The Mental Models Theory of Reasoning

Refinements and Extensions

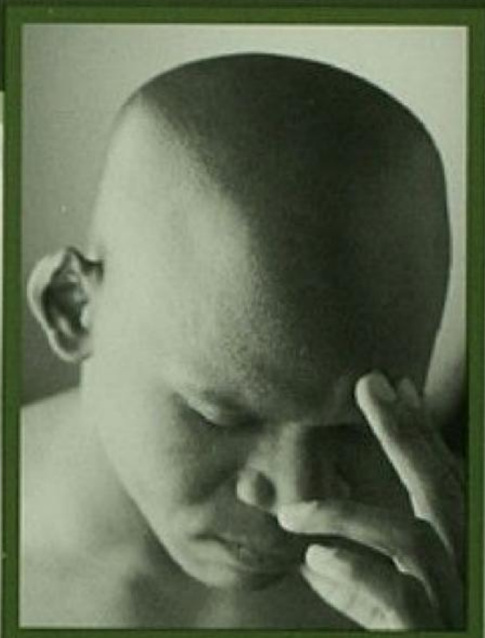

Edited by

Walter Schaeken • André Vandierendonck

Walter Schroyens • Géry d'Ydewalle

The Mental Models Theory of Reasoning

Refinements and Extensions

Edited by

Walter Schaeken
André Vandierendonck
Walter Schroyens
Géry d'Ydewalle

LEA

LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS

2007 Mahwah, New Jersey London

Camera ready copy for this book was provided by the editors.

Lawrence Erlbaum Associates, Inc., Publishers
10 Industrial Avenue
Mahwah, New Jersey 07430
www.erlbaum.com

Cover design by Kathryn Houghtaling Lacey

Library of Congress Cataloging-in-Publication Data

The mental models theory of reasoning : refinements and extensions / edited by Walter Schaeken ... [et al.].
p. cm.

Includes bibliographical references and index.

ISBN 978-0-8058-4183-1 — ISBN 0-8058-4183-0 (cloth)
1. Reasoning (Psychology). I. Schaeken, Walter.
BF442.M47 2007
153.4′3—dc22 2005040691
CIP

Books published by Lawrence Erlbaum Associates are printed on acid-free paper, and their bindings are chosen for strength and durability.

List of Contributors

Pierre Barrouillet, L.E.A.D. - C.N.R.S. UMR 5022, Pôle AAFE, Esplanade Erasme, BP 26513, 21065 Dijon, France

Ruth Byrne, Department of Psychology, Dublin University, Trinity College, Dublin 2, Ireland

Vicky Dierckx, Department of Experimental Psychology, University of Ghent, Henri Dunantlaan 2, 9000 Gent, Belgium

J. St. B. T. Evans, Department of Psychology, University of Plymouth, Plymouth, Devon PL1 5RR, UK

Aidan Feeney, Psychology Department, Durham University, Department of Psychology, South Road, Durham, DH1 3LE, UK

Vittorio. Girotto, Department of Art and Industrial Design, University of Venice,, Dorsoduor 2206, Convento delle Terese, 30123 Venice, Italy

Michel Gonzalez, CNRS, Université de Provence, 29 Av. R. Schumann, F-13621 Aix-en-Provence Cedex 1, France

Eugenia Goldvarg-Steingold, Infant Cognition Laboratory, Department of Brain and Cognitive Science, MIT, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

David Green, Department of Psychology, Centre for Cognitive Science, University College London, Gower Street, London WC1E 6BT, UK

Nelly Grosset, L.E.A.D. - C.N.R.S. UMR 5022, Pôle AAFE, Esplanade Erasme, BP 26513, 21065 Dijon, France

Simon Handley, Department of Psychology, University of Plymouth, Plymouth, Devon PL1 5RR, UK

Philip N. Johnson-Laird, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544, USA

Karl Christoph Klauer, Institut für Psychologie, Albert-Ludwigs-Universität Freiburg, 79085 Freiburg, Germany

Thorsten Meiser, Friedrich-Schiller-Universität Jena, Humboldtstr. 11, D-07743 Jena, Germany

David Over, Business School, University of Sunderland, Reg Vardy Centre, Sir Tom Cowie Campus at St Peter's, St Peters Way, Sunderland, SR6 ODD, UK

Maxwell Roberts, Department of Psychology, University of Essex, Wivenhoe Park, Colchester, Essex, C04 3SQ, UK

Walter Schaeken, Labo Experimental Psychology, University of Leuven, Tiensestraat 102, B-3000 Leuven, Belgium

Walter Schroyens, HEC Montréal, 3000, chemin de la Côte-Sainte-Catherine, Montréal (Québec) H3T 2A7, Canada

Jean-Baptiste Van der Henst, Institut des Sciences Cognitives, 67, Boulevard Pinel, 69675 Bron cedex, FRANCE

André Vandierendonck, Department of Experimental Psychology, University of Ghent, Henri Dunantlaan 2, 9000 Gent, Belgium

Contents

Author Index

Subject Index

Preface

This book has proceeded from the Workshop on mental models and deductive reasoning in Brussels. This workshop was sponsored by the Fund for Scientific Research Flanders and the Federal research project IUAP/PAI P4/19.

We hope that this book will serve the purpose of directing the reader attention to various recent opinions about the status of mental models in the psychology of reasoning.

Mental models are representations in the mind of real or imaginary situations. They can be constructed from perception, imagination, or the comprehension of discourse. Each mental model represents a possibility. Cognitive scientists world-wide have studied how models bring about thoughts and inferences. In the domain of deductive reasoning, the mental model theory is very influential, but definitely not without its criticisms.

We hope to show a sample of the diversity of research related to the role of mental models in reasoning: There are chapters about propositional reasoning, about relational reasoning, about statistical reasoning, and some more meta-theoretical chapters. Moreover, the final collection includes the work of scientists from all over the world: Belgium, England, France, Germany, Ireland, Italy, and the United States.

Since many of the contributions do not deal exclusively with one topic, they have not been explicitly placed in sections. However, the ordering of the chapters reflects the main issue addressed by the authors. Reading the first version of the manuscripts for the book raised the question of what we would do with the presentation of the mental model theory. Indeed, many of the authors gave their own introduction to the theory. We decided that it would be best if the authors produced their chapters to stand alone. As a consequence, each chapter is self-contained and can be read on its own. Some repetition of key theoretical issues therefore occurs.

We owe much to Marleen Devijver. Her secretarial work for the workshop was invaluable. Moreover, she made a major contribution in preparing the manuscripts for publication. We gratefully acknowledge those individuals who contributed to planning and conducting the conference: Vicky Dierckx, Kristien Dieussaert, Niki Verschueren, and Wim De Neys.

Finally, we want to thank all contributors for their enthusiastic cooperation in the realization of the workshop and of this book.

1

Memory Retrieval and Content Effects in Conditional Reasoning: A Developmental Mental Models Account

Pierre Barrouillet
Nelly Grosset

The mental models theory assumes that conditional reasoning is mainly constrained by the number of models individuals can hold and process in working memory and the nature and the accessibility of knowledge used to construct these models. Both constraints result from the limitation of the cognitive resources available to activate knowledge from long-term memory and maintain it active for processing. Because the amount of cognitive resources increases with age, the mental models theory allows precise predictions about the way children, adolescents, and adults interpret conditional sentences and reason from them, and about how contents affect reasoning at different ages. These predictions have been tested in two experiments in which adolescents and adults were asked to reason from conditional sentences that involved either familiar or unfamiliar relations between the antecedent and the consequent. The results confirmed the main developmental predictions of the mental models theory and provided evidence that the fleshing-out process is achieved through a process of retrieval from long-term memory. Both developmental and content effects result from the same processes of constructing and manipulating mental models.

Introduction

Evidence in support of the mental models theory of conditional reasoning (Johnson-Laird, 1999; Johnson-Laird & Byrne, 1991, 2002) has been growing during the last 10 years. According to this theory, understanding an "If *p* then *q*" conditional sentence results in the construction of an initial representation of the following form:

(1) p q
…

that contains only one explicit model representing a state of affairs in which both *p* and *q* propositions are verified. The three dots refer to other possibilities that are kept in an implicit format and in which *p* would be false. Thus, this initial representation supports only affirmative inferences (i.e,. modus ponens [MP], from *p* conclude *q*, and affirmation of consequent [AC], from *q* conclude *p*). To draw denial inferences from a negative minor premise (i.e., either not *p* or not *q*) requires reasoners to flesh out this initial representation with additional models that explicitly represent negated values ($\neg p \cdot \neg q$, and $\neg p \cdot q$).

The mental models theory assumes that this fleshing-out process is demanding and time-consuming. Accordingly, it has been shown that the production of denial inferences that require a fleshing-out process is slower than that of the affirmative inferences supported by the initial model (Barrouillet, Grosset, & Lecas, 2000). In the same way, it has been shown that modus tollens (MT; from not *q* conclude not *p*) is more often endorsed when the minor premise not *q* is presented before rather than after the conditional premise (Girotto, Mazzocco, & Tasso, 1997). Indeed, the preliminary presentation of negative information (i.e., not *q*) leads reasoners to focus on negative values and thus facilitates the explicit representation of the $\neg p \cdot \neg q$ model that supports MT. Furthermore, it has been suggested that the mental models theory provides an account of Wason's selection task (Evans & Handley, 1999; Johnson-Laird & Byrne, 2002), probabilistic reasoning from conditional sentences (Johnson-Laird, Legrenzi, Girotto, Legrenzi, & Caverni, 1999), and the compelling illusory inferences that result from a disjunction of conditional statements (Johnson-Laird & Savary, 1999). Thus, the mental models theory is undoubtedly the most explanatory and heuristic among the available theories of propositional reasoning. However, we would suggest that the standard version of this theory still suffers from two gaps on which this chapter focuses.

First, the theory must account for content effects, which are ubiquitous in conditional reasoning, in a more effective and convincing way than it does. Bonatti (1994a, 1994b) argued that the mental models theory cannot account for content effects in propositional reasoning because this theory relies on a truth-table approach that is formal and ignores both the content and the context. Despite that Bonatti seems to neglect the role of the fleshing-out process in conditional reasoning (see Barrouillet & Lecas, 1998, for a discussion), it should be acknowledged that the standard theory lacks the precise machinery to account for content effects in a predictable way.

Second, the theory must account for the developmental phenomena that have been the focus of past psychological studies on thinking, reasoning, and rationality (Inhelder & Piaget, 1955; Piaget & Inhelder, 1959) because a developmental approach probably constitutes the best way to validate the theory. Indeed, if reasoning is a matter of constructing and manipulating mental models in working memory, these processes should be constrained in three ways. The first of these relates to the limited capacity of the working memory in which mental models are maintained and processed. The second concerns the structure and content of the semantic memory that provides reasoners with knowledge from which mental models are constructed. Finally, the third relates to the relative accessibility of this knowledge in long-term memory. The impact of these constraints on reasoning processes should evolve with age, thus leading to a predictable developmental trend. One of the main developmental changes that could have a direct influence on reasoning skills is the developmental increase in cognitive resources.

It is widely acknowledged that there is an age-related increase in cognitive capacities (Barrouillet & Camos, 2001; Case, 1985; Cowan, 1997; Halford, 1993; Halford, Wilson, & Phillips, 1998; Swanson, 1999). The most recent models of working memory conceive these capacities as a pool of attentional resources available to activate knowledge from long-term memory and to keep it active for processing (Anderson & Lebière, 1998; Cowan, 1995, 2001; Engle, Kane, & Tuholski, 1999; Rosen & Engle, 1997). As a consequence, any increase in cognitive capacities should have an impact both on the number of models that can be processed in working memory and on the accessibility of knowledge from long-term memory (i.e., the two main constraints on reasoning hypothesized by the mental models theory). This theory should thus make it possible to predict the form and content of the representations used at different ages as a function of the amount of available resources and knowledge. Thus, the mental models theory is not only more suited than others to account for content effects, as Johnson-Laird and Byrne (1991) claimed, but it is also developmental in nature.

The Development of Conditional Reasoning: A Model

Markovits and Barrouillet (2002) recently proposed a developmental reformulation of the mental models theory of the conditional based on two main assumptions. First, they assume that although advanced reasoners may develop or learn strategies specific to logical reasoning, children and probably many adults use processes that are general and rely on existing cognitive architectures. Second, they suggest that children (and adults) have an understanding of if-then propositions that is inherently relational and involves the application of a rich linguistic and pragmatic knowledge. More precisely, an "if *p* then *q*" statement is understood as introducing a directional relation between a variable *P*, one value of which is specified by the proposition *p*, and a variable *Q*, one value of which is specified by the proposition *q*. Thus, the "if *p* then *q*" relation defines a semantic space that depends both on the semantic nature of the terms used and on the reasoner's knowledge about the relationship between them. For example, "if he is a postman, then he has a blue cap" would be understood not only as introducing a mapping between different kinds of people and different kinds of hats, but also as referring to hats of professional uniforms because we assume that mental models for conditionals represent not only specific elements but also how they are related (Thompson, 2000; Thompson & Mann, 1995).

In line with Johnson-Laird and Byrne's (1991) theory, it is assumed that children construct an initial representation that contains only one model in which specific tokens represent both *p* and *q* propositions as verified. However, the authors suggest that this model does not represent the mere cooccurrence of *p* and *q* but takes the form of a relational schema in which *p* is understood as a hypothetical state of affairs and *q* as its resulting outcome:

(2) $p \rightarrow q$

The directionality of this relation, already suggested by Evans (1993), accounts for the fact that forward inferences (MP and denial of the antecedent, i.e., DA) are faster than backward inferences (AC and MT) from "if *p* then *q*" forms (Grosset & Barrouillet, 2003).

When a minor premise is given and an inference required, this initial representation could be enriched through a fleshing-out process. Markovits (1993; Markovits & Vachon, 1990) suggested that this fleshing out is the result of an automatic process of the activation and retrieval of knowledge from long-term memory. This process would provide the reasoner with information that makes it possible to construct additional models that represent the values of *Q* that could result from alternative hypotheses on *P* that differ from *p*. In children, at least, these models would then represent specific values of the variables *P* and *Q* rather than negated values using propositional-like tags (represented as ¬ in the standard theory).

The outcome of this fleshing-out process and the resulting representation depend on several factors, including children's ability to maintain complex representations in working memory, the efficiency of the retrieval process, the semantic structure of the concepts the conditional sentence involves, the nature of the relation between the antecedent and the consequent, the amount of available knowledge in long-term memory concerning both these concepts and this relation and finally, the context of enunciation. Though this theoretical framework might seem rather complex, it permits several precise predictions about the way children understand conditional sentences and draw inferences, how this understanding evolves with age, and how contents affect reasoning at different ages.

A first distinction must be made between familiar and unfamiliar relations between the antecedent and the consequent. Indeed, in the first case, retrieval from long-term memory provides the reasoner with knowledge about cases that link possible values of the two variables. For example, on the basis of the conditional premise "if the petrol tank is empty, then the car breaks down," reasoners can retrieve knowledge about the fact that, usually, cars with full tanks run, or that if the spark plugs are dirty cars also fail to start. Such cases constitute ready-made models in which different possible values of *P* (i.e., possible causes of cars breaking down or running) are already linked to their resulting outcome (the car either runs or it does not). On the other hand, when unfamiliar relations are presented, the retrieval process can only provide reasoners with values from the variables *P* and *Q* that must be combined to form models. For example, from an artificial relation such as "if the piece is a square, then it is red," individuals must combine alternative shapes (circle, triangle, etc.) with possible colors. However, there is no available knowledge about this relationship that could direct this construction and help reasoners to keep the constructed models active for processing. Thus, constructing and maintaining models should be easier from familiar rather than unfamiliar relations. As a consequence, the developmental impact of a limitation in cognitive resources depends on the type of conditional relation that is being investigated.

As far as reasoning from artificial relations is concerned, both the construction and maintenance of mental models should be particularly difficult for young children who have limited working memory capacities. As a consequence, the most primitive level in understanding the conditional should involve the construction of only one explicit model of the form $p \cdot q$, the content of which is directly provided by the conditional sentence. This representation leads to a conjunctive-like interpretation of "if *p* then *q*." The developmental increase in cognitive capacities should allow children to construct more complex representations that involve an increasing number of models. Thus, the next step in the development of conditional reasoning should involve the construction of a two-model representation in which a not-*p*–not-*q* model is added to the initial model. Indeed, this additional model maximizes the relevance of the conditional statement (i.e., the amount of information it provides; Sperber & Wilson, 1986). This level corresponds to a biconditional interpretation. Finally, adolescents and adults should be able to construct and process three-model representations that correspond to the complete conditional representation hypothesized by Johnson-Laird and Byrne (1991).

This predicted developmental trend from a conjunctive to a biconditional and then to a conditional interpretation of the conditional sentences involving artificial relations has been observed in many experiments (Barrouillet, 1997; Barrouillet et al., 2000; Barrouillet & Lecas, 1998; Lecas & Barrouillet, 1999). It has also been demonstrated that both developmental and individual differences in understanding conditionals rely on differences in working memory capacity (Barrouillet & Lecas, 1999).

However, as we previously stressed, the way children and adults understand and process conditional sentences does not depend solely on their working memory capacities but also on the semantics of the terms and on the context of enunciation.

Content and Context Effects in Reasoning With Artificial Conditional Relations

We stressed earlier that fleshing out should be the result of a process of retrieval from long-term memory. Thus, even when the conditional sentence that is being considered introduces an artificial relation between the antecedent and the consequent, the nature of the mental models and the resulting interpretation would depend both on the semantics of the terms the conditional statement involves and on its context of enunciation. Indeed, both the semantic and the context determine the possible values that can be combined to construct additional models.

For example, Barrouillet and Lecas (1998) showed that conditional statements that contain binary terms mainly lead to a biconditional interpretation in children. This phenomenon results from the structure of knowledge in long-term memory and can be accounted for by the mental models theory. Suppose reasoners are given a sentence such as "if the light is lit, then the door is open." They should construct an initial model of the following form:

(3) lit → open

When fleshing-out is needed, the retrieval process is both straightforward and highly constrained because this conditional contains binary terms in both the antecedent and the consequent that offer only one alternative value (the light can be either lit or off, and the door can be either open or closed). These alternative values (i.e., "off" and "closed") should then be highly activated and very easy to retrieve. The resulting mental models would be of the following form:

(4) lit → open
off → closed

However, the facility with which this second model is constructed has its own counterpart. Note that this set of models exhausts the possible values on both the antecedent and the consequent, which are linked in a term-by-term correspondence. This type of representation, which Barrouillet and Lecas (1998) referred to as complete, would block any further fleshing-out process because each possible hypothetical value on *P* results in a different outcome, thus maximising the information the conditional sentence conveys. This representation would lead to a biconditional interpretation.

In contrast, when a conditional sentence involves nonbinary terms, several alternative values are possible for both the antecedent and the consequent. For example, a conditional such as "if the piece is a square, then it is red" refers to situations in which the color of a given piece depends on its shape. Additional models should then involve combinations of possible alternative shapes (triangle, circle, etc.) and possible colors (red, blue, white, green, etc.). With these NN conditionals (for nonbinary on both the antecedent and the consequent), the large range of possible values within both variables would make it difficult to determine the precise nature of the possible alternative cases. Thus, the content of the mental models to be added should remain undetermined. The resulting set of models should then take the following form:

(5) square → red
ind shapes → ind colors

in which *ind* refers to indeterminate values. It should be noted that this additional model admits not $p \cdot q$ cases that lead to a conditional interpretation. Thus, according to our theory, the tendency to interpret "if *p* then *q*" sentences as biconditionals should be generally greater with binary than with nonbinary terms.

Now, the number of possible alternative values depends not only on the semantics of the terms the conditional involves but also on the context of enunciation and, more generally, on any factual knowledge that could direct the retrieval process. Barrouillet and Lecas (1998) reasoned that an NN conditional should primarily induce biconditional interpretations provided that it is inserted in a context that permits only two possible values for the antecedent and the consequent. For example, if you were informed that there are only two possible shapes (e.g., a square or a circle), and two possible colors (red or blue), the NN conditional "if the piece is a square, then it is red" would become a binary conditional (Legrenzi, 1970). Thus, if the predominant biconditional interpretation elicited by the binary conditional results from the characteristics of the fleshing-out process, NN conditionals presented in a restricted context would elicit the same type of interpretation.

We tested this hypothesis in a recent experiment in which we asked 12- and 15-year-old children and adults to perform a pencil-paper inference production task (Barrouillet & Lecas, 2002). The participants were presented with a short scenario that introduced a conditional sentence that involved either binary or nonbinary terms. In the latter case, the scenario either restricted to two the possible values on both the antecedent and the consequent or permitted several possible alternatives. This resulted in three experimental conditions defined by the type of conditional presented, either binary (BB conditionals), nonbinary (NN conditionals), or nonbinary restricted (NNR conditionals). For example, the scenario and the conditional sentence in the NN condition took the following form:

> A company's vehicle pool consists of various makes of car: Peugeot, Citroen, Ford, Fiat, Renault, …, present in different colors: red, blue, grey, green, … After looking at all the cars, an observer claims to have found a rule linking the make of car and its color. The rule is "if the car is blue, then it is a Peugeot."

In the NNR condition, the scenario stated that there was only one possible alternative value in both the antecedent and the consequent and took the following form:

> A company's vehicle pool consists of two makes of car: Peugeot and Citroën, present in the colors blue and green. After looking at all the cars, an observer claims to have found a rule linking the make of car and its color. The rule is "if the car is blue, then it is a Peugeot."

In each experimental condition, each conditional was presented with the four canonical forms MP, AC, DA, and MT.

We stressed earlier that the first level of conditional interpretation should be conjunctive in nature and rely on the construction of the initial model only. Thus, younger children should more often produce the affirmative inferences MP and AC on the basis of the initial model than the denial inferences DA and MT that require a fleshing out. Older children should manifest a biconditional interpretation that results in a high production rate of the four canonical inferences. Adults should be able to construct and process three-model representations and thus should manifest a conditional interpretation that corresponds to a high rate of production of the logical inferences MP and MT and a lower production rate of the fallacies AC and DA. We predicted that this developmental trend should be observed with NN conditionals.

However, content and context effects should modulate this developmental trend in a predictable way. Indeed, we hypothesized that conditionals with binary terms should facilitate the construction of a first alternative model of the form not-*p* .not-*q*. As a consequence, younger children should overcome the conjunctive interpretation when presented with binary conditionals and should then reach a biconditional interpretation that results in the increased production of both DA and MT inferences. The same phenomenon should be observed with nonbinary conditionals in restricted contexts that explicitly provide alternative values. On the other hand, content and context effects should have a low impact in those participants who are already able to construct and manipulate two-model representations. Their interpretation should be biconditional in nature, whatever the type of conditional sentence presented. Finally, we expected adults to manifest a conditional interpretation with nonbinary conditionals. However, if binary and nonbinary conditionals in a restricted context impede a complete fleshing-out process, both types of conditionals should induce more biconditional interpretations than nonbinary conditionals, thus resulting in more frequent AC and DA fallacies. In an earlier experiment (Barrouillet & Lecas, 1998), we failed to demonstrate the biconditional interpretation that should result from binary conditionals in adults. Thus, the present experiment used larger samples of participants who were presented with a larger number of conditional syllogisms to solve. Moreover, we had never previously tested the predictions concerning the restricted contexts.

As we predicted, the nonbinary conditionals led to the standard developmental trend described previously. As far as 12-year-old children were concerned, we observed a main conjunctive interpretation, the affirmative inferences MP and AC being produced more often than the denial inferences DA and MT (85% and 47% for affirmative and denial inferences, respectively). The performances in 15-year-old children corresponded to a predominant biconditional interpretation. The production rate of denial inferences was higher than in the 12-year-old children (63% vs. 47% in 15- and 12-year-old children, respectively), whereas the production rates of the affirmative inferences MP and AC remained unchanged. As we predicted, adults manifested the pattern of inferences that should result from a conditional understanding. Indeed, they more often produced the correct inferences MP and MT (98% and 74%, respectively) than the fallacies AC and DA (57% and 44%, see Figure 1.1a and 1.1b).

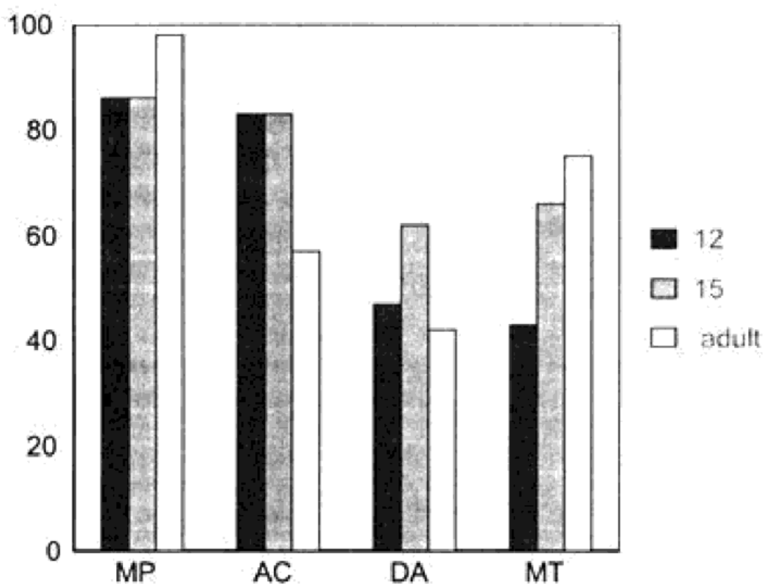

Fig. 1.1.a Production rates of the four inferences (MP: modus ponens; AC: affirmation of the consequent; DA: denial of the antecedent; MT: modus tollens) in the NN condition (conditionals with nonbinary terms) as a function of the age of the

participants.

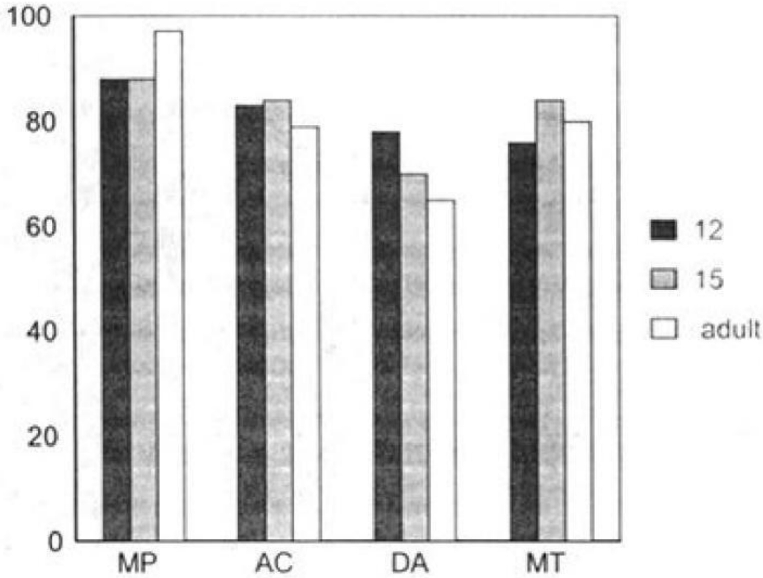

Fig. 1.1.b Production rates of the four inferences (MP: modus ponens; AC: affirmation of the consequent; DA: denial of the antecedent; MT: modus tollens) in the BB condition (conditionals with binary terms) as a function of the age of the participants.

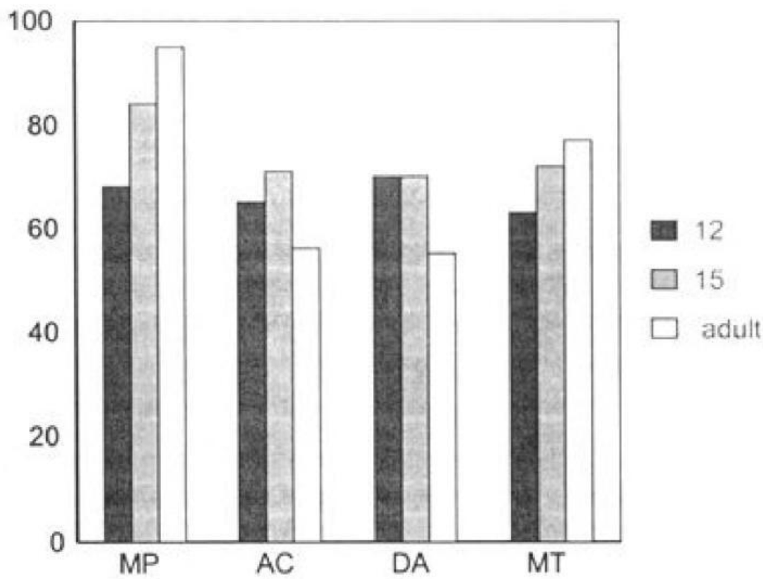

Fig. 1.1.c Production rates of the four inferences (MP: modus ponens; AC: affirmation of the consequent; DA: denial of the antecedent; MT: modus tollens) in the NNR condition (conditionals with nonbinary terms presented in a restricted context) as a function of the age of the participants.

In line with our theory, the effect of age was reduced with binary conditionals that elicited a predominantly biconditional interpretation in all the age groups. The production rate of denial inferences increased from 47% to 77% in the younger participant, whereas adults more often endorsed the fallacies AC and DA (from 51% to 72%). Thus, contrary to what Barrouillet and Lecas (1998) observed, binary terms induced biconditional interpretations in adults as well as in children. However, it should be noted that even with binary conditionals adults obtained slightly better performances than children. They more often produced the correct inferences MP and MT and less often the fallacies AC and DA (see Fig. 1.1).

Nonbinary conditionals presented in restricted contexts led to a more complex pattern of results. As we predicted, this condition elicited a higher production rate of denial inferences in the younger participants and also resulted in lower production rates of the affirmative inferences MP and AC. This could be because the restricted context was introduced by means of a scenario that explicitly presented the possible alternative values for both the antecedent and the consequent. Thus, these conditionals could have elicited initial representations that contained not one but two explicit models right from the start. Of course, these representations not only facilitate the production of denial inferences but also impede the production of affirmative inferences because the initial two-model representation from which these inferences are drawn brings about a higher working memory load. Thus, all of the four inferences were produced at roughly similar rates by 12-year-old children (between 63% and 70%). As we anticipated, the biconditional pattern of results was also predominant in 15-year-old children. However, adults did not exhibit the predicted biconditional pattern. In fact, the restricted contexts had only a small effect on adults. Many of them continued to exhibit a main conditional interpretation and produced more logically correct than fallacious inferences (86% vs. 55%). In summary, conditionals involving binary terms had the same effect in adults as in young children, and participants in both groups mainly interpreted these sentences in a biconditional way. In young children, this tendency resulted in an increased production of the denial inferences MT and DA, whereas in adults it led to an increased production of the fallacies AC and DA. However, and contrary to our hypothesis, restricting the context in which nonbinary conditionals appeared did not have the same effect, and the adults' interpretation remained practically unchanged.

Despite this latter fact, our results lend strong support to the mental models approach of conditional reasoning. The mental models theory assumes that the main constraint on human reasoning is the number of models to be constructed and processed because mental models are held in a limited-capacity working memory (Johnson-Laird & Byrne, 1991). Because children have lower working memory capacity than adults, and because this capacity increases with age, they should be limited in the number of the models they can construct and manipulate, and this number of models should progressively increase (Barrouillet & Lecas, 1999). Thus, the interpretation of conditional sentences evolves from a conjunctive to a biconditional and then to a conditional interpretation. The content effect produced by binary terms reinforces the hypothesis that reasoning is a matter of constructing mental models that link together tokens retrieved from long-term memory (Markovits & Barrouillet, 2002). When the retrieval process directly provides individuals with specific and clearly determined alternative values, mental models are easy to construct and young individuals reach a higher interpretational level (i.e., biconditional). However, the retrieval of only one alternative tends to block the fleshing-out process prematurely, and adults tend to regress to a biconditional interpretation. Note that this effect is akin to Cummins' (1995) observations that causal conditional relations that permit few, if any, alternative causes of the effect presented in the consequent mainly result in biconditional patterns of responses.

That restricted contexts did not have the same effect as binary terms in adults suggests that the effects related to the structure of semantic memory supersede the impact of contextual information. It is possible that the fleshing-out process is mainly directed by semantics and that contextual information only modifies the initial model, as suggested by the performance of the 12-year-old children on affirmative inferences. For example, the automatic retrieval process could elicit, in adults, a fleshed-out representation of NNR conditionals as shown in Diagram 5 in which the initial model already contains two explicit models. Such a process would result in the following set of models

(6)	p	→	q
	p'	→	q'
	ind *P*	→	ind *Q*

in which *p*' and *q*' correspond to the alternative values given by the restricted context. Some results have confirmed this hypothesis. For example, adults produced more denial inferences, DA and MT, with NNR conditionals (57% and 82%, respectively) than with NN conditionals (37% and 72%) when these inferences were cued by minor premises that contained implicit negations (e.g., with the problem of the color of the cars, the DA minor premise was "the car is red" rather than "the car is not blue," and the minor premise of MT was "the car is a Citroën" rather than "the car is not a Peugeot"). Indeed, in the set of models presented in Diagram 6, there is an explicit model that directly matches the implicitly negated minor premise to be considered (i.e., either *p*' or *q*').

Thus, context effects and, to a lesser extent, content effects were more pronounced in children than in adults. These developmental differences could be due to the fact that children have less cognitive resources than adults for the retrieval of knowledge from memory and the processing of complex representations. As a consequence, any semantic or contextual factor that facilitates the retrieval of specific values from long-term memory and the construction of models has a stronger impact in young than in older individuals. The same phenomenon is also observed in reasoning from familiar conditional relations.

Content Effects in Reasoning With Familiar Causal Relations

As we have stressed earlier, the "if *p* then *q*" relation defines a semantic space that is determined by the reasoner's understanding of both the nature of the terms used in the conditional and the relationship between them. We have seen that when unfamiliar relations are used, the nature of the terms and the structure of the semantic space in which they are embedded determines the way individuals understand conditional sentences. When familiar relations are considered, the process of retrieval from memory does not provide the reasoner with alternative values for the antecedent, on the one hand, and the consequent on the other, but with knowledge about cases that link possible values of these two variables. This knowledge underpins the construction of models to be added to the initial representation.

Thus, three classes of cases could be activated that can lead to the construction of three types of models. The first class concerns cases in which both the relationship and the actual objects or events concerned are complementary to those specified in the original conditional, that is, cases where objects or events that are different from *p* are related to not *q* (we refer to this as the complementary class). For example, a reasoner who is given a premise such as "if it rains, then the street will be wet" would activate related events such as "if it is sunny, then the street will not be wet" or "if it is only cloudy, then the street will not be wet." The retrieval of one case from this class leads to the construction of a not *p* · not *q* model that supports the production of the denial inferences DA and MT. The second class concerns possible objects/events that share the same relation to *q* as *p* does, that is, cases in which not-*p* implies *q* (we refer to this as the alternatives class). For example, in "if the street cleaner passes, then the street will be wet," note that the retrieval of such an alternative case results in a not *p* · *q* model that would block the production of the fallacies AC and DA. Indeed, "it rains" no longer follows from the minor premise "the street is wet" for AC, while "the street is dry" cannot be concluded with certainty from the DA premise "it is not raining." Finally, the third class concerns what Cummins (1995) called disabling conditions, that is, conditions which allow the relationship between *p* and *q* to be violated (we refer to this as the disabling class). An example is "if it rains, but the street is covered, then the street will not be wet." The retrieval of disabling conditions should impede the logically correct inferences MP and MT through the construction of a $p \cdot$ not q model. Markovits and Barrouillet (2002) suggested that, at least in children, these three classes should differ in their accessibility: cases from the complementary class should be more accessible than those from either the alternatives class or disabling conditions.

As we stressed earlier, the nature of the conditional relation, either familiar or unfamiliar, has a direct impact on reasoning processes, and models are easier to construct and maintain with familiar relations. Indeed, even young children seem able to construct two-model representations (i.e., $p \cdot q$, and not $p \cdot$ not q) that lead to a biconditional interpretation, as many studies have demonstrated. However, the content of the conditional sentence still has a major impact on reasoning because it directs the retrieval process.

For example, Markovits, Fleury, Quinn, and Venet (1998) demonstrated that inference production from familiar-class based relations depends on the strength of the association the conditional sentence involves. More precisely, they predicted that correct responses of uncertainty to AC and DA should be more frequent from conditional sentences that involve a weak (e.g., "if something is a butterfly, then it has legs"), rather than a strong relation between the antecedent and the consequent (e.g., "if something is a dog, then it has legs"). The authors reasoned that, within the class of animals that have legs, some cases should have a higher base-level activation than others because they are more typical regarding the legs feature and should thus be easier to retrieve (Anderson, 1993 ; Anderson & Lebière, 1998). It should be recalled that the retrieval of alternative cases of the form not $p \cdot q$ leads to a correct uncertainty response on the two fallacies AC and DA. Thus, when the conditional premise involves a weak relation such as "if something is a butterfly, then it has legs," those items that are strongly associated with the relevant feature are highly activated and thus easy to retrieve (e.g., dogs, or cows). Now, these items constitute alternative cases that ensure the correct response of uncertainty to AC and DA. On the other hand, when the conditional involves a strong relation (e.g., "if something is a dog, then it has legs"), alternative cases must be recruited from those items that are less strongly activated and thus more difficult to retrieve. Of course, as far as class-based relations are concerned, this difference should have an effect only on individuals who have low retrieval capacities. Thus, the authors hypothesized that young children should more often endorse the fallacies AC and DA from conditionals involving strong rather than weak relations and that this effect should disappear with development because older children have improved retrieval capacities that allow them to evoke alternative cases whatever the strength of the relation presented. The results confirmed these hypotheses. These facts have been recently extended to causal relations in adults: Strong causal relations lead to more frequent fallacies than weak relations (Quinn & Markovits, 1998).

We recently ran a developmental version of this latter experiment in our lab (Barrouillet, Markovits, & Quinn, 2001). We asked 12- and 15-year-old children and adults to perform a conditional syllogism evaluation task on the basis of either strong or weak causal conditional relations. In the former, an event presented as a consequence was linked to its more frequent and probable cause (e.g., "if a dog has fleas, then it will scratch constantly") whereas in the latter the same consequence was associated with a less frequent cause (e.g., "if a dog has a skin disease, then it will scratch constantly"). Following Quinn and Markovits (1998), we hypothesized that alternative causes of the consequent (scratching) should be easier to evoke when the conditional premise involves a weak rather than a strong relation. Indeed, strong relations have a higher base-level activation than weak relations. Thus, the former are easier to retrieve and constitute alternative cases when weak relations are involved in the conditional premise. Thus, those participants who studied weak relations should be less likely to endorse the AC and DA fallacies than those who studied strong relations. Moreover, we hypothesized that this effect should be stronger the younger the participants are because young participants have lower retrieval capacities.

In this study, the participants were presented with four conditional premises that contained either strong or weak relations between antecedent and consequent (e.g., "suppose that if a dog gets fleas, then he will scratch continuously"), along with minor premises corresponding to the four canonical forms MP, AC, DA, and MT (e.g., "a dog does not have fleas" for DA). For each form, the participants were asked to choose between two conclusions, either of certainty or uncertainty (e.g., "it is certain that it will not scratch constantly" or "one cannot be certain whether it will scratch constantly or not" for DA). Fig. 1.2 displays the endorsement rates (i.e., responses of certainty) of the four canonical forms MP, MT, AC and DA.

As we predicted, the strength of the relation affected the production rates of fallacies. Overall, weak relations significantly elicited lower endorsement rates of both DA and AC than strong relations (38% for weak, 54% for strong relations; Fig. 1.2.a-c). Note that this effect did not affect the evaluation of logical inferences. Indeed, logical inferences do not depend on the retrieval of alternative causes (not $p \cdot q$ models) but on either the initial model for MP or the complementary model not p · not q for MT. As far as 12-year-old children were concerned, the strength effect was only significant for AC but not for DA, whereas this effect was significant for both inferences in 15-year-old children. These effects were no longer significant in adults, contrary to the observations reported by Quinn and Markovits (1998). Though the strength effect did not reach significance for DA in the younger group, this developmental trend was, in fact, in line with our theory.

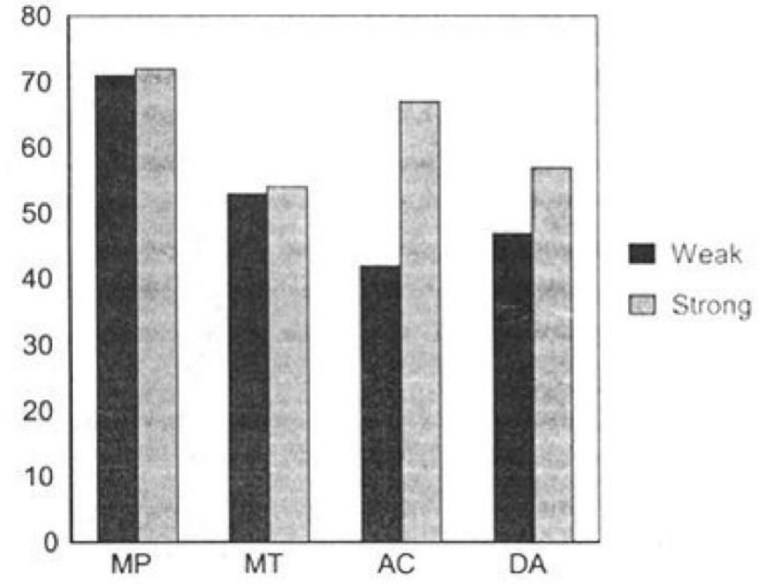

Fig. 1.2.a Production rates of the four inferences (MP: modus ponens; MT: modus tollens; AC: affirmation of the consequent; DA: denial of the antecedent) as a function of the strength of the causal relation involved in the conditional premises used for the 12-year-old children.

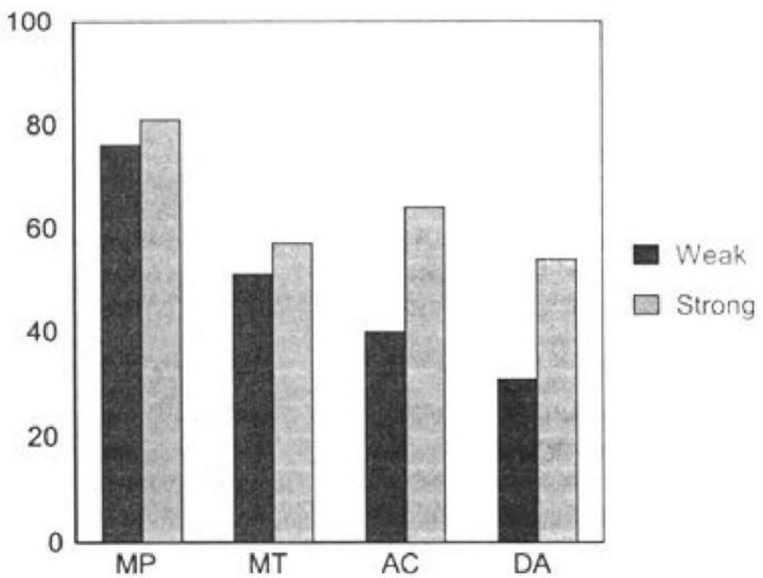

Fig. 1.2.b Production rates of the four inferences (MP: modus ponens; MT: modus

tollens; AC: affirmation of the consequent; DA: denial of the antecedent) as a function of the strength of the causal relation involved in the conditional premises used for the 15-year-old children.

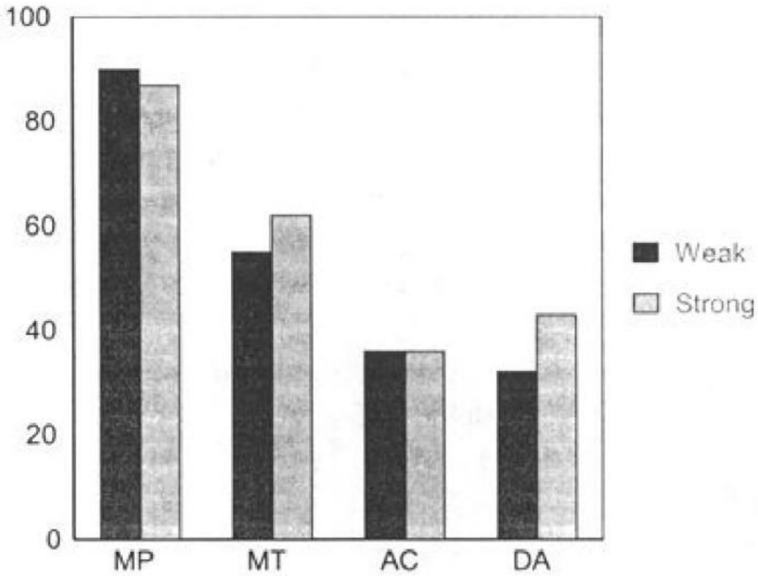

Fig. 1.2.c Production rates of the four inferences (MP: modus ponens; MT: modus tollens; AC: affirmation of the consequent; DA: denial of the antecedent) as a function of the strength of the causal relation involved in the conditional premises used for the adults.

We assumed that the correct response of uncertainty to both AC and DA results from the retrieval of a case from the alternatives class and the construction of a model of the form not $p \cdot q$. This retrieval would be triggered by the content of both the conditional and the minor premises. The content of the conditional premise, the strength of the association it involves, determines the accessibility of the alternative class. However, the minor premises for AC on the one hand and DA on the other differ in their efficiency in activating this relevant knowledge (i.e., the alternative causes). Indeed, the AC minor premise refers to the consequence of the causal relation (e.g., "a dog scratches constantly"), whereas the DA minor premise refers to the absence of the cause involved in the conditional (i.e., "a dog does not have fleas" and "a dog does not have any skin disease" for the strong and weak relations, respectively). Thus, the AC minor premise directly matches alternative cases by evoking the consequence, whereas the DA minor premise should primarily activate cases from the complementary class because it refers to the absence of the cause involved in the conditional. As a consequence, the activation and retrieval of an alternative cause is more probable from AC than from DA. Thus, those participants who have low capacities should be confronted with two difficulties. First, they could access the alternatives class only when the conditional involves a weak association, and second, the efficiency of the retrieval process should then depend on the nature of the retrieval cue provided by the minor premise. As a consequence, the strength effect should be more pronounced with AC than with DA, and this is exactly what we observed in our youngest group. Note that the same phenomenon was also observed by Janveau-Brennan and Markovits (1999). The developmental increase in capacities progressively removes these two constraints. Indeed, the strength effect was observed in older children whatever the type of minor premise being investigated, and the performance of adults remained unaffected either by the strength of the relation or the type of minor premise.

Note that this strength effect differs from the phenomena observed by Cummins (1995) and Janveau-Brennan and Markovits (1999), who observed that reasoning from causal conditional relations depends on the number of available alternative causes that can produce the same effect. More precisely, the rate of biconditional responses (i.e., the endorsement of the fallacies AC and DA) is higher from causal conditional relations that allow few rather than many alternatives. Our model easily accounts for this fact because the probability of retrieving a case from the alternatives class that ensures a conditional response to AC and DA is higher when the causal relation under study allows many rather than few alternative causes. In the present experiment, the number of alternative causes was exactly the same for both the strong and the weak causal relations because they referred to the same resulting outcome (e.g., "a dog scratches constantly"). In fact, the effect of the number of alternative causes Cummins (1995) discovered is similar to the effect produced by binary terms. Both rely on the structure of semantic memory. Here, the strength effect results from the relative accessibility of different items of knowledge pertaining to a given semantic structure. Thus, as we stressed earlier, reasoning depends on the semantic structure of the concepts the conditional sentence involves, the nature of the relation between the antecedent and the consequent, the amount of available knowledge in long-term memory concerning both these concepts and this relation, and the relative accessibility of this knowledge

Mental Models and the Development of Conditional Reasoning

At the beginning of this chapter, we stressed the need for a mental models theory of reasoning that could account for developmental and content effects. As we suggested, the mental models theory is particularly suitable for accounting for developmental phenomena. Indeed, if reasoning is a matter of constructing and manipulating mental models in a limited-capacity working memory (Johnson-Laird & Byrne, 1991), then the complexity of the representations individuals can construct and process should evolve with development. The first available representation should correspond to the initial model postulated by Johnson-Laird and Byrne (1991), and this representation should increase in complexity as cognitive resources evolve with age.

Accordingly, we observed that the interpretation of conditional sentences and the resulting patterns of inferences evolve with age from a conjunctive to a biconditional and then to a conditional interpretation, which is underpinned by the construction of one, two, and then three models, respectively. This fact lends strong support to the mental models framework. Indeed, it has been claimed that adults outperform children and adolescents in laboratory reasoning tasks because they are more able to set aside the conversational principles that lead to a biconditional interpretation (Braine & O'Brien, 1991). In the same way, O'Brien, Dias, and Roazzi (1998) argued that the mental models theory cannot account for development because it predicts conjunctive responses that would never occur, whereas mental logic predicts an initial biconditional developmental level that has often been observed. The results of the first study presented demonstrated that there is an initial developmental level which is conjunctive in nature. An individual response pattern analysis conducted on the production rates of the four classical inferences with nonbinary conditionals indicated that a majority of 12-year-old children manifested a coherent conjunctive pattern of responses (see Fig. 1.3). This predicted conjunctive level has been systematically observed in many previous studies in children (Barrouillet, 1997; Barrouillet et al., 2000; Barrouillet & Lecas, 1998, 1999;; see also Paris, 1973, and Taplin, Staudenmayer, & Taddonio, 1974, for related observations), and even in adults when the fleshing-out process fails (Barrouillet et al., 2000; Girotto et al., 1997). In line with the mental models theory, the predominant interpretation in 15-year-old children was biconditional in nature, whereas many adults manifested a conditional pattern of responses.

regressed to a biconditional interpretation when the conditional premise contained binary terms (73% of them exhibited a coherent biconditional response pattern) and any developmental evolution tended to disappear (the corresponding rates of participants who manifested such a pattern were 70% and 77% in 12- and 15-year-old children, respectively; see Fig. 1.1). Thus, when the structure of semantic knowledge offers clear and readily available alternative values, the fleshing-out process is facilitated, and young participants can go beyond the conjunctive interpretation. However, the resulting representation tends to block any further fleshing out, and most of the adults continue to adhere to a biconditional interpretation.

When the conditional premises contain familiar relations, the structure of the knowledge has more subtle effects that rely on the relative accessibility of different items. Conditionals that involve strong relations between the antecedent and the consequent elicit more fallacies than conditionals that contain weak relations, at least in participants who have limited capacities to activate and retrieve knowledge. When the major premise contains a weak relation, knowledge about strong relations, which is easily retrieved from memory, constitutes alternative not $p \cdot q$ cases that block the classical fallacies. Note that this effect is quite counterintuitive. Indeed, strong relations are probably more familiar than weak relations, and it has been shown that familiarity often improves performance (Markovits, 1984).

Markovits and Barrouillet's (2002) mental models theory accounts for both developmental and content effects. As these authors suggested, there is no need to suppose that reasoning is underpinned by specific cognitive processes or strategies, such as rules or schemas. Development would result from an age-related increase in working memory capacity that allows children to retain and manipulate an increasing number of mental models. Content should modulate this developmental trend through the relative accessibility of the knowledge that provides the building blocks for the construction of models. Thus, both developmental and content effects result from general cognitive constraints that affect the processes by which mental models are constructed and manipulated in working memory. Contrary to what O'Brien et al., (1998) and Bonatti (1994a, 1994b) claimed, the mental models approach has the precise machinery to account for and predict both developmental and content effects.

To decide between competing theories is not the sole object of a developmental approach. Indeed, the developmental evolution we observed sheds light not only on our understanding of children's and adults' reasoning processes but also on the nature of the final state toward which this evolution tends. We suggest that this final state would constitute the norm that human reasoning should conform to.

The Meaning of If Then and the Developmental Process

The real psychological meaning of the if-then connective has long been a subject of controversy. It has been suggested that the way individuals understand if-then does not correspond to the material implication of the formal logic and that logic-based theories of human reasoning, such as the mental models theory, cannot account for human performances (Oaksford & Chater, 1998). More generally, it has even been suggested that formal logic does not constitute, at least in certain cases, the appropriate normative theory for the assessment of human behavior and reasoning. It has been argued that human reasoning would often conform to an adaptive rationality (a rationality 1) that differs from logical rationality (rationality 2, Evans & Over, 1996). However, it should be noted that many of these assumptions are based on adult studies.

We suggest that focusing exclusively on adult reasoning performances might not be the more appropriate way to investigate these difficult problems. Indeed, adult studies involve two pitfalls that can, and actually do, obscure the debate. First, human performances vary from one task to another and, within a given task, from one content to another (Thompson, 2000). For example, there is no doubt that the poor human performance in Wason's selection task has played a critical role in the debate concerning human rationality by suggesting that human reasoning does not conform to logic. However, though the selection task provides psychologists with striking and puzzling data that need to be accounted for, it must nevertheless be borne in mind that it constitutes only one task among others, and many authors have claimed and provided evidence that it does not require conditional reasoning (Markovits & Savary, 1992). More direct evaluations of the human ability to reason from conditional sentences, such as inference evaluation or production tasks, provide results that conform more closely to the logical norm (see Evans, Newstead, & Byrnes, 1993, for a review). Within this latter kind of task, our results demonstrate that striking differences in performance can result from apparently small content variations (e.g., whether a conditional contains binary terms or not). Such variability makes it difficult to determine the meaning of if-then in adults, even when we consider only indicative or causal conditionals. Second, it seems clear that human reasoning, as well as other high-level cognitive processes, is restricted by limited cognitive capacities. Thus, performance cannot be equated with competence. Young adult studies only provide us with a snapshot from which it is difficult to gain an insight into the real distance between performance and competence.

We claim that a better way to solve the problem of the meaning of if-then would be to adopt the respected approach put forward in Piagetian genetic psychology. As pointed out by Karmiloff-Smith (1992), a developmental perspective is essential to the analysis of human cognition, and the way in which

Thompson, V. A., & Mann, J. M. (1995). Perceived necessity explains the dissociation between logic and meaning: The case of «only if». *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *6*, 1554–1567.

argue that the procedures for manipulating models as currently proposed are too deterministic. We suggest that people are flexible in the strategies and approaches that they employ in reasoning tasks. In particular we show that certain inferences are solved using suppositional strategies that bear a close relationship to the suppositional strategies that have been proposed by rule based theorists.

The chapter begins with a brief description of the current model theoretic account of conditionals. We then describe data from some of our own recent work on the role played by background knowledge about temporal relationships in the suppression effect (Byrne, 1989; Byrne, Espino & Santamaria, 1999; Rumain, Connell and Braine, 1983). We suggest that this data illustrates the need for a consideration of the multiple types of information that people represent during reasoning and a closer integration of model based accounts of text comprehension with those of reasoning. Next, we describe other work examining the way in which people interpret and make inferences from everyday even if conditionals. We use this data to illustrate the importance of specifying clearly how background knowledge is combined with information in an assertion to determine people's understanding of a given proposition and the inferences supported by that proposition. In the penultimate section of the chapter, we describe some developmental data on conditional reasoning, which shows that the means by which participants draw the modus tollens (MT) inference changes over time. Evidence is presented that strongly suggests that older participants employ a flexible strategy in MT reasoning that involves the construction and active manipulation of a single model and the generation, coordination, and manipulation of suppositions. In the final section, the general implications of our work are discussed in the context of the future development of model-based theories of deduction.

The Model Theory of Reasoning

According to the model theory of conditionals (Johnson-Laird & Byrne, 1991; Johnson-Laird et al., 1992) reasoning proceeds by the construction of models corresponding to the possible states of affairs that the premises describe. In deduction, reasoners normally construct models of the premises of an argument that correspond to a conception of the way the world would be if the premises were true and represent as little information as possible explicitly to reduce demands on working memory. According to this account, a conditional premise such as the following:

(1) if Jimmy goes fishing, then he will have a fish supper

is represented initially by a single model:

he had a fish supper. In contrast, if reasoners are told that Jimmy did not have a fish supper, they are unable to combine this information with the initial model of the conditional.

In this case, the full set of models must be fleshed-out (Diagram 3); those models in which Jimmy had a fish supper eliminated leaving a single model in which Jimmy did not go fishing. According to the model theory, the fleshing-out of models is prone to error and hence reasoners are less likely to draw MT, which involves consideration of multiple models, than to draw MP which can be drawn from a single model. In a later section, we will return to the case of MT and discuss evidence that suggests an alternative account of the processes that underly MT reasoning.

The explanation for the difference in endorsement rates between AC and DA draws on a similar distinction. According to the acount, the fallacies are endorsed because reasoners adopt a biconditional representation of the rule, representing only two possible states of affairs in the fully explicit model set:

(4)	*F*	*FS*
¬*F*	¬*FS*	

AC is endorsed more frequently than DA because the initial representation of the biconditional consists of the first model discussed previously and hence supports the inference, whereas DA requires the initial representation to be fleshed-out.

Recently Johnson-Laird and Byrne (2002) developed their account of conditionals to include five principles that are drawn on in explaining the ways in which people interpret and reason from a range of conditional forms. The major extension to the theory relates to the introduction of an account of the way in which pragmatics and context affect the representation that is consructed. According to Johnson-Laird and Byrne there are two core meanings of conditionals, the conditional interpretation that consists of the three explicit models presented in Diagram 3 and the tautological interpretation. The tautological meaning is associated with conditionals that specify a modal relationship between antecedent and consequent, such as “if *A* then possibly *B*”. The core tautological meaning is captured in a model set that represents all four truth table cases:

(5)	*A*	*B*
A	¬*B*	
¬*A*	*B*	
¬*A*	¬*B*	

According to the revised account, the basic conditional and tautological meaning can be moderated by semantics or pragmatics. To illustrate the operation of pragmatic modulation, consider the following assertion (Jonhnson-Laird & Byrne, 2002): “Even if the workers settle for lower wages, the company image not available

The argument that mental models are informationally rich representations is not unique to the reasoning field. The situation model approach to text comprehension has long emphasized the multidimensional nature of model-based representations. According to this approach, in understanding any piece of discourse, people not only construct a representation of the words in the sentences but also construct a model corresponding to the situation described by those sentences. Hence, they construct a model corresponding to what the text is about rather than a representation of the text itself (see Zwaan & Radvansky, 1998, for a comprehensive review). Situation models constitute a level of representation associated with a deep understanding of a piece of text, and they allow the integration of information in a text with information from background knowledge. This level of representation is multidimensional in nature, encoding a variety of aspects of the situation described in the text. This might include spatial information (Bransford, Barclay, & Franks. 1972; Zwaan & Radvansky, 1998), information about the causal relationships between events (Singer, Halldorsen, Lear, & Andrusiak, 1992), or the goals and intentions of the protagonist (Graesser, 1981; Schank & Abelson, 1977). The situation model account has been successfully applied equally to single sentences (Rinck, Hahnel, & Becker, 2001) and to full story-based texts (Zwaan, Maglioni, & Graesser, 1995). The important point for our purposes is to highlight that there is good evidence that models encode information over and above the possible situations that a premise may describe. Each possibility may also be augmented by additional information based on our knowledge of events in the world. In the next section we illustrate not only that models for reasoning are much richer than has been proposed but also that the information people encode into their representation has a dramatic effect on the inferences they draw.

Models, Temporal Order and the Suppression of Inference

In addition to information about causality, intentionality, and space, as discussed in the previous section, situation models encode information about both the relative and absolute times at which the events described in the text or discourse occurred. Despite the ubiquity of temporal information in language (see Miller & Johnson-Laird, 1976), time has received surprisingly little attention in the literature on situation models (for a review see, Zwaan & Radvansky, 1998). The work that does exist (see, e.g., Rinck, Hahnel, & Becker, 2001) suggests that people spontaneously encode temporal information in their mental models and that they have strong expectations about the relationship between the order in which events occur in the world and the order in which they appear in text or discourse.

We believe that temporal information will be encoded into the representations that people use for reasoning and that the temporal aspect of