $p < 0.001$.

Second, with regard to the type of conditional (Table 3B), the causally predicted biconditional formulation dominated, although we also found differences between the scenarios; $\chi^2 = 11.7$; $df = 2$; $N = 176$; $p = .003$. The biconditional interpretation predominated in the "alternative cause" condition (84.7% biconditional; $\chi^2 = 28.5$; $df = 1$; $n = 59$; $p < .001$), and in the "inhibitory factor" condition (67.2% biconditional; $\chi^2 = 6.9$; $df = 1$; $n = 58$; $p = .009$). This means that the majority of these participants regarded the causal factor(s) as exhaustive and applied the equivalence principle: If at least one sufficient cause occurs, then the effect will also occur, otherwise not. Interestingly, the support for the equivalence principle from the "one cause" scenario was not as strong as from the other two scenarios (55.9% biconditional; $\chi^2 = .8$; $df = 1$; $n = 59$; $p = .362$).

What might be the reason for this? A possible explanation is that some people did not reformulate the *causal relation* between an accident and the airbag reaction, but reformulated the original *conditional statement* instead, since there was no need to go beyond this. In the other two conditions, however, participants did have to go beyond the original conditional, as the reformulation required the integration of a second causal factor; simply re-stating the original conditional was no longer sufficient.

Altogether, the reformulation data indicate that people's biconditional interpretations often observed in causal conditional reasoning tasks do not result from interpreting the conditional relation per se, but from a specific understanding of causal relations that includes an equivalence relation.

Discussion

The inferential results reveal the typical effect of causal conditional reasoning: People's answers systematically depend on knowledge about alternative causes and inhibitory factors. It was argued that these effects result from a specific understanding of causal relations: Participants regarded the presented causal factor(s) as exhaustive for the effect in question (*exhaustive principle*), they integrated them according to their causal background knowledge either disjunctively or conjunctively, they regarded them as, together, being necessary and sufficient for the effect (*equivalence principle*), and they inferred from this which situations are possible and which are not (*modal principle*). Further evidence for these principles is provided by three experiments that replicated the present results and extended them to pure causal tasks without conditional premise and to abstract causal tasks (Beller & Kuhnmünch, 2006), as well as by an experiment that demanded a differentiation between causes and enabling conditions (Kuhnmünch & Beller, 2005).

The data also reveal a modulating influence of the conditional formulation. This again confirms the necessity of *dual source* analyses (Beller & Spada, 2003): Both the domain-general understanding of the conditional and the domain-specific understanding of the causal relation are needed to explain the data. The experiment shows in particular how these two sources of information play together: People represent causal relations according to specific principles and reason from these relations with remarkable precision (*content competence*). In conditional reasoning tasks, this competence leads to *content effects*: People's conditional inferences are "biased" towards the respective causal interpretation. This does not mean, however, that their reasoning is deductively invalid. It only means that people supplement the conditional premise with relevant information from long-term memory serving as additional premises (cf. Henle, 1962, for a similar argument). With regard to the underlying causal relations, the inferences are deductively valid. Interestingly, our analysis of inference patterns uncovered that many people consistently use only one source of information. Finally, the reformulation task required people to think explicitly about conditional statements in relation to the causal content – and, in such tasks, the majority of them mastered the logical connectives *if-then*, *and*, *or* correctly (*form competence*). Evidence for all three dual source phenomena – content competence, content effects, and form competence – was also found in other content domains: with conditional promises and threats (Beller, Bender & Kuhnmünch, 2005; Beller & Spada, 2003), with deontic rules that define what is forbidden and allowed (Beller, 2003), and with various social and conceptual relations (Neth & Beller, 1999).

As argued in the introduction, the data on causal conditional reasoning tell us something about how people understand causal relations. Three representational principles could be generated from these data. However, the same principles appeared to also be involved in the interpretation of deontic rules (Beller, 2003); these principles are therefore not sufficient to characterize people's notion of causality unambiguously. Further aspects are necessary to distinguish the concept of causality from other concepts, including the following: Causality is concerned with systematic changes in the physical world; causes are typically assumed to precede their effects (cf. Goldvarg & Johnson-Laird, 2001); and mechanisms are assumed to transmit some sort of force, energy or "causal power" in order to produce the change (e.g., Ahn & Kalish, 2000). Only a detailed analysis of how people understand the concept of causality can help to explain the full range of phenomena in causal reasoning and thus in causal conditional reasoning.

Acknowledgements I am grateful to Andrea Bender and Stefan Kleinbeck for helping to collect the data, and to Andrea Bender, Christoph Klauer, Gregory Kuhnmünch, Sarah Mannion, and Mike Oaksford for valuable comments on an earlier version of this paper.

References

Ahn, W., & Kalish, C. W. (2000). The role of mechanism beliefs in causal reasoning. In F. C. Keil, & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 199-225). Cambridge, MA: The MIT Press.

Beller, S. (2003). The flexible use of deontic mental models. In R. Alterman, & D. Kirsh (Eds.), *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society* (pp. 127-132). Mahwah, NJ: Erlbaum.

Beller, S., Bender, A., & Kuhnmünch, G. (2005). Understanding conditional promises and threats. *Thinking and Reasoning*, *11*, 209-238.

Beller, S., & Kuhnmünch, G. (2006). What causal conditional reasoning tells us about people's understanding of causality. Under review.

Beller, S., & Spada, H. (2003). The logic of content effects in propositional reasoning: The case of conditional reasoning with a point of view. *Thinking and Reasoning*, *9*, 335-378.

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

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

De Neys, W., Schaeken, W., & D'Ydewalle, G. (2002). Causal conditional reasoning and semantic memory retrieval: A test of the semantic memory framework. *Memory & Cognition*, *30*, 908-920.

De Neys, W., Schaeken, W., & D'Ydewalle, G. (2003). Inference suppression and semantic memory retrieval: Every counterexample counts. *Memory & Cognition*, *31*, 581-595.

De Neys, W., Schaeken, W., & D'Ydewalle, G. (2005). Working memory and counterexample retrieval for causal conditionals. *Thinking and Reasoning*, *11*, 123-150.

Fairley, N., Manktelow, K. I., & Over, D. E. (1999). Necessity, sufficiency, and perspective effects in causal conditional reasoning. *Quarterly Journal of Experimental Psychology, 52A*, 771-790.

Goldvarg, E., & Johnson-Laird, P. N. (2001). Naive causality: A mental model theory of causal meaning and reasoning. *Cognitive Science*, *25*, 565-610.

Henle, M. (1962). On the relation between logic and thinking. *Psychological Review*, *69*, 366-378.

Johnson-Laird, P. N., & Byrne, R. M. J. (2002). Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review*, *109*, 616-678.

Kuhnmünch, G., & Beller, S. (2005). Distinguishing between causes and enabling conditions – through mental models or linguistic cues? *Cognitive Science*, *29*, 71-84.

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 & Cognition*, *29*, 736-744.

Markovits, H., & Quinn, S. (2002). Efficiency of retrieval correlates with "logical" reasoning from causal conditional premises. *Memory & Cognition*, *30*, 696-706.

Neth, H., & Beller, S. (1999). How knowledge interferes with reasoning – Suppression effects by content and context. In M. Hahn, & S. C. Stoness (Eds.), *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society* (pp. 468-473). Mahwah, NJ: Erlbaum.

Quinn, S., & Markovits, H. (2002). Conditional reasoning with causal premises: Evidence for a retrieval model. *Thinking and Reasoning*, *8*, 179-191.

Thompson, V. A. (1994). Interpretational factors in conditional reasoning. *Memory & Cognition*, *32*, 742-758.

Wilson, F. (1937). *Hume's defence of causal inference*. Toronto: University of Toronto Press.